

On the Model-Based Interpretation of Filters and the Reliability of Trend-Cycle Estimates

Tommaso Proietti*

Dipartimento di Scienze Statistiche, Università di Udine

Abstract

The paper is concerned with a class of trend cycle filters, encompassing popular ones, such as the Hodrick-Prescott filter, that are derived using the Wiener-Kolmogorov signal extraction theory under maintained models that prove unrealistic in applied time series analysis. As the maintained model is misspecified, inference about the unobserved components, and in particular their first two conditional moments, given the observations, are not delivered by the Kalman filter and smoother or the Wiener-Kolmogorov filter for the maintained model.

The paper proposes a model based framework according to which the same class of filters is adapted to the particular time series under investigation; via a suitable decomposition of the innovation process, it is shown that any linear time series with ARIMA representation can be broken down into orthogonal trend and cycle components, for which the class of filters is optimal. Finite sample inferences are provided by the Kalman filter and smoother for the relevant state space representation of the decomposition.

In this framework it is possible to discuss two aspects of the reliability of the signals' estimates: the mean square error of the final estimates and the extent of the revisions. The paper discusses and illustrates how the uncertainty is related to features of the series and the design parameters of the filter, the role of smoothness priors, and the fundamental trade-off between the uncertainty and the magnitude of the revisions as new observations become available.

Keywords: Signal Extraction, Revisions, Kalman filter and Smoother.

*Address for Correspondence: Via Treppo 18, I-33100 Udine. e-mail: proietti@dss.uniud.it.

1 Introduction

The separation of the trend from the cycle is a major issue in the analysis of the dynamic behaviour of macroeconomic variables, such as output, unemployment and inflation. Recent contributions, and in particular Orphanides and van Norden (2002), have focussed on the issue of the uncertainty with which signals are estimated in macroeconomics: for instance, given the relevance that measures of the output gap are assigned for the conduct of monetary policy, the econometric profession should provide a clear assessment of the reliability of such measures, including, *inter alia*, the evaluation of features that are related to the properties of the signal extraction filter, such as the final estimation error and the process of revision.

When the signals are estimated within a parametric approach, as in Harvey and Jäger (1993), this assessment is a natural by product of the modelling effort. Often, however, those measures are provided by the application of *ad hoc* filters that select certain features of the series without entertaining a model of the series dynamics; in other occurrences, which are the ones considered in this paper, the filter has a genuine model based interpretation, but the underlying model is clearly misspecified for the series under investigation. In all these occurrences it may not be immediately clear how the reliability of the corresponding signals should be evaluated.

This is the case for the Hodrick-Prescott filter (Hodrick and Prescott (1997), HP henceforth): the underlying local linear trend model, that decomposes the series into uncorrelated components represented by an integrated random walk trend plus pure white noise (see section 2 below), is usually inadequate for macroeconomic time series such as real gross domestic product. If the signal to noise ratio were estimated, rather than fixed, experience suggests that its value would result so large to render the trend indistinguishable

from the series; furthermore, the usual residual based diagnostics would definitively speak out against the maintained model.

The objective of this paper is to assess two important aspects of the uncertainty of the trend-cycle estimates arising from a class of filters, considered in Pollock (2000) and Gómez (2001), and nesting popular filters such as HP and rational square wave filters: the final estimation error mean square error (MSE) and the magnitude of the revision of the estimates at the end of the sample, as new observations become available.

This assessment is allowed for by the fact that the filters admit an interpretation within a model based framework: extending the approach initiated by Gómez (2001) and Kaiser and Maravall (2001), we show that it is possible to define a trend-cycle decomposition of any ARIMA process via a suitable decomposition of the ARIMA innovation process. The trends and cycles emerging from the decomposition are artificial, as they do not necessarily correspond to a mechanism that has generated the data; nevertheless, the decomposition furnishes the theoretical underpinning for framing the filters within the general theory of linear estimation. This assumes that the filters have autonomous justification, eg. as bandpass filters, an interpretation that we review in the course of the discussion.

Within the model-based framework, the class of filters yields the Wiener-Kolmogorov optimal filters of the components, given the availability of a doubly infinite sample. However, although the impulse responses for the central sample points are invariant, the MSE of the smoothed estimates depends on the time series model for the series. The paper provides an upper bound for it and discusses its dependence upon the filter design parameters. Moreover, the filtered estimates and the MSE of the components depend on the properties of the series under investigation, in that they vary according to the ARIMA process considered.

In sum, the model based framework allows correct inferences on the reliability of the

estimates of trends and cycles, and the paper discusses how the estimation MSE depends on both the features of series (for instance, the order of integration), and the parameters that regulate the design of the filter, discussing also the the role of smoothness priors.

The paper is organised as follows: section 2 introduces the class of filters that we concentrate upon, presenting the local trend model for which it is optimal and discussing the role of the main parameters. The frequency domain arguments which enforce the interpretation of as bandpass filters is also reviewed. Section 3 sets up the decomposition of any ARIMA process into trends and cycles that yield the same filters as the minimum mean square estimators of the components for a doubly infinite sample. In finite samples inferences are provided by the Kalman filter and smoother for the state space representation of the decomposition, which is given in the appendix. In section 4 we derive an upper bound for the MSE of the final estimate and discuss how it depends on features of the series, namely the order of integration, and the design parameters of the filter. Section 5 discusses further aspects of the uncertainty of the signal estimates. It presents an empirical example, referring to the U.S. real gross domestic product, a well known case study in the application of the HP filter, illustrating how the estimates of the cycle depend on the time series model adapted to the series, how the uncertainty is understated by the MSE outputted by the Kalman filter and smoother for the misspecified local linear trend model at the basis of the HP filter, and finally how the uncertainty depends on the cutoff frequency, and thus on the bandpass nature of the filter. Finally, the revision issue is addressed when the true model is ARIMA(1,1,0) and illustrate the fundamental trade-off between the reliability and the extent of the revision process. In section 6 some conclusions are drawn.

2 A class of trend-cycle filters

The class of filters considered in this paper arises from the application of the Wiener-Kolmogorov (WK) optimal signal extraction theory to the signal plus noise, or local trend, model:

$$\begin{aligned}
 y_t &= \mu_t + \psi_t, & t = 1, 2, \dots, T, \\
 \Delta^m \mu_t &= (1 + L)^n \zeta_t, & \zeta_t \sim \text{NID}(0, \sigma_\zeta^2), \\
 \psi_t &\sim \text{NID}(0, \lambda \sigma_\zeta^2), & \text{E}(\zeta_t, \psi_{t-j}) = 0, \forall j,
 \end{aligned} \tag{1}$$

where μ_t is the signal, or trend, component, ψ_t is the noise, Δ is the difference operator, $\Delta = 1 - L$ and L is the lag operator such that $L^j y_t = y_{t-j}$ for integer j .

The (pseudo) autocovariance generating functions (ACGF) of the components and the series are:

$$g_\mu(L) = \frac{|1 + L|^{2n}}{|1 - L|^{2m}} \sigma_\zeta^2, \quad g_\psi(L) = \lambda \sigma_\zeta^2, \quad g_y(L) = g_\mu(L) + g_\psi(L),$$

where $|1 + L|^2 = (1 + L)(1 + L^{-1})$ and $|1 - L|^2 = (1 - L)(1 - L^{-1})$. Assuming a doubly infinite sample, the minimum mean square estimators (MMSE) of the components are respectively $\tilde{\mu}_t = w_\mu(L)y_t$ and $\tilde{\psi}_t = y_t - \tilde{\mu}_t = w_\psi(L)y_t$, where $w_\mu(L) = g_\mu(L)/g_y(L)$ and $w_\psi(L) = g_\psi(L)/g_y(L)$; see Whittle (1983). Hence, the WK filters can be written:

$$w_\mu(L) = \frac{|1 + L|^{2n}}{|1 + L|^{2n} + \lambda|1 - L|^{2m}}, \quad w_\psi(L) = \frac{\lambda|1 - L|^{2m}}{|1 + L|^{2n} + \lambda|1 - L|^{2m}} = 1 - w_\mu(L). \tag{2}$$

The above trend filter can be equivalently derived by solving the following penalised

least square problem:

$$\min_{\mu_t} \text{PLS} = \sum_t [(1 + L)^n (y_t - \mu_t)]^2 + \lambda \sum_t (\Delta^m \mu_t)^2,$$

as can be shown by direct differentiation. Also, after a transformation and with a change of sign, the PLS above coincides with the kernel of the joint Gaussian density of the observations and the trend, when y_t is generated according to (1). The connection with the signal-noise ratio makes clear that the Lagrange multiplier, λ , measures the variability of the noise component relative to that of the trend, and regulates the smoothness of the long-term component.

Using Whittle's result (1983, page 58), the ACGF of the final estimation error, $e_t = \mu_t - \tilde{\mu}_t = -(\psi_t - \tilde{\psi}_t)$, is equal to

$$g_e(L) = \frac{g_\mu(L)g_\psi(L)}{g_y(L)} = \frac{\lambda|1 + L|^{2n}}{|1 + L|^{2n} + \lambda|1 - L|^{2m}}\sigma_\zeta^2$$

The estimators $\tilde{\mu}_t, \tilde{\psi}_t$, are also known as *smoothed* or *final* estimators. From the operational standpoint, given a time series y_t , available at times $t = 1, 2, \dots, T$, the MMSE estimates of the components using information up to and including time $t + l$, denoted $\tilde{\mu}_{t|t+l}$ and $\tilde{\psi}_{t|t+l}$, along with their mean square errors, are computed by the Kalman filter and the associated smoothing algorithms for the model (1), see Harvey (1989). For $l = 0$ the estimators are also known as *filtered* or *real time* estimators. The treatment of initial conditions in the presence of nonstationarity is dealt with in de Jong (1991), Ansley and Kohn (1985) and Koopman (1997), and de Jong (1989) presents various smoothing algorithms; the connection with the WK signal extraction theory is discussed in Burridge and Wallis (1988).

The class of filters depends on the order of integration of the trend (m , which regulates

its flexibility), on the order of the unit root at the Nyquist frequency (n , which *cæteris paribus* regulates the smoothness of $\Delta^m \mu_t$), and λ , which measures the relative variance of the noise component. The filter proposed by Hodrick and Prescott (1997), enjoying large popularity in economics, arises for the combination $m = 2, n = 0, \lambda = 1600$ for quarterly data. Gómez (2001) consider two types of Butterworth filters for which $n = 0$ or $m = n$. Rational square wave trend-cycle filters have been introduced by Pollock (2000) using 5 ideal conditions (phase-neutrality, complementarity, symmetry, high- and lowpass conditions); as Pollock shows, they constitute the optimal filters for the decomposition (1) with the noise replaced by the process $\psi_t = \Delta^{n-m} \kappa_t$; our framework thus encompasses rational square wave filters with $n = m$, which is perhaps the most interesting case, as it postulates a stationary and invertible representation for ψ_t . Finally, the multiresolution Haar scaling and wavelet filters (see Percival and Walden (1999)) occur for $m = n = 1, \lambda = 1$, in which case the trend filter and the cycle filter are both finite impulse response filters: $w_\mu(L) = 0.25L^{-1} + 0.5L + 0.25L^{-1}$, $w_\psi(L) = -0.25L^{-1} + 0.5L - 0.25L^{-1}$.

The trend filter can also be characterised as a lowpass filter whose cutoff frequency depends on the three parameters. Frequency domain arguments can be advocated for designing the parameters so as to select the fluctuations that are in a specified periodicity range.

In particular, let $w_\mu(\omega)$ denote the Fourier transform of the trend filter (2), $w_\mu(\omega) = w_\mu(e^{-i\omega})$, $\omega \in [0, \pi]$; as the latter is real and positive, it is coincident with the gain of the filter. The gain of the trend is a monotonically decreasing with λ , it is unit at the zero frequency and it is zero if n is greater than zero. The trend filter will preserve to a great extent those fluctuations at frequencies for which the gain is greater than 1/2 and reduce to a given extent those for which the gain is below 1/2. This simple argument justifies the definition of a lowpass filter with cutoff frequency ω_c if the gain halves at that frequency;

see Gómez (2001), section 1.

Solving the equation $w_\mu(\omega_c) = 1/2$, the parameter λ is expressed as a function of the cutoff frequency and the orders m and n :

$$\lambda = 2^{n-m} \left[\frac{(1 + \cos \omega_c)^n}{(1 - \cos \omega_c)^m} \right]. \quad (3)$$

In the light of this relationship, $f(m, n, \omega_c)$ will denote the trend filter corresponding to the orders m and n and the cutoff ω_c .

Figure 1 displays the gains of various trend filters for $m, n = 1, 2, 3$ and two cutoff frequency, the first corresponding to $\omega_c = \pi/20$ (a period of 10 years of quarterly observations) and $\pi/2$. The upper panels illustrate that for low cutoff frequency $\omega_c = \pi/20$, the gain is invariant to the values of n , whereas this is not the case for higher frequencies. The lower panels consider the effects of increasing the parameter m , given the others. Sharper gains are obtained with more appreciable differences at $\pi/2$. Notice that $f(1, 1, \pi/2)$ corresponds to the gain of the Haar scaling filter ($\lambda = 1$) and that $f(2, 0, \pi/20)$ is close to the HP filter for quarterly observations, as the smoothness parameter corresponding to this cutoff frequency is $\lambda = 1649$. As m and n increase the gain gets closer to the ideal lowpass filter, with unit gain for $\omega \leq \omega_c$ and zero for $\omega > \omega_c$.

3 Model based interpretation: embedding the trend-cycle decomposition

Economic time series only rarely admit the representation (1); nevertheless, applications of the filters 2 is widespread, as the popularity of the HP filter testifies; the bandpass nature could provide a justification for their use (see Gómez (2001)). However, when

the available series y_t cannot be modelled as (1) it is not immediately clear how trends and cycles should be defined and how inferences on them should be made. In particular, the Kalman filter and the associated smoothing algorithms for the model (1) no longer provide the MMSE estimators of the components nor their MSE.

In this section we propose an embedding strategy that defines artificial trend and cyclical components whose optimal signal extraction filters are provided by (2) when a doubly infinite sample is available, and that rely on the Kalman filter and smoother associated to an appropriately specified state space model for the computation of the MMSE of the components and their MSE in real time. As a result, the optimal filter varies with the properties of the series under investigation.

This approach was initiated by Gómez (2001) and Kaiser and Maravall (2001). In the present section we provide a novel derivation of the model based interpretation of the filters (2) based on the decomposition of the innovation process into ARMA components with noninvertible roots and common stationary AR polynomial.

Let y_t denote a univariate time with ARIMA(p, d, q) representation, that we write

$$\phi(L)(\Delta^d y_t - c) = \theta(L)\xi_t, \quad \xi_t \sim \text{NID}(0, \sigma^2),$$

where c is a constant, $\phi(L) = 1 - \phi_1 L - \dots - \phi_p L^p$ is the AR polynomial with stationary roots, $\Delta = 1 - L$ and $\theta(L) = 1 + \theta_1 L + \dots + \theta_q L^q$ is invertible. The (pseudo) autocovariance generating function of y_t is

$$g_y(L) = \frac{|\theta(L)|^2}{|1 - L|^{2d} |\phi(L)|^2} \sigma^2,$$

where $|\theta(L)| = \theta(L)\theta(L^{-1})$, and $|\phi(L)| = \phi(L)\phi(L^{-1})$.

Let us now introduce the following decomposition of the white noise disturbance ξ_t :

$$\xi_t = \frac{(1+L)^n \zeta_t + (1-L)^m \kappa_t}{\varphi(L)}, \quad (4)$$

where ζ_t and κ_t are two mutually and serially independent Gaussian disturbances, $\zeta_t \sim \text{NID}(0, \sigma^2)$, $\kappa_t \sim \text{NID}(0, \lambda\sigma^2)$, and

$$|\varphi(L)|^2 = \varphi(L)\varphi(L^{-1}) = |1+L|^{2n} + \lambda|1-L|^{2m}. \quad (5)$$

We refer to (5) as the spectral factorisation of the lag polynomial on the right hand side; the existence of the polynomial $\varphi(L) = \varphi_0 + \varphi_1 L + \dots + \varphi_{q^*} L^{q^*}$, of degree $q^* = \max(m, n)$, is guaranteed by the fact that the Fourier transform of the rhs is never zero over the entire frequency range; see Sayed and Kailath (2001).

In the light of (4)-(5), the series can be decomposed into orthogonal trend and cyclical components:

$$\begin{aligned} y_t &= \mu_t + \psi_t, \\ \phi(L)\varphi(L)(\Delta^d \mu_t - c) &= (1+L)^n \theta(L) \zeta_t, \quad \zeta_t \sim \text{NID}(0, \sigma^2) \\ \phi(L)\varphi(L)\psi_t &= \Delta^{m-d} \theta(L) \kappa_t, \quad \kappa_t \sim \text{NID}(0, \lambda\sigma^2) \end{aligned} \quad (6)$$

such that the trend has the same order of integration as the series (regardless of m) and the cycle is stationary provided that $m \geq d$. An interesting case arises for $m = d + n$, for which the trend and the cycle have the same number of unit roots in the MA representation, at the π and zero frequency, respectively.

The decomposition is an artifact, as it does not necessarily correspond to a characteris-

tic of the phenomenon under investigation (if it does the components would be estimated with the minimum MSE among all alternative decompositions); nevertheless, nothing prevents that artificial components are introduced and measured with the intent of selecting some fluctuations of interest.

The ACGFs of the components are respectively:

$$g_\mu(L) = w_\mu(L)g_y(L), \quad g_\psi(L) = w_\psi(L)g_y(L), \quad (7)$$

Obviously, $g_y(L) = g_\mu(L) + g_\psi(L)$. Given the availability of a doubly infinite sample, the optimal signal extraction filters are obtained from the ratio of the ACGFs of the components to that of y_t . Thus, $\tilde{\mu}_t = w_\mu(L)y_t$ and $\tilde{\psi}_t = w_\psi(L)y_t$, with impulse response function given by (2), are the WK estimators of the components.

This simple argument shows that the signal extraction filter for the central data points will continue to be represented by (2), regardless of the properties of y_t , but this is the only feature that is invariant to those properties. The MSE of the smoothed components, as a matter of fact, depends on the ACGF of y_t as will be shown in the next section. Furthermore, the estimators $\tilde{\mu}_{t+l|t}$, $\tilde{\psi}_{t+l|l}$, and the corresponding MSEs will be provided by the Kalman filter and smoother (if $l \leq 0$) associated to the model (6), whose state space representation is presented in the appendix.

4 An upper bound for the estimation error variance

We can apply general principles and in particular Whittle's formula for obtaining the ACGF of the components' estimation error, $e_t = \psi_t - \tilde{\psi}_t = -(\mu_t - \tilde{\mu}_t)$:

$$g_e(L) = \frac{g_\mu(L)g_\psi(L)}{g_y(L)} = w_\mu(L)w_\psi(L)g_y(L) = \frac{\lambda|1 - L|^{2m}|1 + L|^{2n}}{|\varphi(L)|^4}g_y(L). \quad (8)$$

Let us denote $g_x(L)$ the ACGF of the stationary process $x_t = \Delta^d y_t$ and consider the factorisation:

$$g_e(L) = \frac{\lambda|1 - L|^{2(m-d)}|1 + L|^{2n}}{|\varphi(L)|^4}g_x(L) \equiv v(L)g_x(L).$$

Applying the Cauchy-Schwartz inequality, the estimation mean square error has an upper bound that can be broken down as follows:

$$\text{MSE}(e_t) = \frac{1}{\pi} \int_0^\pi g_e(\omega) d\omega \leq \frac{1}{\pi} \left[\int_0^\pi v(\omega)^2 d\omega \right]^{1/2} \left[\int_0^\pi g_x(\omega)^2 d\omega \right]^{1/2}.$$

The last factor depends on the ACGF of the stationary representation of the process and it is invariant to the trend-cycle filter; the first factor; the first, on the other hand, depends solely on the properties of the filter and the true order of integration d .

We now consider how different values of m, n and ω_c affect the uncertainty of the estimated components components, distinguishing three cases, according as to whether the series is stationary ($d = 0$) or integrated up to the second order, $d = 1, 2$. Figure (2) displays the logarithm of $\int v(\omega) d\omega$ versus the cutoff frequency used in the determination of the smoothing parameter according to formula (3).

In the stationary case the components defined using high and low cutoff frequencies

are estimated with lower uncertainty. Moreover, for a given ω_c , the upper bound decreases as m and n increase: this important feature holds also in the nonstationary case; however, the sensitivity to these parameters decreases quite rapidly, as can be seen from the second and third panels as we move from $m = 2$ to $m > 2$. For nonstationary series, $d = 1, 2$, the upper bound decreases monotonically with the cutoff frequency, ω_c . This implies that components defined using a lower cutoff, i.e. the trend preserves the longer periodicities, are estimated with greater uncertainty.

5 Uncertainty and revisions

The previous section discussed how the nature of the filter affects the upper bound of the components MSE. We turn now to two case studies that illustrate how the reliability of the trend-cycle estimates depends on the cutoff frequency and the other parameters that regulate the flexibility and the smoothness of the filter, and the extent of the revision process¹.

5.1 The decomposition of U.S. GDP

Our first illustration deals with the popular HP filter ($m = 2, n = 0, \lambda = 1600$) adapted to the logarithm of the U.S. quarterly real gross domestic product (GDP), available at the time of writing for the sample period 1947.q1-2003.q3. We consider three ARIMA models, with parameter estimates presented below, along with the Akaike and Bayesian information criteria (AIC and BIC, respectively), and the Ljung-Box portmanteau autocorrelation test with 8 lags (p-value in parenthesis): the first is a simple random walk

¹The computations in the paper were performed using the programming language Ox by Doornik (2001), and the library of state space function SsfPack by Koopman, Shephard and Doornik (1999). The ARIMA models were estimated using E-views

with drift, the second is an ARIMA(1,1,0) selected on the grounds of parsimony, and the third is the model adapted by Morley, Nelson and Zivot (2002), which provides a smaller AIC, but greater BIC. The Ljung-Box statistic clearly points out that the RW model is misspecified.

Random walk [AIC = -6.09, BIC = -6.07, $Q(8) = 33.13(0.00)$]

$$\Delta y_t = \underset{(0.0010)}{0.0092} + \xi_t, \quad \xi_t \sim \text{NID}(0, 0.0172^2)$$

ARIMA(1,1,0) [AIC = -6.18, BIC = -6.14, $Q(8) = 8.98(0.25)$]

$$(1 - \underset{(0.0833)}{0.3260} L)\Delta y_t = \underset{(0.0014)}{0.0092} + \xi_t, \quad \xi_t \sim \text{NID}(0, 0.0109^2)$$

ARIMA(2,1,2) [AIC = -6.20, BIC = -6.10, $Q(8) = 2.12(0.71)$]

$$(1 - \underset{(0.0853)}{1.4432} L + \underset{(0.0753)}{0.8527} L^2)\Delta y_t = \underset{(0.0011)}{0.0092} + (1 - \underset{(0.1124)}{1.2240} L + \underset{(0.0938)}{0.6914} L^2)\xi_t$$

$$\xi_t \sim \text{NID}(0, 0.0106^2)$$

We consider now the trend-cycle decomposition with cutoff frequency $\omega_c = 0.16$ corresponding to a period of about 10 years (39.7 quarters) and $\lambda = 1600$; each ARIMA model implies a different representation for the components; estimates of the latter, computed by the Kalman filter and smoother for the corresponding state space model, are displayed in figure 3, where the trend component refers to the ARIMA(2,1,2) model. The HP cyclical component is also displayed in the second and the third panel, which

presents $\tilde{\psi}_{t|T}$ for the last years in the sample, in order to appreciate better the differences among the estimates. The HP cycle is the MMSLE of the cycle under the IMA(2,2) model $\Delta^2 y_t = (1 - 1.7771L + 0.7994L^2)\xi_t$, which would not be selected for the series under investigation, being manifestly misspecified.

As figure 3 shows the HP estimates at the end of the sample differ from those that would be obtained from the model based decompositions. The optimal filter varies with the times series model for y_t ; however, the estimates for the ARIMA(1,1,0) and ARIMA(2,1,2) are indistinguishable, and those for the RW model are quite close. Differences arise with respect to the estimation MSE, plotted in the last panel. That arising from the HP filter is a clear underestimation of the MSE that would arise from models that provide a better representation of the series. We notice also that the latter is quite sensitive to the model selected, being greater for the ARIMA(1,1,0) model.

Leaving the other parameters unchanged ($m = 2$ and $n = 0$), we next consider the model-based filter that arise for the cutoff frequency equal to 1.26, corresponding to a period of 5 quarters and $\lambda = 0.52$. This filter has been adopted by Artis, Marcellino and Proietti (2003) in order to extract a lowpass component reducing the amplitude of those fluctuations with periodicity less than the minimum cycle duration (one year and a quarter), which is employed used to date the peaks and troughs of the business cycle; as a matter of fact, in dating the business cycle, we should abstract from those high frequency movements that cannot qualify as cyclical because they are too short lived. The estimates of the components and the estimation error MSE are presented in figure 4. The evidence, given the properties of the series under investigation, such that the high frequency components of Δy_t have little amplitude, the estimates of the highpass component do not vary with the model selected and are virtually coincident with the HP filter with $\lambda = 0.52$. The difference pertains the MSE, which in turn is very small and close to zero. This reflects

the fact that for an integrated series, components with high cutoff frequency are estimated with greater reliability. The example fosters the conclusions that long run trends are estimated less reliably than short run ones. We notice in closing that both cyclical components are affected by a change in volatility, as documented in Stock and Watson (2002).

5.2 The ARIMA(1,1,0) case

The extent of the revision process is assessed comparing the final estimation error MSE with the real time one. This ratio is important for characterising the magnitude of the revision process as future observations become available, as it is shown below.

Since the MSE of the filtered and smoothed estimates is identical for the two components (recalling $\mu_t - \tilde{\mu}_{t|t} = \tilde{\psi}_{t|t} - \psi_t$), let us concentrate on the cycle and denote $\text{MSE}(\tilde{\psi}_{t|t}) = \mathbf{E}[(\psi_t - \tilde{\psi}_{t|t})^2]$, $\text{MSE}(\tilde{\psi}_t) = \mathbf{E}[(\psi_t - \tilde{\psi}_t)^2]$. The MSE of the filtered estimates admits the following decomposition:

$$\text{MSE}(\tilde{\psi}_{t|t}) = \text{MSE}(\tilde{\psi}_t) + V_R,$$

where $V_R = \mathbf{E}[(\tilde{\psi}_t - \tilde{\psi}_{t|t})^2]$ is the variance of the revisions. Thus the ratio

$$R = \frac{\text{MSE}(\tilde{\psi}_{t|t})}{\text{MSE}(\tilde{\psi}_t)} = 1 + \frac{V_R}{\text{MSE}(\tilde{\psi}_t)}, \quad (9)$$

measures the relative importance of the revision process; the larger the the more the ratio moves away from unity, which is the reference value that would be achieved were the components fully estimated in real time.

The ratio (9) clearly depends on the ARIMA model for y_t ; in this paper we limit the

analysis of its behaviour to the fairly realistic and representative case when the series follows the ARIMA(1,1,0) process $(1 - \phi L)(\Delta y_t - c) = \xi_t$. If ϕ is positive, the dynamics of Δy_t are dominated by low frequency components and viceversa; we mention that similar results hold for the ARIMA(0,1,1) process. Figure 5 displays the values of R against the cutoff frequency for different combinations of the pair m, n . Both the numerator and the denominator are evaluated by the steady state Kalman filter and smoother for the state space representation given in the appendix.

The main evidence can be summarised as follows:

- For given ϕ, ω_c and $m = 1, 2$, increasing n enhances the relative magnitude of the revision process.
- In the RW case ($\phi = 0$) we have the interesting property that choosing $n = m$ makes R invariant to the cutoff frequency.
- For given ϕ, ω_c and n , the magnitude of the revision process increases with m . Hence the choice of a more flexible trend results *cæteris paribus* in larger revisions.
- Long run trends (ω_c is low) are subject to comparatively smaller revisions if Δy_t is dominated by low frequencies fluctuations ($\phi > 0$).

6 Conclusions

The paper has focussed on a class of trend-cycle filters that is optimal for a particular local trend model and that depends on the order of integration of the trend (trend flexibility), on the order of the unit root at the Nyquist frequency (trend smoothness), and the relative variance of the cyclical component. The trend filter can be characterised as a lowpass filter whose cutoff frequency depends on the three parameters.

By embedding the trend-cycle decomposition within the ARIMA time series model for a univariate time series we provide a model-based framework that enables inferences on the unobserved components to be conducted by means of the Kalman filter and smoother associated to the relevant state space model. The components select certain features of the series and in this respect they represent an artificial construct. Nevertheless, the model-based interpretation allows to formalise the discussion on the relevant issue of the uncertainty by which certain signals or components are estimated.

In particular, it has been shown that for the type of nonstationary time series usually encountered in macroeconomics more stable components and long run trends and cycles, characterised by lower cutoff frequencies are estimated with less reliability. Moreover, in designing the filter the analyst faces some trade-offs. The purpose of this paper was that of illustrating them. In particular, if on the hand increasing the flexibility of the trend and its smoothness via the introduction on noninvertible for a given cutoff frequency enhances the reliability of the signal, in real time the signal will be subject with large revisions. Hence, smoothness priors have their costs, as the quest for smooth signals results in larger revisions and greater discrepancies between real time and final estimates. Intuitively, the smoother the trend, the more observations are needed before the estimate settles down to its final value, and thus the greater the amount of revision involved.

References

- Ansley, C.F., Kohn, R., 1985. Estimation, filtering and smoothing in state space models with incompletely specified initial conditions. *The Annals of Statistics*, 13, 1286-1316.
- Artis, M., Marcellino, M., Proietti, T., 2003. Dating the Euro Area Business Cycle. *Centre for Economic Policy Research*, Discussion paper n. 3696.
- Burridge, P., Wallis, K.F., 1988. Prediction theory for autoregressive-moving average processes. *Econometric Reviews*, 7, 65–9.
- de Jong, P., 1989. Smoothing and interpolation with the state space model. *Journal of the American Statistical Association*, 84, 1085–1088.
- de Jong, P., 1991. The diffuse Kalman filter. *Annals of Statistics* 19, 1073–83.
- Doornik, J.A., 2001. *Ox. An Object-Oriented Matrix Programming Language*. Timberlake Consultants Press, London.
- Gómez, V., 2001, The Use of Butterworth Filters for Trend and Cycle Estimation in Economic Time Series. *Journal of Business and Economic Statistics*, 19, 365–373.
- Hodrick J.R. and Prescott E.C., 1997. Postwar U.S. business cycles: an empirical investigation, *Journal of Money, Credit and Banking*, 29, 1–16.
- Harvey, A.C., 1989. *Forecasting, Structural Time Series and the Kalman Filter*, Cambridge University Press, Cambridge.
- Harvey, A.C., and Jäger, A., 1993. Detrending, stylized facts and the business cycle. *Journal of Applied Econometrics*, 8, 231-247.

- Kaiser, R. Maravall A., 2001. *Measuring Business Cycles in Economic Time Series*, Lecture Notes in Statistics, 154, Springer-Verlag, New York.
- Koopman, S.J., 1997. Exact initial Kalman filter and smoother for non-stationary time series models. *Journal of the American Statistical Association* 92: 1630–8.
- Koopman, S.J., Shephard, N., Doornik, J.A., 1999. Statistical algorithms for models in state space using SsfPack 2.2. *Econometrics Journal*, 2, 113–66.
- Morley, J.C., Nelson, C.R., Zivot, E., 2002. Why are Beveridge-Nelson and Unobserved-Component Decompositions of GDP So Different? *Review of Economics and Statistics*, 85, 235–243.
- Orphanides, A., van Norden, S., 2002. The Unreliability of Output Gap Estimates in Real Time. *Review of Economics and Statistics*, 84, 569–583.
- Percival, D.B., and Walden, A.T., 1999. *Wavelet Methods for Time Series Analysis*, Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press.
- Pollock, D.S.G., 2000. Trend Estimation and De-trending via Rational Square-wave Filters. *Journal of Econometrics*, 99, 317–334.
- Sayed, A.H., and Kailath, T., 2001. A Survey of Spectral Factorization Methods. *Numerical Linear Algebra with Applications*, 8, 467–496.
- Stock J.H., and Watson, M., 2002. Has the business cycle Changed and Why? *NBER Working Paper* No. w9127.
- Whittle P., 1983. *Prediction and Regulation by Linear Least Squares Methods*, Second edition. Basil Blackwell, Oxford.

A State space representation

The spectral factorisation $\varphi(L)\varphi(L^{-1}) = a_0 + a_1(L + L^{-1}) + \dots + a_k(L^k + L^{-k})$, $k = \max(m, n)$ with

$$a_j = I(j \leq n) \binom{2n}{n+j} + I(j \leq m) \binom{2m}{m+j},$$

where $I(\cdot)$ is the indicator function, taking value 1 if the argument is true and 0 otherwise, can be achieved via the Riccati equation method presented in Sayed and Kailath (2001).

Let now $\phi(L)^*$ denote the AR polynomial $\phi(L)^* = \phi(L)\varphi_0^{-1}\varphi(L) = 1 - \phi_1^*L - \dots - \phi_{p+q^*}L^{p+q^*}$, common to the representation of the trend and the cycle in (6), and let $\theta_\mu(L) = (1 + L)^n\theta(L)$, the MA polynomial of the trend component. Further, define the orders $q_\mu = \max(p + q^*, n + q + 1)$, $q_\psi = \max(p + q^*, q + 1)$.

The state space representation of 6 consists of the measurement equation $y_t = \mathbf{z}'\boldsymbol{\alpha}_t$, with

$$\mathbf{z}' = [\mathbf{z}'_\mu, \mathbf{z}'_\psi], \quad \boldsymbol{\alpha}_t = [\boldsymbol{\alpha}'_{\mu,t}, \boldsymbol{\alpha}'_{\psi,t}]', \quad \mathbf{z}'_\mu = [\mathbf{i}'_{d+1}, \mathbf{0}'_{q_\mu}], \quad \mathbf{z}'_\psi = [1, \mathbf{0}'_{q_\psi}];$$

and the transition equation $\boldsymbol{\alpha}_{t+1} = \mathbf{T}\boldsymbol{\alpha}_t + \mathbf{c} + \mathbf{R}\boldsymbol{\epsilon}_t$ with $\boldsymbol{\epsilon}_t = \varphi_0^{-1/2}[\eta_t, \kappa_t]'$,

$$\mathbf{T} = \text{diag}(\mathbf{T}_\mu, \mathbf{T}_\psi), \quad \mathbf{c} = [c\mathbf{z}'_\mu, \mathbf{0}]', \quad \mathbf{R} = \text{diag}(\boldsymbol{\theta}_\mu, \boldsymbol{\theta}_\psi),$$

$$\mathbf{T}_\mu = \begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{0} & \mathbf{T}_\mu^* \end{bmatrix},$$

where \mathbf{A} is a $d \times d$ upper triangular matrix with elements $a_{ij} = 1, j \geq i$, and $a_{ij} = 0$, \mathbf{B} is $d \times q_\mu$ matrix with zero elements except for the first column, which contains unit

elements, $\mathbf{B} = [\mathbf{i}_d, \mathbf{0}]$, and

$$\mathbf{T}_\mu^* = \begin{bmatrix} \boldsymbol{\phi}_\mu^* & \mathbf{I}_{q_\mu-1} \\ & \mathbf{0}' \end{bmatrix}, \mathbf{T}_\psi = \begin{bmatrix} \boldsymbol{\phi}_\psi^* & \mathbf{I}_{q_\psi-1} \\ & \mathbf{0}' \end{bmatrix},$$

$\boldsymbol{\theta}_\mu = [1, \theta_{\mu,1}, \theta_{\mu,2}, \dots, \theta_{\mu,q_\mu}]'$, $\boldsymbol{\theta}_\psi = [1, \theta_{\psi,1}, \theta_{\psi,2}, \dots, \theta_{\psi,q_\psi}]'$, $\boldsymbol{\phi}_\mu^* = [\phi_{\mu,1}^*, \phi_{\mu,2}^*, \dots, \phi_{\mu,q_\mu}^*]'$, $\boldsymbol{\phi}_\psi^* = [\phi_{\psi,1}^*, \phi_{\psi,2}^*, \dots, \phi_{\psi,q_\psi}^*]'$. The nonzero AR coefficients in $\boldsymbol{\phi}_\mu^*$ and $\boldsymbol{\phi}_\psi^*$ are the coefficients of $\phi(L)^*$, the nonzero MA coefficients in $\boldsymbol{\theta}_\mu$ are those of the polynomial $\theta_\mu(L)$, and finally, those of $\boldsymbol{\theta}_\psi$ are the coefficients of $\theta(L)$.




Figure 1: Gain function of various trend filters $f(m, n, \omega_c)$.

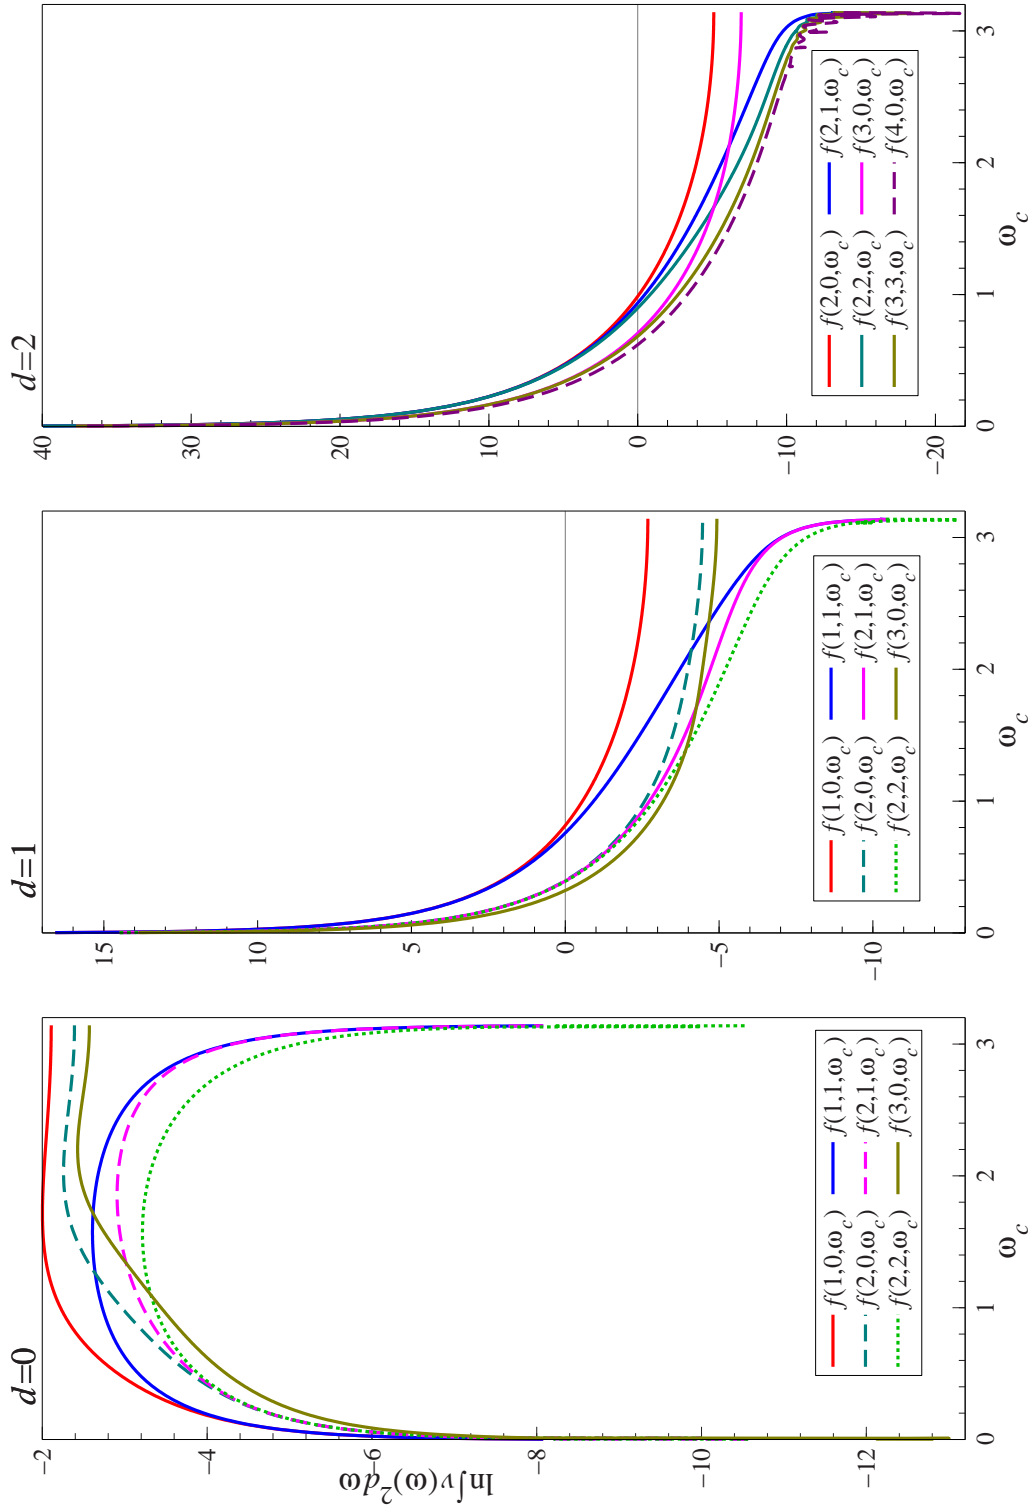


Figure 2: Plot of $\ln \int_0^\pi v(\omega)^2 d\omega$ versus the cutoff frequency ω_c .

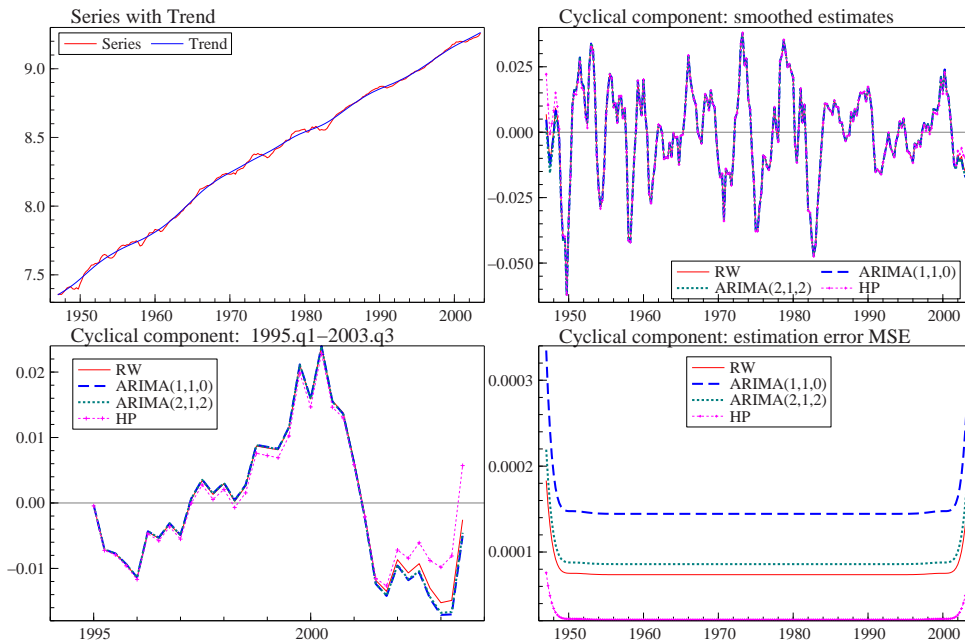


Figure 3: U.S. Gross Domestic Product, 1947.q1 - 2003.q3. Trend cycle decomposition with $m = 2$, $n = 2$ and $2\pi/\omega_c = 39.7$ ($\lambda = 1600$) corresponding to the HP filter

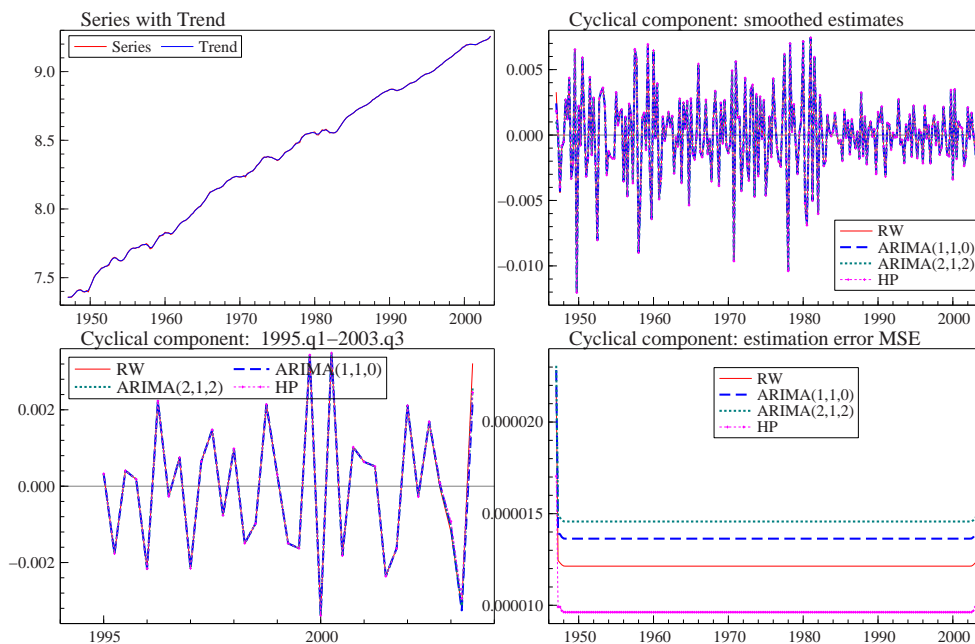


Figure 4: U.S. Gross Domestic Product, 1947.q1 - 2003.q3. Trend cycle decomposition with $m = 2$, $n = 2$ and $2\pi/\omega_c = 5$ quarters ($\lambda = 0.52$)

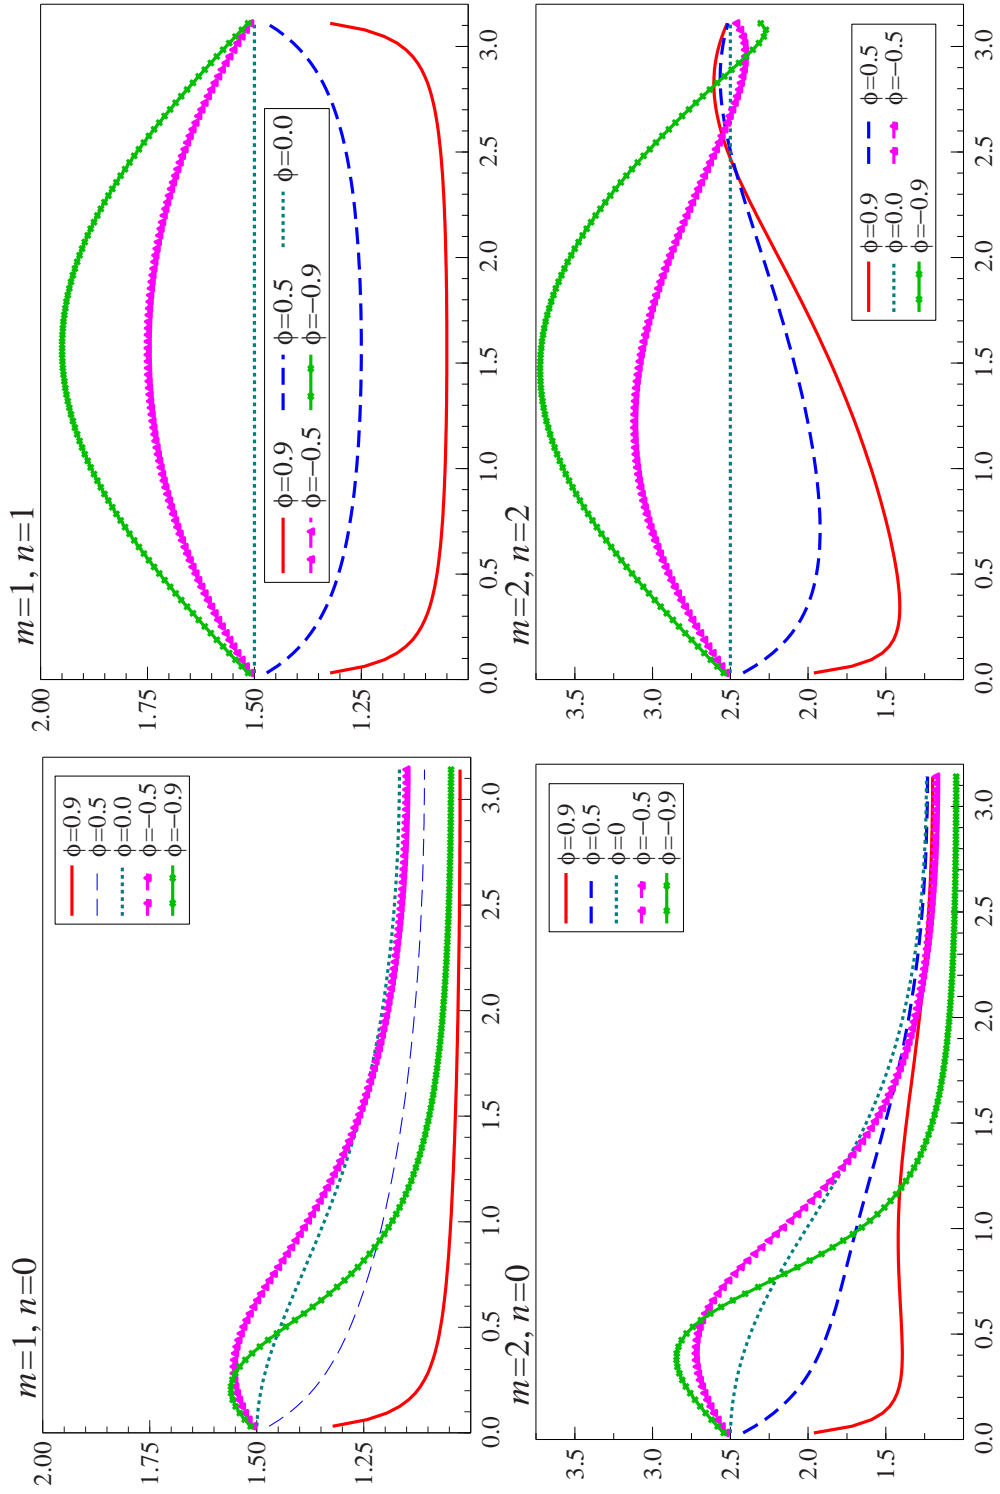


Figure 5: MSE ratio of filtered and smoothed estimates for the ARIMA(1,1,0) model $\Delta y_t = \phi \Delta y_{t-1} + x_t i_t$ versus cutoff frequency ω_c for different values of m, n .