# A Zero Phase Shift Band Pass Filter

Hens Steehouwer[†,*] Kai Ming Lee[†]

October 11, 2011

**Abstract**

Many economic time series can be decomposed into transitory phenomena, cyclical movements and permanent shifts. Band pass filters can be used to decompose such series without necessitating strong model assumptions on the time series dynamics. Baxter and King (BK) and Christiano and Fitzgerald (CF) are time domain filters which are derived by approximating the frequency domain properties of ideal band pass filters, and can be used for frequency decompositions. However, both filters have some known drawbacks. The BK filter can remove trends at the zero frequency effectively, but does not adequately suppress other undesired low frequency fluctuations. Furthermore, improving the approximation of the frequency response of ideal filters requires sacrificing more observations. The CF filter improves the frequency response approximation, but has deteriorating performance near the end points of the data, and induces phase shift distortions.

In this paper we present an approximate band pass filter which overcomes the drawbacks of the previous filters. The filter works by recursively fitting multiple trigonometric functions in the time domain to observed series, and filtering the remaining residual directly in the frequency domain. By only retaining the frequencies in the desired pass band, a band pass filter with desired properties is defined. In particular, the filter does not induce phase shifts and does not require down weighting or removal of any data.

We provide two empirical applications of the filter. In the first application we filter macroeconomic series with business cycle pass bands and make a comparative analysis with the filter output of the BK and CF filters. In the second application we consider stock market returns at different horizons or frequencies. We show that when returns with different horizons or frequencies are of simultaneous interest (i.e., when returns of different frequencies need to be considered), optimal boundaries of the pass bands can be derived in order to approximate these mixed frequency returns by filtering the underlying (log) index with the presented filter.

*Keywords*: time series; frequency domain decomposition; band pass filter; business cycle

[*]Affiliated with the Econometric Institute, Erasmus University Rotterdam, PO Box 1738, 3000 DR Rotterdam, The Netherlands

[†]Ortec Finance Research Center, PO Box 4074, 3006 AB, Rotterdam, The Netherlands; email corresponding author: Kai.MingLee@ortec-finance.com

# 1 Introduction

In economic analysis several different techniques are commonly used to decompose time series into components which can be analysed separately. Often policy or decision makers are focussed on mainly one aspect of the time series dynamics. Growth policy is mainly concerned with long term trends, while stabilisation policy usually focusses on the more rapidly fluctuating business cycles. Seasonal and other high frequency fluctuations are usually of little interest in themselves, and are often removed prior to or during analysis. Moving average filters, (seasonal) differencing, deterministic trend fitting, unobserved components modelling and Hodrick-Prescott filters can all be considered examples of decomposition tools.

Spectral analysis provides a firm foundation for time series decomposition based explicitly on frequency characteristics. A key tool in frequency domain analysis is the band pass filter, which suppresses fluctuations outside a specified range of frequencies called the pass band, while leaving those inside the range intact. Using such a filter allows the researcher to precisely isolate the frequencies of interest. Furthermore, using non-overlapping pass bands provides an orthogonal decomposition. Since the different components are uncorrelated by construction, this simplifies using separate models for each component.

An ideal band pass filter has the properties that it removes *all* fluctuations outside the pass band, leaves fluctuations inside the pass band unaltered, and does not affect the series in any other way. Such a filter is unfortunately only realisable in a theoretical setting with unlimited observations. Two well know algorithms, the Baxter-King and the Christiano-Fitzgerald filters attempt to approximate the ideal filter in the time domain with a limited number of observations. We show that both these approximations have some undesirable properties, and propose an alternative method, based on a combination of iteratively fitting trigonometric series in the time domain, and directly filtering the residual in the frequency domain. This procedure circumvents the problems with the time domain approximations, as well as the finite sample artefacts of a direct frequency domain filter.

To demonstrate possible applications of the proposed filter, as well as band pass filtering in general, we provide two empirical illustrations. First, we filter economic production series of the Netherlands and Finland with various lengths, sampling frequencies and pass bands, using our proposed filter and with the approximating time domain filters. We show that our filter gives similar output as the Christiano-Fitzgeral filter, except near the end points of the sample. Second, we use stock market returns over different horizons as an aide to choose appropriate pass band boundaries. Calculating returns over different horizons can be considered ad-hoc filters to infer long, medium and short term dynamics from the time series. However, these filters do not provide a proper decomposition, since returns over, for instance long horizons also contain medium and short movements. We show that given a set of different return horizons, we can choose pass bands such that the maximum variance of the returns is captured using a proper decomposition, in which the total return can be modelled as the sum of the component returns.

Our proposed filter is very largely based on the ideas from Bloomfield (1976) and the filtering approach described in Schmidt (1984). Our main contributions lie in a detailed and refined development of the filter algorithm, a comparison of the filter to the better known BK and CF filters and an application to mixed frequency return modelling. The proposed filter is a central element of the frequency domain methodology described in Steehouwer (2010).

# 2 Approximating band pass filters

In this section, we briefly review some standard frequency domain filter theory, along with descriptions of the BK and CF filters. The section mainly serves to establish notation and collect some basic results. For more details and derivations, we refer to Koopmans (1974), Bloomfield (1976) or Priestley (1976).

## 2.1 The ideal band pass filter

A linear filter $G(L)$ is a linear transformation of a time series $x_t$ characterised by weights $g_l$ at lag $l$:

$$G(L) = \sum_{l=a}^{b} g_l L^l, \qquad a \leq 0 \leq b, \tag{1}$$

where $L$ is the lag operator $L^k = x_{t-k}$. The filter is applied to a time series $x_t$ to produce the filtered series $y_t$:

$$y_t = G(L)x_t = \sum_{l=a}^{b} g_l x_t. \tag{2}$$

The effect of applying the filter is summarised by the frequency response function (FRF) of the filter, given by

$$G(e^{-i\omega}) = \sum_{l=a}^{b} g_l e^{-i\omega l}. \tag{3}$$

A linear filter induces a multiplication of the amplitudes of $x_t$ of

$$Gain(\omega) = |G(e^{-i\omega})| \tag{4}$$

and a shift of its position in time of

$$Phase(\omega) = \arg(G(e^{-i\omega}))/(2\pi) \tag{5}$$

at frequency $\omega$. These effects are known respectively as the gain and the phase shift of the filter. A linear filter induces no phase shifts if and only if the weights are symmetrical, i.e., $g_l = g_{-l}$ for $l > 0$.

An *ideal* band pass filter with pass band $[\omega_1, \omega_2]$ is a filter which leaves all fluctuations with frequencies within the pass band unaltered, removes all fluctuations with frequencies outside the pass band, and induces no phase shifts. The the zero phase shift requirement is especially important for analysis of multiple time series, since phase shift inducing filters can produce distorted or spurious lead-lag relationships between variables. These requirements imply that he gain function of the ideal band pass filter is the perfectly rectangular function

$$Gain(\omega) = \begin{cases} 1 & \text{for } \omega_1 \leq \omega \leq \omega_2 \\ 0 & \text{for } \omega < \omega_1 \text{ or } \omega > \omega_2 \end{cases}, \tag{6}$$

and the phase shift function is a constant zero.

The weights of the ideal band pass filter can be derived as

$$g_l = \begin{cases} \frac{\sin(\omega_2 l) - \sin(\omega_1 l)}{\pi l} & \text{for } l \neq 0 \\ \frac{\omega_2 - \omega_1}{\pi} & \text{for } l = 0 \end{cases}, \tag{7}$$

see, e.g. Steehouwer (2005, Ch. 5.3.1). Unfortunately, the ideal band pass filter cannot be applied to practical datasets. The weights $g_l$ extent in the infinite future and past, and differ substantially from zero at long lags. Since a proper application of the ideal filter would require infinite observations, some form approximation is needed.

## 2.2 Baxter-King and Christiano-Fitzgerald filters

In the Baxter and King (1999) approach, the ideal filter weights (7) are used up to $a$ lags, after which the weights are truncated. Truncation of weights at higher lags can be shown to be the optimal procedure when a squared distance metric between the ideal and the (finite sample) approximating FRF is adopted as a loss function. Furthermore, a restriction is added that the FRF at frequency

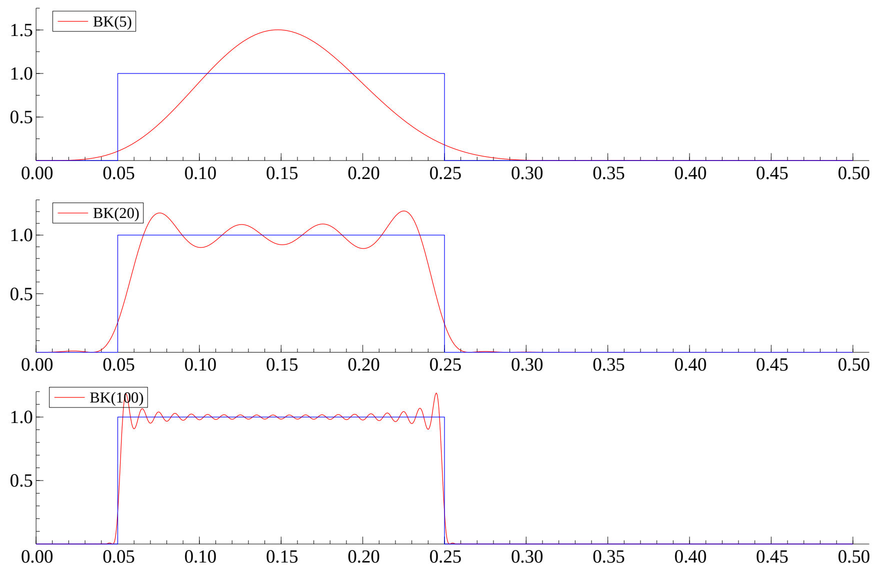


Figure 1: Squared gain of BK($a$) and ideal filter with pass band [1/20, 1/4] (frequency as number of cycles per period).

$\omega = 0$ is exactly zero. This ensures that linear, quadratic and stochastic trends up to order two are completely removed, which is desirable for economic analysis. The resulting filter has weights

$$g_l^* = \begin{cases} g_l + \theta & \text{for } -a \leq l \leq a \\ 0 & \text{for } |l| > a \end{cases}, \qquad \theta = \frac{-\sum_{l=-a}^{a} g_l}{2a+1}, \tag{8}$$

where $g_l$ are the weights of the ideal filter (7). We will refer to this as the BK or BK($a$) filter.

The BK filter has the desirable properties that it induces no phase shifts (due to the symmetry of the weight), removes most common trends and approximates the ideal band pass filter when using high truncation lags. This is illustrated in figure 1, which plots the squared gain of the BK filter with various maximum lags $a$, together with the squared gain of the ideal filter. The most obvious drawback of the filter is that increasing filter accuracy (in terms of approximating the ideal FRF) is attained at the cost of losing observations at the start and the end of the sample. In particular, a BK($a$) filter gives filtered output $y_t$ for only $t = a+1, a+2, \ldots, T-a$. As we can see from the illustration, even sacrificing 20 observations at the start and end of the sample yields a frequency response that is quite different from that of an ideal filter.

A second popular filter based on approximating the ideal band pass filter was devised by Christiano and Fitzgerald (2003). Using an alternative loss criterion and assumptions on the underlying process of the $x_t$ they derive procedures to adjust the weights to take account of the "missing" observations. Essentially, the observations are extrapolated beyond the observed sample using the assumed model, the full (infinite) ideal filter weights are applied and the results are incorporated in modified weights for the first and last observations. Using a Random Walk assumption for the process $x_t$ leads to particularly simple adjustments:

$$\tilde{g}_0 = g_0/2 \tag{9}$$

$$\tilde{g}_1 = -g_0/2 \tag{10}$$

$$\tilde{g}_l = -g_0/2 - \sum_{k=1}^{l-1} g_l, \quad l \geq 2 \tag{11}$$

where again $g_l$ are the weights of the ideal filter (7). The adjusted weights $\tilde{g}_l$ are applied on the end points $x_1$ and $x_T$; the observations in between are weighted with the unmodified weights $g_l$.

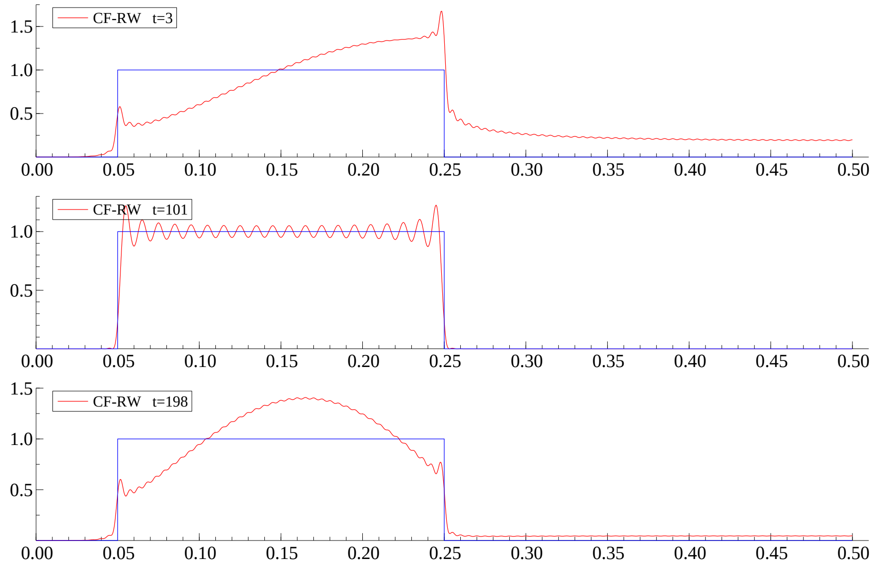


Figure 2: Squared gain of ideal and CF-RW filter with $T = 201$ observations at various times $t$, with pass band [1/20, 1/4] (frequency as number of cycles per period).

We will refer to this filter as the CF-RW filter, and the general case (with an unspecified model for $x_t$) as the CF filter.

Compared to the BK filter, the CF filter has advantages and drawbacks. It does not require sacrificing observations to gain accuracy, and experimentation shows that it usually outperforms the BK($a$) filter with low truncation values $a$. However, the accuracy can depend on the Random Walk assumption. Christiano and Fitzgerald show that for many economic time series the CF-RW filter yields good results, even when a Random Walk assumption is unlikely to hold. However, this does not apply universally. Furthermore, the filter is not symmetric, and therefore causes phase shifts. Finally, the filter is time varying, since calculating the filtered value $y_t$ requires different adjustments to the end point weights for each $t$. This implies that the FRF of the filter for $y_t$ also depends on $t$. We illustrate this in figure 2, which shows the squared gain of the CF-RW filter near the start, middle and end of the sample. In the majority (middle) part of the sample, the filter performs quite well. However, the accuracy deteriorates near the end points of the sample.

### 2.3 Direct frequency filtering

Given that the desired properties of band pass filters are formulated in the frequency domain, it appears natural to perform the filtering in the frequency domain. A straightforward way to do so is as follows.

1. Calculate the discrete Fourier transform of the time series

$$J_j = \frac{1}{T}\sum_{t=0}^{T-1} x_t e^{-i\omega_j t}, \qquad \omega_j = 2\pi j/T, \tag{12}$$

   for $j = 0, \ldots, T-1$. Note that we index the series $x_t$ here with $t = 0, 1, \ldots, T-1$, as common in Fourier analysis.

2. Multiply the Fourier coefficients $J_j$ with the desired frequency response function to obtain $\tilde{J}_j$. In the case of a band pass filter, this amounts to retaining only the coefficients in the pass band.

3. Apply the inverse discrete Fourier transform

$$x_j = \sum_{j=0}^{T-1} \tilde{J}_j \exp(i\omega_j t) \tag{13}$$

on the filtered coefficients to transform the series back into the time domain.

This is a *direct frequency filter*, in contrast to the BK and CF filters, which take place in the time domain using coefficients that approximate the desired frequency response function. The obvious advantage of a direct frequency filter is that the frequency response is exact. A further advantage is that no observations at the start or end of the sample are lost.

Note that for small samples it may be desirable to apply the Fourier transforms on the expanded sample $x_0, x_1, \ldots, x_{T-1}, x_T, x_{T+1}, \ldots, x_{T^*-1}$ with $x_t = 0$ for $T+1 \leq t \leq T^* - 1$. As we increase $T^*$ this allows us to calculate the Fourier transform on arbitrary frequencies $\omega \in [0, 2\pi]$

$$J(\omega) = \frac{1}{T} \sum_{t=0}^{T-1} x_t e^{-i\omega t}. \tag{14}$$

Although direct frequency filtering seems appealing given that the exact FRF can be applied, it does not perform well with small sample data. Whereas the time domain approximations from the previous subsections give distortions in the FRF due to the finite samples, direct frequency filtering suffers from the related *spectral leakage* phenomenon. Leakage occurs when applying a Fourier transform to a finite sample of a continuous function. Since a perfectly periodic function extends in the infinite future and past, it cannot be unambiguously resolved from a finite sample. In a discrete Fourier transform, this translates into a large peak around the true periodic frequency, and several smaller peaks spread out over adjacent frequencies.

Note that both the time domain approximations and direct frequency filter suffer from approximation errors due to the finiteness of the data. In filtering, they manifest themselves in different places, namely in the FRF approximation of the time domain filters, and in the Fourier transform in the direct frequency filter. Both methods gain accuracy when the sample size can be increased. However, for economic analysis this is usually not possible. Furthermore, economic time series are often dominated by trending behaviour, which typically result in very larger leakage errors. To make this explicit, an upper boundary of squared error of the direct frequency filter can be derived as

$$|y_t - \hat{y}_t|^2 \leq M^2 \Big( \sum_{l \in Q} |g_l| \Big)^2, \qquad Q = \{-\infty, \ldots, t-T\} \cup \{t+1, \ldots, \infty\}, \tag{15}$$

where $g_l$ are the implied weights and

$$M = \max_{0 \leq \omega \leq 2\pi} |T \cdot J(\omega)|, \tag{16}$$

the maximum of the modulus of the Fourier transform times $T$ (see e.g., Steehouwer (2005, Ch. B.15)). Thus the filter error (15) is bounded by the product of two factors, $M^2$, which is proportional to the periodogram of the data

$$P(\omega) = \frac{T}{2\pi} \big| J(\omega)^2 \big| \quad \text{for } \omega \in [0, 2\pi] \tag{17}$$

and $\big(\sum_{l \in Q} |g_l|\big)^2$, which decreases with the sample size.

# 3 A zero phase shift filter

In the previous section we discussed some common frequency based filtering approaches that attempt to approximate an ideal band pass filter. The BK filter and the CF filter are time domain approximations, which can show substantial deviations in the FRF compared to the ideal filter, with few obvious ways to improve them. A direct frequency filter comes close to the optimal properties, but shows leakage errors when using finite data. However, because of leakage effect errors when using small sample sizes, and the fact that the periodograms of macroeconomic time series will typically have a large maximum value caused by trending behaviour, this technique also seems inadequate. The leakage effect as the source of the filter errors in case of the direct frequency filters however suggests adjustments that can lead to an appropriate filtering technique.

The leakage effect originates from transforming a finite sample time series into the frequency domain. If we are able to (partially) avoid this transformation, the leakage effect and thereby the filter error would decrease. The leakage effect decreases as the sample size increases, or as the maximum value of the periodogram decreases. Although it is usually not possible to increase the number of observations of an empirical time series, reduction of the periodogram is achievable. This can be done by removing the dominant components at certain frequencies from the time series. These components should then be filtered separately while the remainder of the time series is filtered using a direct frequency filter. The more we would be able to reduce the maximum value of the periodogram without significantly changing the filter output, the bigger the reduction in the filter error would be. It turns out that such an optimal filtering technique can indeed be constructed. This section describes the workings of what will be called a Zero Phase (ZP) frequency filter which can be used to filter any, possibly short sampled, time series following any desired gain function while leaving the phase properties of the time series unchanged with as little a filter error as possible. Furthermore, the filter does not lead to any loss of observations at the beginning and end of the sample.

## 3.1 Basic algoritm

Suppose that we are able to split some infinite time series $x_t$ that needs to be filtered into a periodic component with frequency $\theta$ and a remainder part $r_t$,

$$x_t = \alpha \cos(\theta t) + \beta \sin(\theta t) + r_t. \tag{18}$$

For linear filters, filtering $x_t$ can be performed separately on both components:

$$\begin{aligned} y_t &= \sum_{l=-\infty}^{\infty} g_l x_{t-l} && (19) \\ &= \sum_{l=-\infty}^{\infty} g_l \big(\alpha \cos(\theta(t-l)) + \beta \sin(\theta(t-l))\big) + \sum_{l=-\infty}^{\infty} g_l r_{t-l}, && (20) \end{aligned}$$

Since $x_t$ is assumed to be an infinite time series, the full set of weights $g_l$ of the ideal filter can be applied. For the periodic component, the output of applying an ideal filter is known. If $\theta$ lies within the pass-band of the filter, the periodic component remains unaltered, otherwise the component is completely suppressed by the filter, that is,

$$\sum_{l=-\infty}^{\infty} g_l \big(\alpha \cos(\theta(t-l)) + \beta \sin(\theta(t-l))\big) = \begin{cases} 0 & \text{if } G(\theta) = 0 \\ \alpha \cos(\theta t) + \beta \sin(\theta t) & \text{if } G(\theta) = 1 \end{cases}. \tag{21}$$

The periodic component can therefore be exactly filtered without making a transition into the frequency domain, thereby avoiding leakage effect errors. Furthermore, if we choose the parameters $\theta$, $\alpha$ and $\beta$ of the periodic component such that the periodogram of the remainder $r_t$ has a smaller maximum than the periodogram of the original time series $x_t$, then the error of filtering this remainder time series will be smaller than the error of filtering the original time series. In that case, the total leakage effect and hence the total filter error will have decreased.

In order to minimise the filter error, we use a least squares criterion to estimate the periodic parameters. Since

$$\frac{1}{T} \sum_{t=0}^{T-1} r_t^2 = \frac{T}{2\pi} \int_0^{2\pi} |J_r(\omega)|^2 \, d\omega, \tag{22}$$

it follows that

$$|J_r(\omega)|^2 \, d\omega \le \frac{1}{T^2} \sum_{t=0}^{T-1} r_t^2. \tag{23}$$

Note that we assume a finite sample here. Since $d\omega$ is fixed, by reducing the sum of squares of $r_t$ the periodogram of the remainder time series at all frequencies is bounded from above by a lower value. For a given $\theta$, the minimisation problem can be solved by OLS applied to (18). The full optimisation problem requires minimisation of the objective function

$$V(\alpha, \beta, \theta) = \sum_{t=0}^{T-1} \Big( x_t - \alpha \cos(\theta t) + \beta \sin(\theta t) + r_t \Big)^2, \tag{24}$$

which is non-linear in $\theta$. Most standard procedures for solving non-linear least squares problems are based on Newton methods, using the gradient of the objective function to iteratively move towards to optimum. Unfortunately, for practical data sets the objective function in our problem usually contains many local extrema. Bloomfield (1976) suggests to use the algorithm of Brent (1973) for the minimization of (24). The Brent method is not based on derivatives, but requires as input a starting interval $[a, b]$ which is known to contain the optimum. Such an interval should not be chosen too wide in order to prevent obtaining a local minimums as the solution, but wide enough to contain the actual optimum. Practically, we first search over a fine grid on the interval $[0, \pi]$ to obtain a suitable starting interval for the Brent algorithm. Such an interval runs from the first point to the left of the optimal point on the grid until the first point to the right of the optimal point on the grid. Although the Brent algorithm cannot be guaranteed to find the global minimum of arbitrary non-linear functions, we have found the algorithm to be quite robust for the criterion functions of the form (24) with economic time series $x_t$, when using a grid with hundreds of points.

For a fixed $\theta$, the minimisation of (24) is a simple OLS problem, with solutions for $\alpha$ and $\beta$ that can to a large extent be written explicitly as a function of $\theta$, see appendix A. Therefore, we only need to use the Brent algorithm for the optimisation of (24) over $\theta$. This will be important in practice, as in the final algorithm the minimisation procedure has to be performed many times in two nested iterations.

## 3.2 Extensions

Three extensions to the basic algorithm from the previous subsection are described here, before proceeding to the complete filtering algorithm.

First, aside from a periodic component, a constant term $\mu$ can be estimated and filtered separately. Since a constant term can be regarded as a special case of a periodic component (with a frequency $\theta = 0$) it is not a priori clear whether the inclusion of an explicit constant term improves the overall filtering results. Using simulated data we can investigate this by comparing the filter error with and without an assumed constant term.

As a second extension, rather than estimating a single periodic component, we estimate $m$ periodic component simultaneously. Including both a constant term and multiple periodic components, the basic decomposition (18) becomes

$$x_t = \mu + \sum_{i=1}^{m} \Big( \alpha_i \cos(\theta_i t) + \beta_i \sin(\theta_i t) \Big) + r_t, \tag{25}$$

with the corresponding objective function

$$V(\mu, \alpha_1, \beta_1, \theta_1, \ldots, \alpha_m, \beta_m, \theta_m) = \sum_{t=0}^{T-1} \Big( x_t - \mu - \sum_{i=1}^{m} \big( \alpha_i \cos(\theta_i t) + \beta_i \sin(\theta_i t) \big) \Big)^2. \tag{26}$$

Minimising this function given $\theta_1, \ldots, \theta_m$ is again a simple linear least squares problem, but full optimisation is a difficult high-dimensional non-linear problem. To find a solution, we use a *cyclical descent* algorithm (Bloomfield (1976)). In the cyclical descent procedure, the optimisation is performed iteratively, solving the optimisation over a small subset of parameters in each step, conditional on all the remaining parameters. The iterations terminate when the improvement in the overall objective function becomes very small. In our problem, we iterate over the subsets

$\{\mu\}, \{\alpha_1, \beta_1, \theta_1\}, \ldots, \{\alpha_m, \beta_m, \theta_m\}$, that is, in each cyclical descent iteration, the function

$$V(\alpha_k, \beta_k, \theta_k) = \sum_{t=0}^{T-1} \Big( x_t^* - \alpha_k \cos(\theta_k t) + \beta_k \sin(\theta_k t) \Big)^2. \tag{27}$$

is minimised, with

$$x_t^* = x_t - \mu - \sum_{i \neq k} \big( \alpha_i \cos(\theta_i t) + \beta_i \sin(\theta_i t) \big). \tag{28}$$

Optimising over $\{\alpha_k, \beta_k, \theta_k\}$ is done by using the Brent method to find a minimising $\theta_k$ and solving the OLS problem for $\alpha$ and $\beta$, as described in the previous subsection.

As a third extension, we perform another level of iteration. Recall that we attempt to reduce the direct frequency filter error by trying to reduce the variance of the remainder $r_t$ after fitting trigonometric series. Since $r_t$ will likely still contain periodic elements, we can apply the same argument on $r_t$ as on the original $x_t$. Rather than applying the direct frequency filter immediately on the remainder $r_t$ in the decomposition (25), we treat $r_t$ as a new series $x_t$ to be filtered. By iterating this several times, the final remainder can be reduced to a very small level, thereby rendering the leakage effect of the direct filtering step negligible.

Both the second and third extensions are aimed at reducing the variance $r_t$ before using the direct frequency filter. It is not clear theoretically whether either, or a combination of both iterations are preferable to achieve this. Again this can be investigated using simulated data.

## 3.3 Algorithm summary and discussion

The final algorithm of the ZP frequency filter is a combination of the basic algorithm and the three extensions from the previous subsections. Filtering a time series $\{x_t, t = 0, \ldots, T-1\}$ based on a frequency response function $G(\omega)$ consists of the following steps.

**Initialisation**

1. Calculate $M_0$ as the maximum of the periodogram of $x_t$.
2. Choose the number of periodic components $m$ and decide whether or not a constant term $\mu$ is to be included.

**Iteration** $j$

1. Estimate the parameters $\{\mu_j, \alpha_{j,1}, \beta_{j,1}, \theta_{j,1}, \ldots, \alpha_{j,m}, \beta_{j,m}, \theta_{j,m}\}$ of (26) using cyclical descent over $\{\mu\}, \{\alpha_1, \beta_1, \theta_1\}, \ldots, \{\alpha_m, \beta_m, \theta_m\}$, and the Brent algorithm for the sub-problems and calculate the remainder time series $r_t$.
2. Calculate $M_j$ as the maximum of the periodogram of $r_t$.
3. Stop if the upper boundary on the filter error is reduced far enough in a relative sense. This can be seen by verifying whether $M_j/M_0$ is small enough. If not, start iteration $j+1$ using the time series $r_t$ as input.

**Finish**

1. Use the direct frequency filter to calculate the filtered version of the remainder time series $r_t$ from the last iteration and call this series $s_t$.
2. Calculate the total filtered time series $y_t$ as the sum of $s_t$ and the constant terms and estimated periodic components of all, say $k$, iterations multiplied by the value of the FRF at the corresponding frequencies. That is

$$y_t = s_t + \sum_{j=1}^{k} \Big( G(0)\mu_{j,i} + \sum_{i=1}^{m} G(\theta_{j,i}) \big( \alpha_{j,i} \cos(\theta_{j,i} t) + \beta_{j,i} \sin(\theta_{j,i} t) \big) \Big) \tag{29}$$

For practical applications, the following choices remain to be made:

(a) The number of periodic components $m$ to be estimated simultaneously.

(b) Whether or not to include a constant term $\mu$.

(c) The minimal (relative) reduction of the upper boundary on the filter error $M_j/M_0$ which stops the iterations.

(d) The minimal (relative) reduction of the objective function in the cyclical descent algorithm which stops the iterations.

(e) The number of points on the grid before applying the Brent algorithm.

(f) The resolution of the Fourier transform when applying the direct frequency filter on the final remainder time series.

From the results of several simulation experiments (see Steehouwer (2005, Ch. 5.5), and from our experience with filtering a large variety of economic time series, including production, unemployment, price indices for consumer goods, stocks, real estate, commodities, interest rates, and credit spreads we recommend the following settings as a good default:

(a) Using $m = 20$ or more simultaneously components per iteration helps in accurate filtering. That is, in each iteration stop, we fit 20 or more trigonometric series, and proceed to the next iteration step with the least squares residual.

(b) A constant term $\mu$ should not be included, unless there is specific evidence or motivation to do so. Its inclusion is essentially an a priory parameter restriction, which detracts from the filter fit if it is not appropriate.

(c) A minimal value of $M_j/M_0 = 10^{-5}$ is satisfactory.

(d) A stopping criterion of $10^{-4}$ in the cyclical descent iteration is satisfactory.

(e) Set the number of points on the grid to 500 before applying the Brent algorithm.

(f) The resolution of the Fourier transform when applying the direct frequency filter on the final remainder time series can be set to $2^{16}$.

Our experience shows that these settings produce robust filter output over hundreds of economic time series. Although the filter appears to depend on more user specified settings than the BK and the CF filters, we regard this generally as an advantage. Increasing grid points, simultaneous periodic components and tightening stopping criteria are easy ways decrease the approximation error, at the cost of increased computing time. The BK and CF filters, by contrast require removing data or choosing better approximating models to decrease the error.

To conclude this section, we summarise here the arguments that motivated the ZP filter design.

1. The ideal way of designing a filter with the desired frequency response characteristics is in the frequency domain.
2. Transforming small sample time series (as typically encountered in economics) into the frequency domain induces significant errors due to leakage effects, especially at the boundaries of the sample.
3. Leakage can be reduced by increasing the sample size (rarely possible in practical economic applications) or by decreasing the maximum of the periodogram of the series.
4. An upper boundary on the filter error can be minimised by minimising the sum of squared residuals when fitting a trigonometric function to the time series.
5. For the perfectly periodic trigonometric function, the filter output is known without having to transform it into the frequency domain. Thereby leakage errors can be avoided.

Note that in the ZP filter, we iteratively fit trigonometric functions until the remainder is very small. The vast majority of the filtering is therefore actually performed in the time domain. However, since we fit perfectly periodic functions, the *exact* FRF can be applied, in contrast with other time domain approximations.

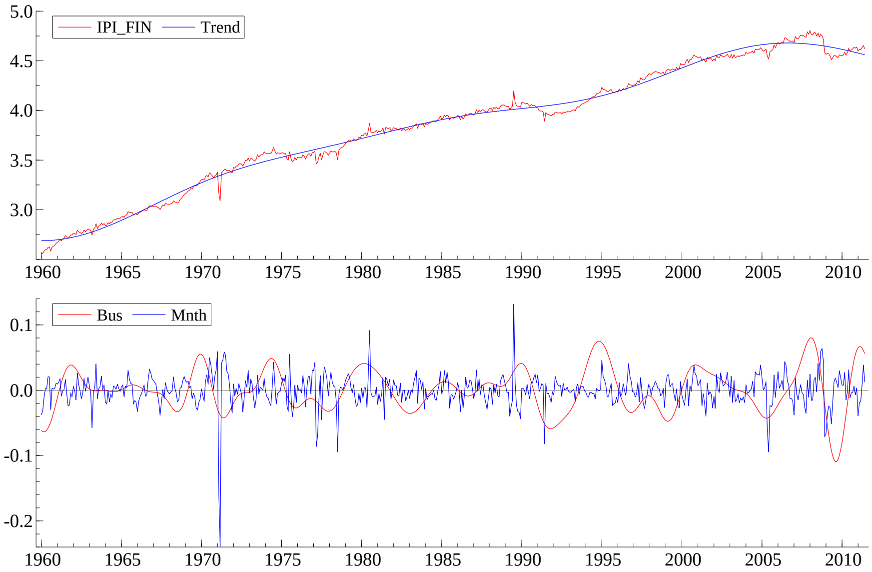


Figure 3: Finland IPI decomposition into trend, cycle (Bus, from 24 to 96 months), and high frequency (Mnth)

# 4 Empirical illustrations

## 4.1 Frequency decomposition of macro-economic time series

In this subsection, we compare the results of the ZP-filter with those of the popular BK and CF filters. We apply the three filter procedures to industrial production index (IPI) series and GDP series for the Netherlands and Finland.

The IPI data are monthly series starting in January 1960, obtained from OECD Main Economic Indicators. For GDP we use a long series of annual observations of real (volume) production, dating back to the nineteenth century. The early part of the data (1820–1948) was taken from (Maddison, 1995). More recent numbers were obtained from GlobalFinancialData and Bloomberg, which are based on national accounts.

Figures 3 and 4 show the log IPI series for Finland and the Netherlands, together with a trend, business cycle (Bus) and high frequency (Mnth) decomposition. The pass band for the business cycle is defined to have period lengths between two and eight years; faster and slower movements are assigned to the high frequency and the trend components respectively. The Finland series shows a number of large outliers, especially noticeable in 1971 and 1989. Nevertheless, these appear to be well isolated by the high frequency filter. The Dutch series does not show large outliers. In figure 5 series of both countries are plotted together, together with their isolated business cycles, which have similar magnitudes. The cycle for both countries in the recent past (post-2007) seem very similar.

Figure 6 shows the business cycle for Finish IPI, extracted by the ZP, the BK(24) and the CF-RW filters. The top panel shows the cycle over the entire sample, while the bottom panel focusses on the recent past. We can see that generally the ZP and the CF filter provide similar results, while the BK shows more discrepancy. The largest differences between ZP and CF filter output are seen near the ends of the sample, which is clearly seen in the bottom panel. This should not come as a surprise, since in the middle part of the sample the CF filter is largely identical to the ideal band pass filter, while near the end points, it relies more heavily on the Random Walk extrapolation assumption.

In our next example, we show the application of filters to Dutch annual GDP series starting from 1820. While recent business cycle literature usually adopts a business cycle definition of fluctuations between 6 and 32 quarters periods, essentially following Burns and Mitchell (1946),

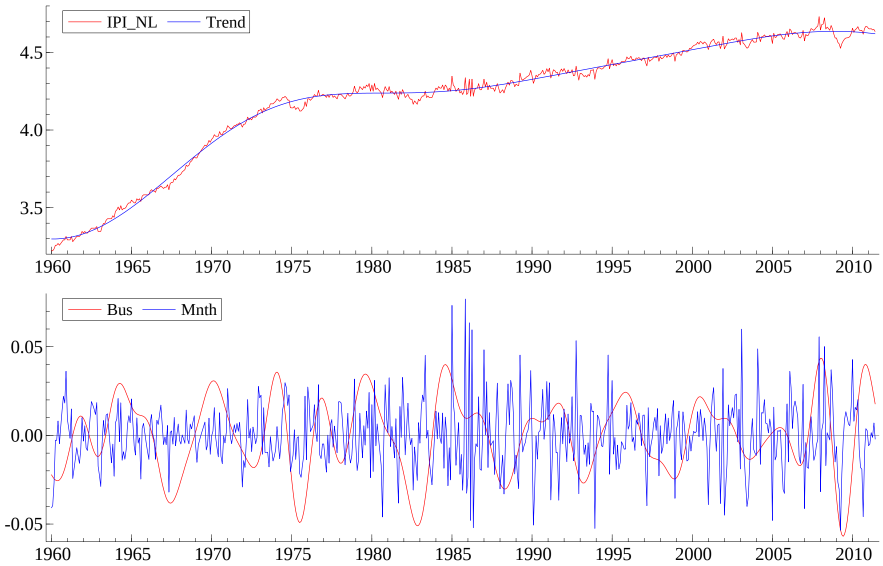


Figure 4: Netherlands IPI decomposition into trend, cycle (Bus, from 24 to 96 months), and high frequency (Mnth).

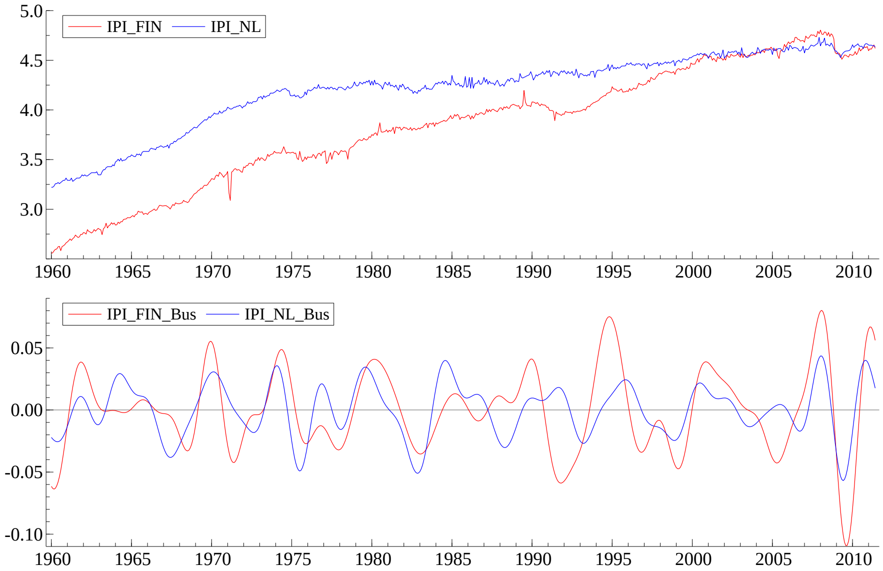


Figure 5: IPI series and business cycle for Finland and the Netherlands.

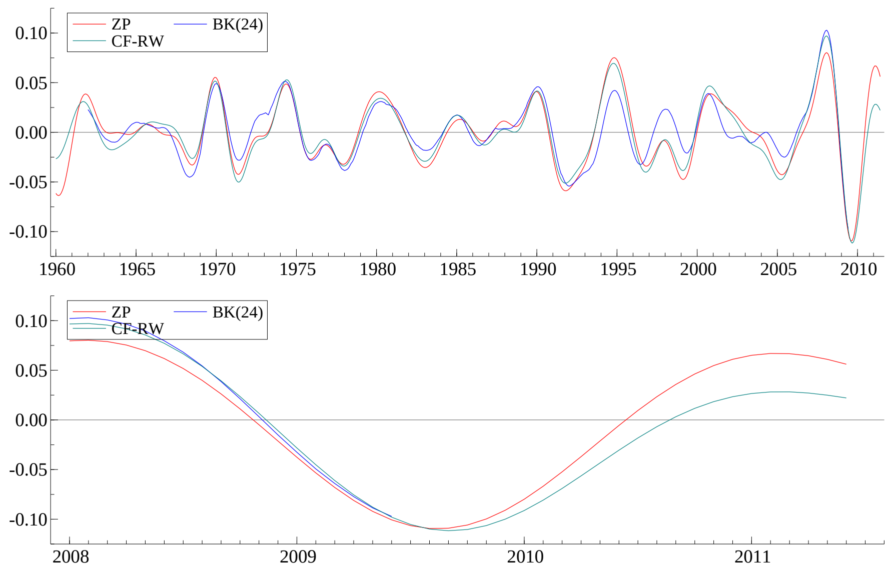


Figure 6: Zero Phase, Baxter-King and Christiano-Fitzgerald business cycle for Finland IPI.

on annual data with a long history it seems interesting to investigate the presence of longer cycles. To this end, we filter the series on fluctuations with periods between 30 and 70 years. This filter may capture longs waves, which are often associated with the works of Kondratieff. Our position is agnostic concerning economic rationales for the existence of such waves. From a frequency domain perspective, filtering on this pass band is a relatively straightforward exercise. Here, we are mainly interested in the output of different filter approaches.

The filtered long waves using the ZP, BK and CF filters are shown in figure 7. As in the IPI decomposition, the ZP and CF filter provide quite similar output, especially in the middle part of the sample. Near the end points, where the ZP and CF disagree, we advocate the ZP filter on the grounds that it provides by construction the proper frequency response, while the CF frequency response is quite distorted at the edges, as we saw in section 2.2.

Separate from the long wave, we show the output of filtering the post-1950 sample on periodicities faster than 15 years in figure 8. In this example, all three filters show very closely matching output, again differing mainly near the end points.

## 4.2 Pass band selection using multi-period returns

In this subsection, we consider the modelling of a mix of different returns calculated at different horizons from a single index. Consider a U.S. stock market index series (in log), decomposed into three distinct components with the ZP filter:

1. A trend with period longer than 16 year
2. A cycle with periods between 2 and 16 years
3. A high-frequency component with periods between 2 years and 2 months

Figure 9 shows the index and the filtered components, which display long, medium and short term dynamics in the series. Practitioners often extract information of long, medium and short run dynamics using log-returns over different horizons. For example, the top panels in figures 10–12 show 8 year, annual and monthly returns. The figures also contain the same multi-period returns calculated from the the filtered components using the trend, cycle and high frequency filters defined above. As can be seen, the multi-period returns show similar long, medium and short term, dynamics as the filtered components. However, they also show much extraneous

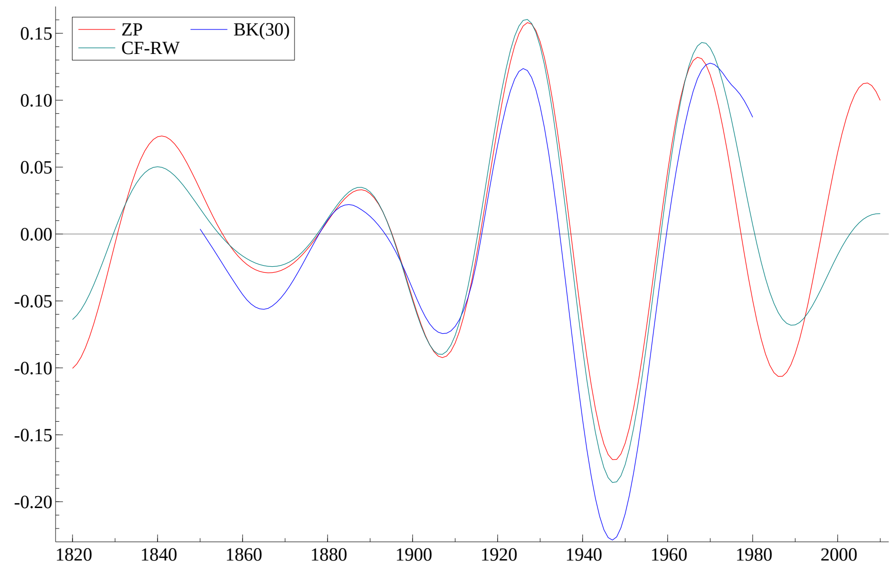


Figure 7: Netherlands production, long wave (30–70 years).

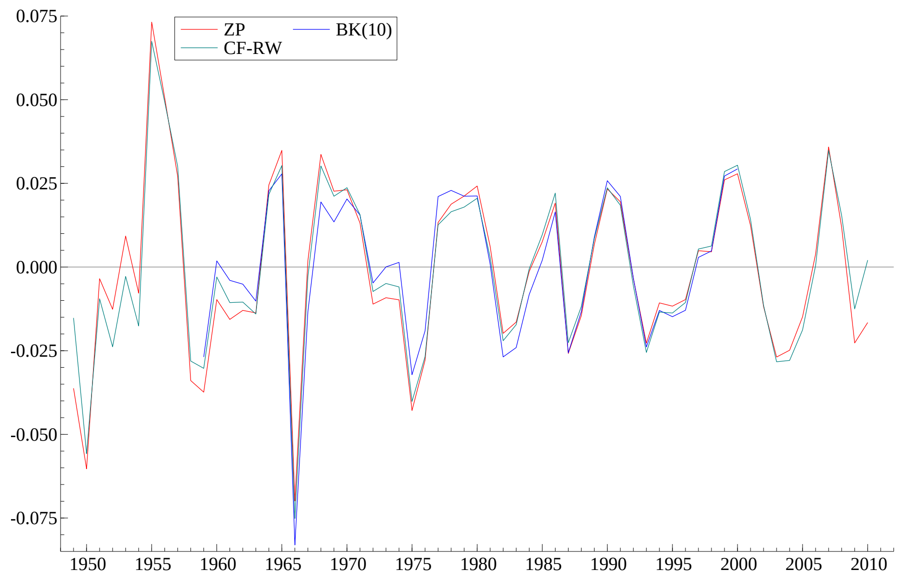


Figure 8: Netherlands production, short fluctuations (15 year and faster).

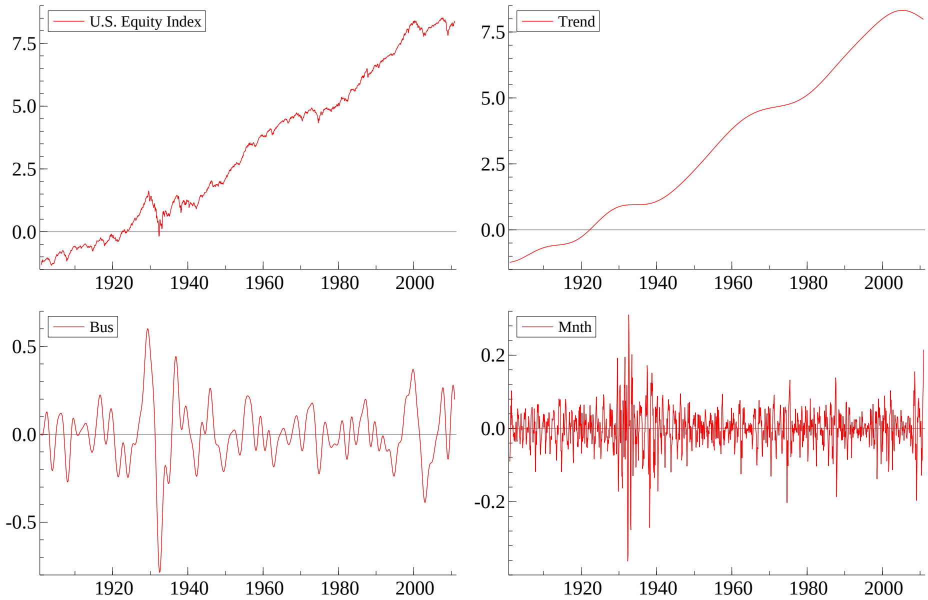


Figure 9: Frequency decomposition of U.S. Equity Index, Trend (slower than 192 months), cycle (Bus, 24 to 192 months), high frequency (Mnth, faster than 24 months).

variance, corresponding with frequencies outside frequencies of interest. Unlike when using non-overlapping band pass filtering, different multi-period returns do not provide a proper orthogonal decomposition. This makes it difficult to devise different models for returns over different horizons simultaneously in a consistent manner. However, the similarity of the dynamics of multi-period and filtered returns suggests an alternative approach. Given three separate return horizons of interest (long, medium and short term movements), we can choose pass bands boundaries for a band pass decomposition by minimising the difference between the multi-period returns and their filtered outputs. This will allow us to use separate models for different components, which can be summed to provide a model of the total index. The returns calculated on different horizons of these filtered components match those of the returns calculated from the original series as close as possible.

With the previous motivation, we attempt select pass band boundaries which render filtered returns close to multi-period returns using a minimum sum of squared difference criterion. This can in principle be done in the time domain, using a good quality filter and multiple iterations of least squares estimation. For this problem, however, a direct frequency domain solution seems more suitable. In order to do so, we need to examine the frequency characteristics of both the data and the employed filters.

A $k$-period (log-)return is the difference between the current (log-)index value and its value $k$ periods earlier. The operator is identical to a seasonal difference $\Delta_k = (1 - L^k)$, which is a linear filter with weight 1 at lag $l = 0$, $-1$ at lag $l = k$ and zero elsewhere. The square gain of this filter is straightforwardly derived as

$$\begin{aligned} Gain^2(\omega) &= |1 - e^{-ik\omega}|^2 && (30)\\ &= (1 - e^{-ik\omega})(1 - e^{+ik\omega}) && (31)\\ &= 2 - 2\cos(k\omega), && (32) \end{aligned}$$

while its phase shift is

$$Ph(\omega) = \tan^{-1}\left(\frac{-\sin k\omega}{1 - \cos k\omega}\right). \quad (33)$$

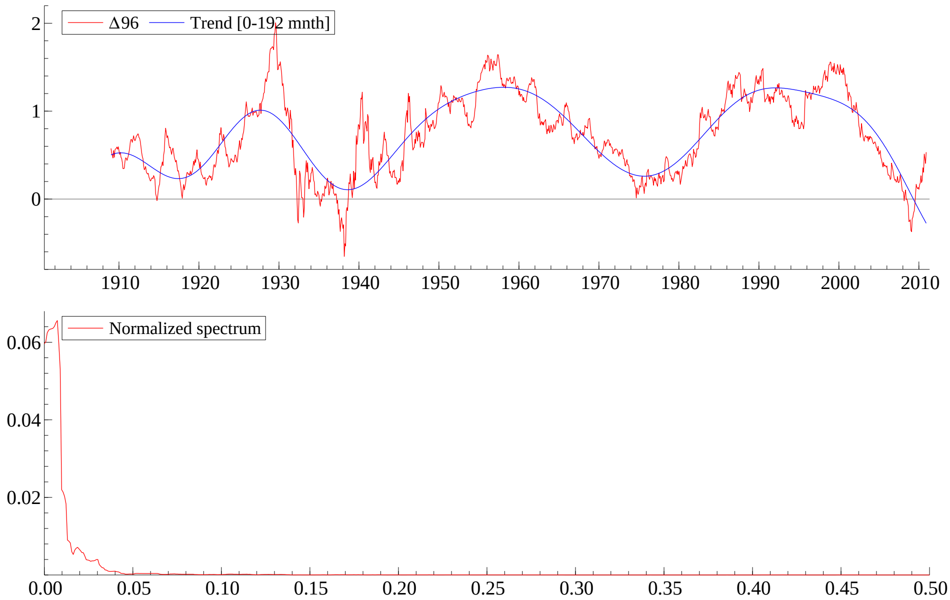


Figure 10: 96 month returns, filtered returns, spectrum (in cycles per period).

We plot the square gain function in the top panel of figure 13 for $k = 1, k = 12$ $k = 96$. From the gain function we can see that the filter suppresses the trend and periodic fluctuations at periods $j, j/2, j/3, \ldots$. Due to the linearity of the filters, the $k$-period return on a band pass filtered component can be obtained from the band pass filter applied to the $k$-period return of the index. The combined gain of applying three ideal band pass filters with pass bands $[0, a], [a, b], [b, 0.5]$ (given in units of cycles per period) on the original series is shown in the bottom panel of figure 13.

The spectrum of the multi-period returns of the filtered series can also be obtained by multiplying the spectrum of the data with the square gain of the ideal bandpaaa filter. In figures 10–12 we show the returns, the returns calculated on the filtered components, and the estimated return spectrum on 96, 24 and 1 month horizons. We use a two-sided moving average filter with 5 lags on the sample periodogram to obtain a simple smoothed sample spectrum of the data. As can be seen, the majority of the variance resides in the low frequency components.

Using the sample spectrum and the gain functions of the different band pass filters, we can adjust the pass band boundaries $a$ and $b$ such that the maximum amount of variance is retained over all three return horizons. This is done by moving $a$ and $b$ such that the smoothed spectrum multiplied by the different ideal (rectangular) pass bands, summed over all three components is maximised. Using Fast Fourier Transform routines to estimate the spectra, this can done rapidly with a simple grid search. Figure 14 shows the optimised pass band boundaries, which correspond to 93 and 16 months respectively. The spectra are enlarged versions of those in figures 10 and 11 that show more details of the interesting parts of the spectra. After band pass filtering, the left part of the 96 month return (smaller than 0.01070), the middle of the 12 month return spectra and the part to the right of 0.06116 of the 1 month return are retained. In figures 15 we show the returns filtered with the previous and the optimised pass bands. We can see that in the optimised case, the trend captures much more variation, which was previously assigned to the cycle component. Figure 16 shows original monthly returns and the filtered monthly returns in the last part of the sample. As can be seen, the movements of the series are very similar. The filtered returns show slightly less variation, which is as we intend, as the low frequency movements have been assigned to the filter versions of the filtered components on the lower frequency pass bands.

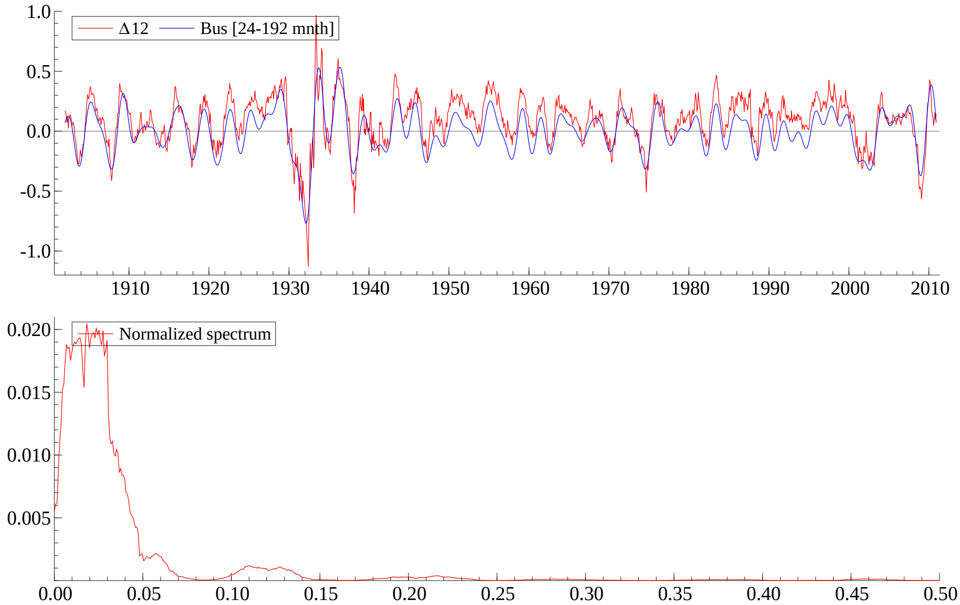


Figure 11: 12 month returns, filtered returns, spectrum (in cycles per period).

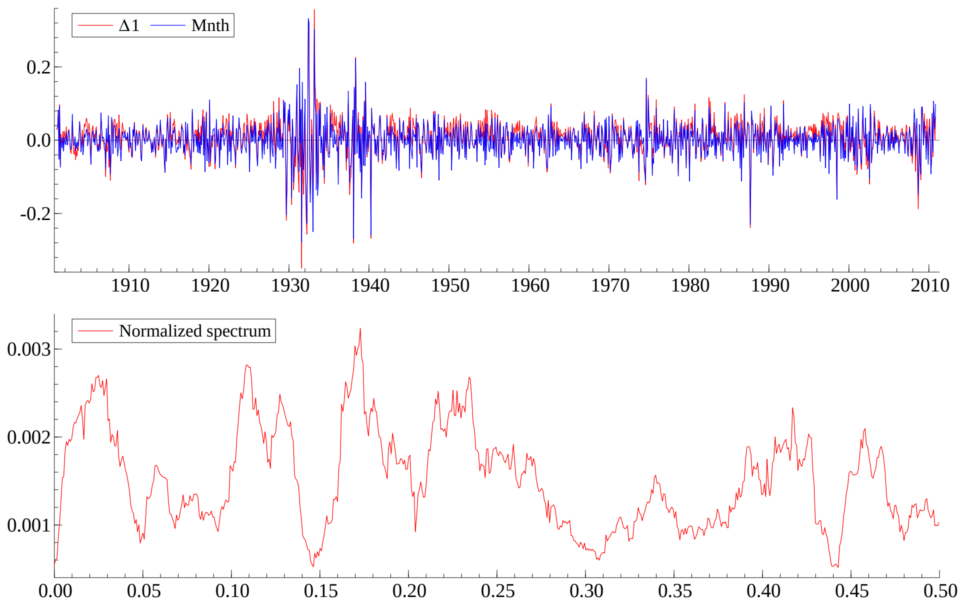


Figure 12: 1 month returns, filtered returns, spectrum (in cycles per period).

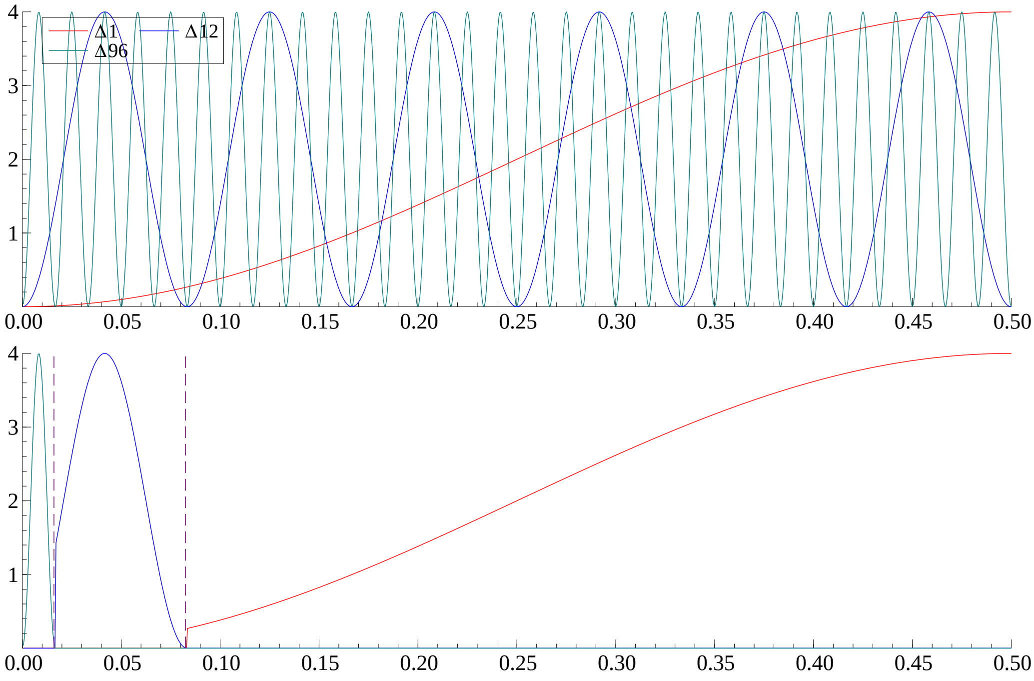


Figure 13: Squared gain of multi-period difference operator and band pass filter multi-period difference (frequencies in cycles per period).

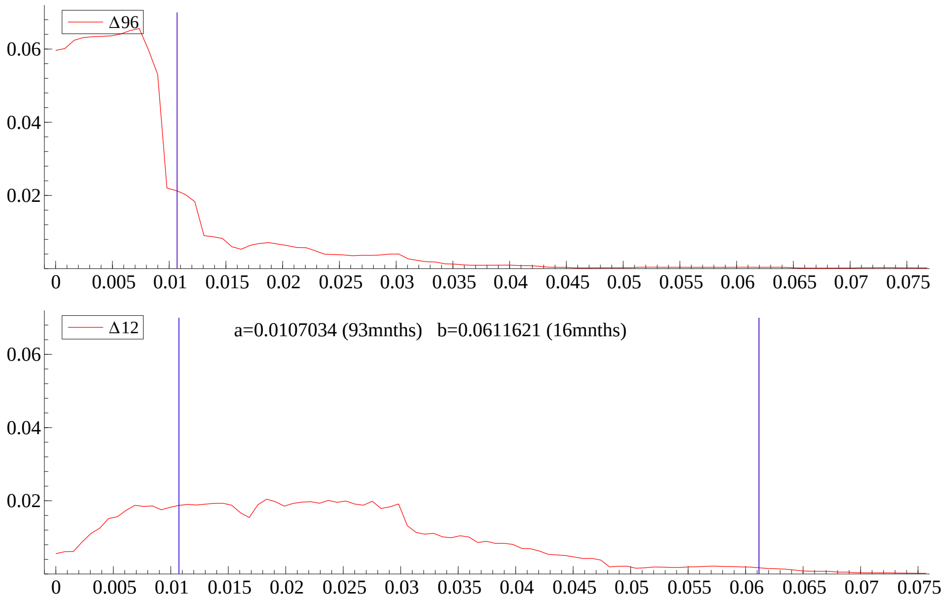


Figure 14: Optimised pass bands for filtering 96, 24 and 1 month returns.

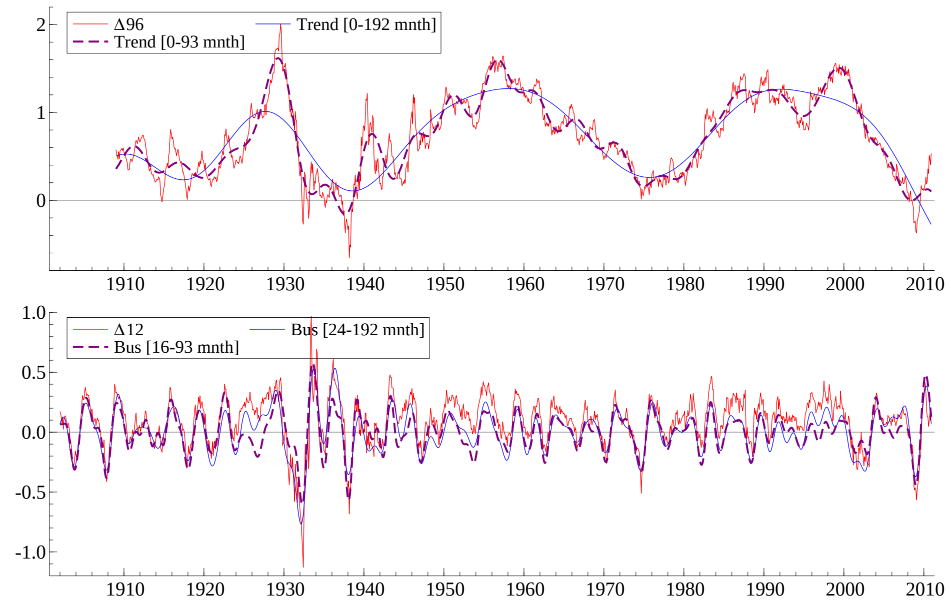


Figure 15: 96 and 12 month returns filtered with different pass bands.

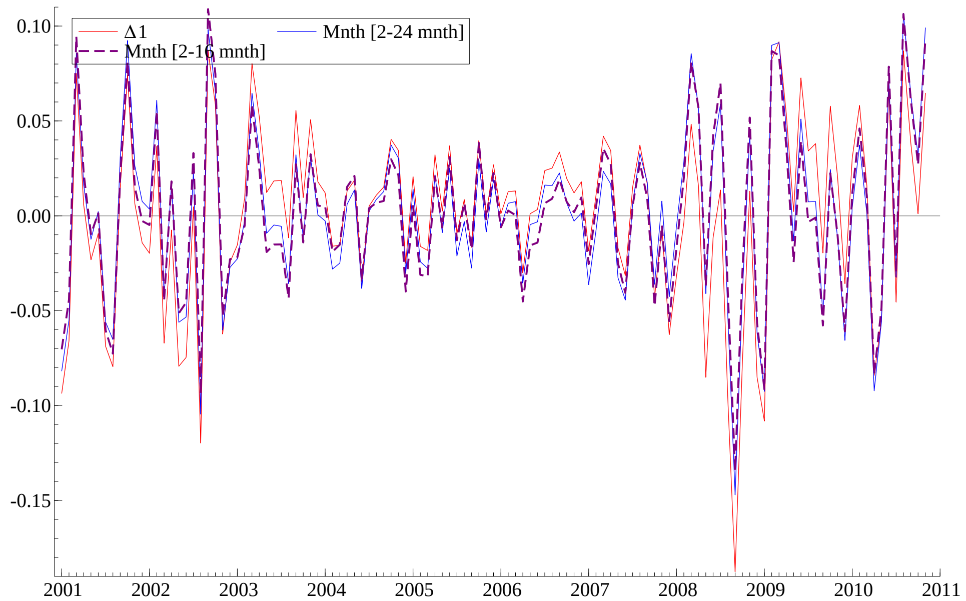


Figure 16: 1 month returns filtered with different pass bands.

## 5 Conclusion

In this paper we presented a frequency domain based band pass filter that allows decomposition of times series in trends, business cycle and high frequency fluctuations strictly on the basis of a frequency domain characteristics. The filter is an iterative algorithm that successively fits trigonometric functions the observed data in the time domain, and filters a small residual directly in the frequency domain. Since the frequency characteristics in both parts are known, the exact frequency response can be obtained.

The filter overcomes shortcomings of both direct frequency domain filters, and the Baxter-King and Christiano-Fitzgerald time domain approximating filters. In particular, spectral leakage is minimised, the filter does not induce phase shifts, uses all observations, and applies the correct frequency response on the entire sample.

We provide a detailed algorithmic description, and give two empirical illustrations. In comparing filtered macro-economic time series, we find that the filter gives similar results as the Christiano-Fitzgerald in most of the sample, but differs near at the end. This is especially pertinent when timely decisions are to be taken, based on recent data. In that case it is essential that accurate filtering at the end of the sample can be performed.

As a more general example of frequency domain methodology, we consider using band pass filtering to improve modelling of stock market returns at different horizons. Since such returns also contain fluctuations at frequencies which are also present at other horizons, it is not easy to model each of them separately in a consistent manner with each other. By applying band pass filters, we can provide a proper decomposition of the series on which the returns are based. We show that pass bands can be chosen such that the maximum amount of variation in the returns at different horizons are captured by the returns of the filtered series. This results in a decomposition from which returns closely match those of the original returns calculated at different horizons.

## References


Baxter, M. and R. G. King (1999). Measuring business cycles: Approximate bandpass filters for economic time series. *The Review of Economics and Statistics 81(4)*, 575–593.

Bloomfield, P. (1976). *Fourier analysis of time series: An introduction.* New York: Wiley.

Brent, R. (1973). *Algorithms for minimization without derivatives.* Englewood Cliffs, NJ: Prentice-Hall.

Burns, A. M. and W. C. Mitchell (1946). *Measuring Business Cycles.* New York: NBER.

Christiano, L. J. and T. J. Fitzgerald (2003). The band pass filter. *International Economic Review 44(2)*, 435–465.

Koopmans, L. H. (1974). *The spectral analysis of time series.* New York and London: Academic Press.

Maddison, A. (1995). *Monitoring the World Economy 1820-1992.* Paris: OECD.

Priestley, M. B. (1976). *Fourier analysis of time series: An introduction.* New York and London: Academic Press.

Schmidt, R. (1984). Konstruktion von digitalfiltern and ihre verwendung bei der analyse konomischer zeitreihen. Bochumer wirtschaftswissenschaftliche Studien 100.

Steehouwer, H. (2005). *Macroeconomic Scenarios and Reality: A Frequency Domain Approach for Analyzing Historical Time Series and Generating Scenarios for the Future.* Ph. D. thesis, VU University Amsterdam. `http://hdl.handle.net/1871/9058`.

Steehouwer, H. (2010). A frequency domain methodology for time series modeling. In A. Berkelaar, J. Coche, and N. K. (Eds.), *Interest Rate Models, Asset Allocation and Quantitative Techniques for Central Banks and Sovereign Wealth Funds.* Palgrave Macmillan.

# A Least squares solution

For a given $\theta$, the least squares solution to

$$\min_{\alpha,\beta} \sum_{t=0}^{T-1} \Big( x_t - \alpha \cos(\theta t) + \beta \sin(\theta t) + r_t \Big)^2 \tag{34}$$

is obtained by solving the normal equations

$$-2 \sum_{t=0}^{T-1} \big( x_t - \alpha \cos(\theta t) - \beta \sin(\theta t) \big) \cos(\theta t) = 0, \tag{35}$$

$$-2 \sum_{t=0}^{T-1} \big( x_t - \alpha \cos(\theta t) - \beta \sin(\theta t) \big) \sin(\theta t) = 0. \tag{36}$$

The solution is given by

$$\hat{\alpha}(\theta) = \Big( \sum_{t=0}^{T-1} x_t \cos(\theta t) \sum_{t=0}^{T-1} x_t \sin^2(\theta t) - \sum_{t=0}^{T-1} x_t \sin(\theta t) \sum_{t=0}^{T-1} x_t \sin(\theta t) \cos(\theta t) \Big) / \gamma, \tag{37}$$

$$\hat{\beta}(\theta) = \Big( \sum_{t=0}^{T-1} x_t \sin(\theta t) \sum_{t=0}^{T-1} x_t \cos^2(\theta t) - \sum_{t=0}^{T-1} x_t \cos(\theta t) \sum_{t=0}^{T-1} x_t \sin(\theta t) \cos(\theta t) \Big) / \gamma, \tag{38}$$

$$\gamma = \sum_{t=0}^{T-1} \cos^2(\theta t) \sum_{t=0}^{T-1} \sin^2(\theta t) \Big( \sum_{t=0}^{T-1} \sin(\theta t) \cos(\theta t) \Big)^2. \tag{39}$$

The calculations of the solution can be accelerated by using the relationships

$$\sum_{t=0}^{T-1} \cos^2(\theta t) = \frac{T}{2} + \frac{\sin(T\theta) \cos((T-1)\theta)}{2 \sin \theta} \tag{40}$$

$$\sum_{t=0}^{T-1} \sin^2(\theta t) = \frac{T}{2} - \frac{\sin(T\theta) \cos((T-1)\theta)}{2 \sin \theta} \tag{41}$$

$$\sum_{t=0}^{T-1} \cos(\theta t) \sin(\theta t) = \frac{\sin(T\theta) \sin((T-1)\theta)}{2 \sin \theta} \tag{42}$$

which hold for all $\theta \neq 0$.