



EUROPEAN CENTRAL BANK

WORKING PAPER SERIES

NO. 499 / JULY 2005

**A TREND-CYCLE(-SEASON)
FILTER**

by Matthias Mohr



EUROPEAN CENTRAL BANK



WORKING PAPER SERIES

NO. 499 / JULY 2005

A TREND-CYCLE(-SEASON) FILTER ¹

by Matthias Mohr ²

In 2005 all ECB publications will feature a motif taken from the €50 banknote.

This paper can be downloaded without charge from <http://www.ecb.int> or from the Social Science Research Network electronic library at http://ssrn.com/abstract_id=726694.

¹ I want to thank Andrew C. Harvey, Karl-Heinz Tödter, Jean-Pierre Vidal and an anonymous referee for helpful comments. The views expressed in this paper do not necessarily reflect those of the European Central Bank.

² European Central Bank, Kaiserstraße 29, D-60311 Frankfurt am Main, Germany; e-mail: matthias.mohr@ecb.int

© European Central Bank, 2005

Address

Kaiserstrasse 29
60311 Frankfurt am Main, Germany

Postal address

Postfach 16 03 19
60066 Frankfurt am Main, Germany

Telephone

+49 69 1344 0

Internet

<http://www.ecb.int>

Fax

+49 69 1344 6000

Telex

411 144 ecb d

All rights reserved.

Reproduction for educational and non-commercial purposes is permitted provided that the source is acknowledged.

The views expressed in this paper do not necessarily reflect those of the European Central Bank.

The statement of purpose for the ECB Working Paper Series is available from the ECB website, <http://www.ecb.int>.

ISSN 1561-0810 (print)
ISSN 1725-2806 (online)

CONTENTS

Abstract	4
Non-technical summary	5
1 Introduction	6
2 The HP filter	7
2.1 The stochastic model of the HP Filter	7
2.2 The value of λ	9
3 The TC filter	10
3.1 A general stochastic trend model	10
3.2 A stochastic model for the cycle	11
3.3 Putting it all together: the TC filter	13
3.4 Properties of the TC filter in the time domain	15
3.5 Properties of the TC filter in the frequency domain	16
4 An application to real GDP in selected countries	18
4.1 The endpoint problem and the predictive capabilities of filters	20
4.2 The forecasting performance of the HP and the TC filter	22
4.3 The real time reliability of the TC filter	25
5 Testing the stochastic cycle model of the TC filter and adjusting it to the data	29
6 The Trend-Cycle-Season filter	32
6.1 Integrating seasonal components in the TC filter: general remarks	32
6.2 A stochastic model for the seasonal component	32
7 Conclusion	34
Annexes and figures	36
References	43
European Central Bank working paper series	45

Abstract

This paper proposes a new univariate method to decompose a time series into a trend, a cyclical and a seasonal component: the *Trend-Cycle* filter (*TC* filter) and its extension, the *Trend-Cycle-Season* filter (*TCS* filter). They can be regarded as extensions of the Hodrick-Prescott filter (HP filter). In particular, the stochastic model of the HP filter is extended by explicit models for the cyclical and the seasonal component. The introduction of a stochastic cycle improves the filter in three respects: first, trend and cyclical components are more consistent with the underlying theoretical model of the filter. Second, the end-of-sample reliability of the trend estimates and the cyclical component is improved compared to the HP filter since the pro-cyclical bias in end-of-sample trend estimates is virtually removed. Finally, structural breaks in the original time series can be easily accounted for.

Keywords: economic cycles, time series, filtering, trend-cycle decomposition, seasonality.

JEL Classification: C13, C22, E32.

Non-technical summary

Univariate trend-cycle decompositions suffer from all-too simple implicit models of the data generating process, while more elaborated approaches—as for instance unobserved components models—are not always easily applicable. This paper develops an intermediate approach by generalising the HP filter and incorporating a cyclical component into the model representation of the filter in the time domain. The resulting trend-cycle (TC) filter has better end-of-sample properties than the HP filter or the related Extended Exponential Smoothing (EES) procedure. In particular, the pro-cyclicality in end-of-sample trend/cycle estimations, characterising one-component filters such as the HP filter and the EES, is virtually removed.

The incorporation of a cycle model turns out crucial for the favourable properties of the TC filter. Furthermore, structural breaks or exogenous variables to identify the trend and the cyclical component can be easily incorporated in the TC filter. Finally, the Trend-Cycle filter can be expanded towards a Trend-Cycle-Season (TCS) filter in a straightforward way. With the TCS filter, a trend, cyclical and seasonal component can be simultaneously extracted from a time series.

Basic assumptions of the stochastic model underlying the TC filter can be tested, and the model can to some extent be adjusted to the data. As a consequence of the more flexible model-structure, the TC filter produces results, which are more model-consistent than those obtained with the one-component counterparts, the EES and the HP filter.

1 Introduction

The decomposition of macroeconomic time series into trend and cyclical components is crucial to many macroeconomic concepts such as potential output, p-star, or the natural interest rate, which imply that short- and long-term movements can be separated. Typically, the components are theoretical concepts and therefore not observable. Rather, they have to be identified on the basis of a theoretical model or plausible ad-hoc assumptions.

Several tools for trend extraction have been developed in the literature¹. Some of them allow building multivariate economic models and adjusting the model parameters to the data such as models with *unobserved components* (UC), others are purely mechanical transformations of the original data such as the Baxter-King filter (Baxter and King 1999) and the Hodrick-Prescott filter (Hodrick and Prescott 1997). From a theoretical perspective, complex unobserved components models are clearly superior to the simpler methods. From a more practical point of view, the estimation of unobserved component models—which is usually carried out using recursive estimation methods such as the Kalman filter—can be difficult: The results depend on well specified initial conditions for unobserved variables and their variances. The final model chosen is usually the outcome of a relatively elaborate procedure of model selection². Furthermore, in many cases the Kalman filter approach does not work with annual data.

While simple trend extraction methods are more convenient to use, the economic interpretation of their results may pose problems. This is mainly because it is not possible to adjust the filter to properties of the time series to be filtered. Such mechanical approaches may also give rise to “spurious cycles” (Harvey and Jäger 1993; Jäger 1994; Cogley and Nason 1995) which reflect more the properties of the filter used rather than those of the time series. An additional problem, which all approaches—including UC models—have in common, concerns the instability of trend estimations at the end of the data sample. The trend values of the last sample periods can change significantly when the sample is extended with the arrival of new data.³

This paper follows an approach between the two polar methods of trend extraction—UC models on the one hand and mechanical filters on the other. Our filter can be interpreted as an extension of the well known Hodrick-Prescott filter (HP filter). It is based on explicit stochastic models for both the trend and the cycle—hence the name “trend-cycle filter” (TC filter)—allowing the simultaneous extraction of the trend and the cyclical process.

Compared with other common univariate filters, the TC filter has several advantages:

¹ Comprehensive overviews over trend-cycle decompositions are given in Dupasquier et al. (1997) or in Chagny and Döpke (2002).

² As Planas and Rossi (2004, 130) note in an investigation of the real time reliability of UC Phillips curve models:

“...recursive estimation requires a close monitoring of the parameter values, as sudden jumps can strongly increase the revisions. For instance, we found that the proper handling of the Kalman filter starting conditions is critical to the stability of model parameter estimates over time.”

³ The trend also changes if past data are revised ex post. Empirically, the instability due to the revision of past data is less problematic than the instability stemming from new data (Döpke 2004; Rünstler 2002).

first, it has better real time properties than other common univariate filters, as for instance the HP filter. Second, as both, trend and cyclical component, are explicitly modelled, it has a better foundation in the time domain than common univariate filters. Third, it can to some extent be adjusted to the data. Fourth, it can be easily extended to incorporate structural breaks. Finally, it is more convenient to use than unobserved components models.

The paper proceeds as follows. Section 2 discusses general properties of the HP filter. In section 3, the trend cycle filter is developed by generalising the underlying trend model of the HP filter and by amending it with an explicit stochastic model for the cycle. Section 4 discusses the instability of trend/cycle estimations at the end of the sample—the so-called end-point problem of filters. Second, it assesses the endpoint reliability of the TC filter empirically by applying it to real GDP in selected countries and the euro area. Section 5 presents some tests of the TC filter and shows how it can be adjusted to the data. In section 6, the stochastic model of the filter is extended with a seasonal component. The resulting *Trend-Cycle-Season* filter (TCS filter) can be applied to time series which contain seasonal patterns. Section 7 concludes.

2 The HP Filter

The HP filter is obtained by minimising the objective function

$$\sum_{t=1}^N (x_t - x_t^T)^2 + \lambda \sum_{t=2}^N ((x_t^T - x_{t-1}^T) - (x_{t-1}^T - x_{t-2}^T))^2 \quad (1)$$

for x_t^T . It is convenient to express the objective function in matrix form:

$$(X - X^T)'(X - X^T) + \lambda X^{T'} \nabla^{2'} \nabla^2 X^T \quad (2)$$

where X and X^T are $T \times 1$ vectors of the original data and the trend and ∇^2 denotes the 2nd difference matrix⁴. The solution⁵ of this optimisation problem follows from the first order conditions in matrix form:

$$\begin{aligned} X^T &= (I + \lambda \nabla^{2'} \nabla^2)^{-1} X \\ X^C &= X - X^T \end{aligned} \quad (3)$$

2.1 The stochastic model of the HP Filter

For a more general interpretation of the HP filter one may start with the implicit stochastic trend model, a second order random walk. Let us write the model in matrix notation:

$$\begin{aligned} X - X^T - X^C &= 0 \\ \nabla^2 X^T &= \eta, & E(\eta_t) &= 0 & E(\eta_t^2) &= \sigma_\eta^2 & \forall t = 1 \dots N, & E(\eta\eta') &= \sigma_\eta^2 I_N \\ X^C &= \zeta, & E(\zeta_t) &= 0 & E(\zeta_t^2) &= \sigma_\zeta^2 & \forall t = 1 \dots N, & E(\zeta\zeta') &= \sigma_\zeta^2 I_N \\ & & E(\eta\zeta') &= \mathbf{0}_N. \end{aligned} \quad (4)$$

⁴ Lag- and difference operators in matrix form are explained in Appendix A.

⁵ For a more detailed derivation of the solution see for instance Danthine and Girardin (1989).

The residuals η and ζ are typically referred to as signal and noise. We assume that these processes have a zero-mean and that their variances exist. Furthermore, they are assumed to be mutually uncorrelated. The signal variable η is a white noise error term, whereas ζ may follow an unspecified stationary ARMA-process.

When inspecting the stochastic model of the filter and the definition of the trend in equation (4), several points are worth mentioning. First, the objective function in equation (2) is a weighted sum of the inner products of the residuals $\zeta'\zeta + \lambda\eta'\eta$ with the weight parameter λ .

Second, the stochastic model of the trend process as a second order random walk is a prior which may or may not be appropriate, depending on the properties of the series being filtered⁶.

Third, equation (3) implies that the cycle is proportional to the fourth difference of the trend, shifted forwards by two periods. To show this, observe that $X - X^T \equiv X^C = \lambda\nabla^{2'}\nabla^2 X^T = \lambda L^{2'}\nabla^2\nabla^2 X^T$ can be derived⁷ from equation (3). For data points in the middle of the sample, $2 < t < N - 2$, this is identical⁸ to $X^C = \lambda L^{2'}\nabla^4 X^T$. This property is highly implausible—as stated by Reeves et al. (1996, 4)

Fourth, the trend and the cycle add up to the original series, meaning that there is no residual component capturing non-cyclical random impacts. According to the time domain representation of the filter in equation (4), the cycle is not explicitly modelled. Rather, it is defined as a residual process so that an additional residual component cannot be identified.

Finally, under the additional assumptions that the cycle process ζ is white noise and that η and ζ are distributed normally, maximising $\zeta'\zeta + \lambda\eta'\eta$ gives an optimal filter for the underlying stochastic process⁹ if the parameter λ is set equal to the inverse signal-to-noise variance ratio: $\lambda = \sigma_\zeta^2/\sigma_\eta^2$. This interpretation is also consistent with an unobserved components model in which the parameter λ would be estimated as the inverse signal-to-noise variance ratio. These additional assumptions are usually not met in practice. In addition, the choice of the value of λ is based on prior assumptions and not on the concept of an optimal filter. Therefore, the HP filter is in general not an optimal filter in practical applications¹⁰. Furthermore, the cyclical component obtained from filtering is not a white noise process but follows some auto-correlated process, the properties of which depend on λ .

⁶ Many macroeconomic time series are assumed to be I(1) which contradicts the local linear trend model underlying the HP filter.

⁷ See *Property 5* of lag- and difference-matrices in Appendix A.

⁸ The first four rows of the 4th difference matrix are zeroes as $\nabla^4 = I_4\nabla^2\nabla^2$. The lead operator $L^{2'}$ shifts all rows by two rows upwards. Hence, $L^{2'}\nabla^2\nabla^2 = L^{2'}\nabla^4$ holds for $2 < t < N - 2$.

⁹Whittle (1983). A filter is optimal if the sum of squared differences between the true and the estimated cyclical component take a minimum.

¹⁰ It follows also that the fixed value of λ is unequal to the *observed* inverse signal-to-noise variance ratio $\frac{(X-X^T)'(X-X^T)}{N-1} / \frac{X^{T'}\nabla^{2'}\nabla^2 X^T}{N-3}$.

2.2 The value of λ

Since the parameter λ is key for the properties of the HP filter, much has been written about the proper value without, however, providing clear indications as to how to choose the appropriate value of λ . Ideally, the choice of λ should be adjusted so that it reflects prior knowledge on the length of the cycle. However, the smoothing parameter does not only affect the cycle but the volatility of trend growth as well—a consequence of the fact that the HP filter does not contain an explicit model of the cycle. Therefore, many practitioners tend to choose high values for λ when filtering annual data because they feel that lower values—as suggested in the econometrics literature—would give rise to implausibly volatile trend growth rates. Thus, the value of λ is often based on a prior assumption of an acceptable trend volatility.

Values of 1600 for quarterly data and of 100 for annual data are commonly used. Ravn and Uhlig (2002) argue on the basis of frequency domain considerations that $\lambda = 1600$ for quarterly data is inconsistent with $\lambda = 100$ but would rather correspond to $\lambda = 6.5$ for annual data. Kaiser and Maravall (1999) propose a value of 8 for annual data, and Pedersen (2001) argues for a value of 1000 for quarterly data and for 3–5 for annual data. In Bouthevillain et al. (2001) the filter is applied with $\lambda = 30$ and in Mohr (2001) with $\lambda = 20$ to annual data.

The impact of the value of λ can be best demonstrated in the frequency domain. As the gain functions of the trend and the cyclical component for different λ -values in

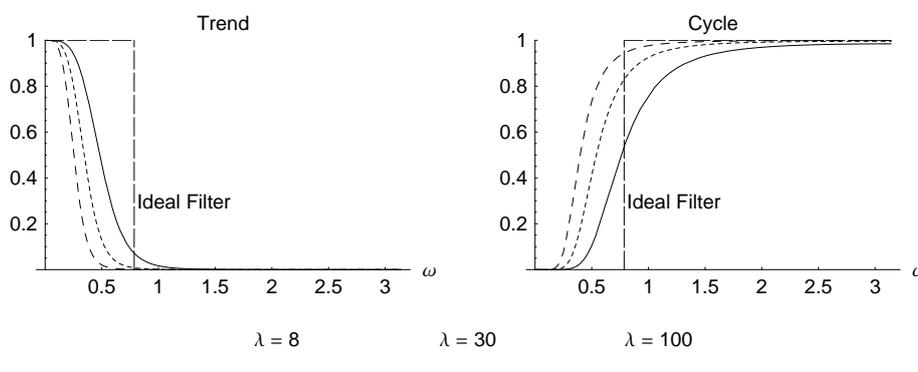


Figure 1: Gain function of the trend and the cyclical component of the HP filter for different values for λ

Figure 1 show, low frequency components are allocated to the trend while high frequency components are allocated to the cycle. Higher values of λ shift the gain function of the trend to the left so that the trend contains less of the higher frequencies, thereby becoming smoother. If $\lambda \rightarrow \infty$, the extracted trend approaches a linear trend. With lower values of the smoothing parameter, the trend becomes more volatile as it contains a larger part of the high-frequency spectrum. In the extreme case of $\lambda = 0$, the trend is equal to the original series¹¹.

¹¹ It is possible, to translate the value of lambda into a corresponding critical frequency ω_c , determined by $\omega_c : G_{HP}^{-1}(\lambda, \cdot) = 0.5$. In this way, the filter can be characterised by a reference cycle of frequency ω_c .

The frequency domain characteristics of the HP filter have well-known implications: First, the volatility of the cycle is controlled by the smoothing parameter λ . However, as λ defines the trend-volatility as well, there is no way to model the trend and the cycle independently from each other. Extracting shorter cycles comes automatically at the cost of a more volatile trend.

Second, the missing model for the cyclical component has important consequences when additional, new data at the end of the sample are processed. There is no other choice than to allocate the information contained in a new data either to the trend or to the cycle, even though it may represent an outlier not generated by the data generating process underlying the HP filter.

Finally, the HP filter is often used as an approximation to an ideal filter. Suppose, for instance, that the objective is to filter out a cycle of 8 or less periods length implying an ideal filter as shown in Figure 1: all frequencies below the critical frequency of $\frac{2\pi}{8}$ are cut off. By adjusting λ , the HP filter can approximate the desired ideal filter to some extent. However, there is a trade off in the choice of λ : while decreasing λ gives a better approximation to the ideal filter in the low frequency range, it worsens the approximation in the higher range. Therefore, either the trend contains frequencies which ideally should be fully captured in the cycle and is therefore overly volatile, or longer waves which—according to the ideal filter—belong to the trend have too much weight in the cycle.

In short, a third component capturing irregular random influences is missing in the HP filter model. This tends to increase the instability of the trend estimate in real time as random influences are partly forced to contribute to the trend variability. This issue will be discussed further in section 4.1.

3 The TC filter

This section extends the HP filter first by allowing for stochastic trends of arbitrary order and second by adding a stochastic model for the cycle to the filter. The resulting trend-cycle filter provides simultaneous, model-based estimates of the trend and the cyclical component.

3.1 A general stochastic trend model

In the HP filter model, the stochastic trend is restricted to a second order random walk. We generalise the trend model to a stochastic trend of any order. In this way, the order of the stochastic trend can be adjusted to the original series. For instance, many economic time series are $I(1)$ and a first order stochastic trend—possibly with a deterministic drift—would be more appropriate than the second order trend embodied in the HP filter.

The generalised trend model in matrix form can be described as

$$\nabla^{d-1}(\nabla X^T - Ub) = \eta \quad (5)$$

where U denotes the $(T \times 1)$ vector $[0, 1, \dots, 1]'$, b stands for the drift parameter to be determined endogenously, and d denotes the order of the trend. The expression Ub

accounts for a deterministic drift if the trend is of first order ($d = 1$). For a higher order trend ($d > 1$), the drift term vanishes as $\nabla^{d-1}Ub = 0$.

Replacing the second line in equation (4) by equation (5) leads to the following objective function of the generalised trend filter in matrix form:

$$(X - X^T)'(X - X^T) + \lambda(\nabla^{d-1}(\nabla X^T - Ub))'(\nabla^{d-1}(\nabla X^T - Ub)). \quad (6)$$

In the case of the first order random walk with drift the objective function has to be maximised for both the trend vector X^T and the drift parameter b , yielding

$$\begin{aligned} X - (I + \lambda\nabla'\nabla)X^T + \nabla'Ub &= 0 \\ b &= (U'U)^{-1}U'\nabla X^T. \end{aligned} \quad (7)$$

Thus, the drift term is computed as the average change in the trend: $b = \frac{x_N^T - x_1^T}{N-1}$. Note, however, that the drift term b and the trend X^T are determined simultaneously. Merging the solutions for b and X^T yields

$$X^T = \left(I + \lambda\nabla' \left(I - U(U'U)^{-1}U' \right) \nabla \right)^{-1} X, \quad d = 1. \quad (8)$$

The solution for $d > 1$ is straightforward, as in this case the trend reduces to a d -th order stochastic trend $\nabla^d X^T = \eta$, and the solution is similar to that of the original HP filter in equation (3):

$$X^T = \left(I + \lambda\nabla^{d'}\nabla^d \right)^{-1} X, \quad d \geq 2. \quad (9)$$

The solution in equation (9) can also be applied to a first order random walk with drift if the linear trend is removed from the time series before filtering. The result should not differ too much from the trend as given in equation (8), in which the deterministic and the stochastic trend components are simultaneously determined.

The generalisation of the trend order is well-known in the literature. The case of $d = 1$ without simultaneous determination of the deterministic drift is known as *exponential smoothing* and was used by Lucas (1980) in an empirical analysis of the quantity theory of money. The simultaneous determination of the drift was first proposed in Tödter (2002) as the *Extended Exponential Smoothing (EES)*. Furthermore, the *Butterworth filter*, which is primarily known in the engineering literature, depicts the general case of a stochastic trend of order d (Gomez 2001).

For macroeconomic time series, stochastic trends of order higher than two do not make much sense. In the following sections, we will therefore concentrate on the EES, the HP filter and on TC filters with first- and second-order stochastic trends.

3.2 A stochastic model for the cycle

In this subsection, the stochastic model for the HP filter is extended by an explicit model for the cycle. The cyclical process is now assumed to follow a stationary ARMA-process,



which is not left implicit as in the HP filter. Thus, we amend the stochastic model in equation (7) by the equation $AX^C = B\zeta$, in which the elements of the matrices A and B are determined by the parameters of an appropriately specified stationary ARMA process.

A convenient approach to model cyclical movements are stochastic cycles as suggested in Harvey (1989) or in Harvey and Jäger (1993). The original stochastic cycle approach in Harvey (1989) was extended towards stochastic cycles of order c in Harvey and Trimbur (2003). A stochastic cycle of order 2 is a stochastic cycle of order 1 with an error process that itself follows a stochastic cycle. Stochastic cycles of higher order are defined respectively. Stochastic cycles of order c give rise to ARMA($2c, c$) processes as shown in Harvey and Trimbur (2003).

The model for the c -th order stochastic cycle can be specified in state-space form as

$$\begin{aligned} \begin{bmatrix} x_{1,t}^C \\ x_{1,t}^{C*} \end{bmatrix} &= \rho \begin{bmatrix} \cos(\mu) & \sin(\mu) \\ -\sin(\mu) & \cos(\mu) \end{bmatrix} \begin{bmatrix} x_{1,t-1}^C \\ x_{1,t-1}^{C*} \end{bmatrix} + \begin{bmatrix} \zeta_t \\ 0 \end{bmatrix} \\ \begin{bmatrix} x_{i,t}^C \\ x_{i,t}^{C*} \end{bmatrix} &= \rho \begin{bmatrix} \cos(\mu) & \sin(\mu) \\ -\sin(\mu) & \cos(\mu) \end{bmatrix} \begin{bmatrix} x_{i,t-1}^C \\ x_{i,t-1}^{C*} \end{bmatrix} + \begin{bmatrix} x_{i-1,t}^C \\ 0 \end{bmatrix} \end{aligned} \quad (10)$$

$i = 2 \dots c$

where $x_{i,t}^{C*}$ is an auxiliary variable needed to write the model in state space form. The properties of the cycle are obtained by writing

$$\begin{bmatrix} x_{i,t}^C \\ x_{i,t}^{C*} \end{bmatrix} = \begin{bmatrix} 1 - \rho \cos(\mu) & -\rho \sin(\mu) \\ \rho \sin(\mu) & 1 - \rho \cos(\mu) \end{bmatrix}^{-1} \begin{bmatrix} x_{i-1,t}^C \\ 0 \end{bmatrix}$$

$i = 2 \dots c$

from which one obtains

$$\begin{aligned} x_{i,t}^C &= (\alpha(L)/\beta(L))^{-1} x_{i-1,t}^C \\ \alpha(L) &= 1 - 2\rho \cos(\mu)L + \rho^2 L^2 \\ \beta(L) &= 1 - \rho \cos(\mu)L \end{aligned} \quad (11)$$

$i = 2 \dots c$.

The parameter ρ should be chosen from the open interval $]0, 1[$. It dampens the cycle, and $\rho < 1$ ensures that the cyclical process is stationary. In practice, ρ will be assigned a value close to 1, for instance $\rho = 0.975$. The parameter μ , which defines the ‘critical’ frequency that dominates the stochastic cycle, is more important. As with the value for ρ , the parameter μ can be determined on the basis of prior knowledge on the length of the cycle. Alternatively, the parameters ρ and μ can be estimated from the data in an iterative procedure discussed later in subsection 5.

By iterative substitution, one obtains

$$(1 - 2\rho \cos(\mu)L + \rho^2 L^2)^c x_{c,t}^C = (1 - \rho \cos(\mu)L)^c \zeta_t \quad (12)$$

for the c -th order stochastic cycle which we will incorporate in the TC filter: $x_t^C = x_{c,t}^C$.

The stochastic cycle model can be easily transformed to its matrix form $AX^C = B\zeta$ where A and B denote $(N - 2c) \times N$ matrices representing the AR and the MA process, respectively:

$$A = \begin{bmatrix} a_{2c} & \dots & a_1 & 1 & 0 & \dots & 0 \\ 0 & a_{2c} & \dots & a_1 & 1 & 0 & \dots & 0 \\ \vdots & & & & & & & \vdots \\ 0 & \dots & & a_{2c} & \dots & a_1 & 1 & \vdots \end{bmatrix} \quad B = \begin{bmatrix} 0 & \dots & 0 & \overset{\text{column } c+1}{\downarrow} b_c & \dots & b_1 & 1 & 0 & \dots & 0 \\ \vdots & \vdots & & 0 & b_c & \dots & b_1 & 1 & 0 & \dots & 0 \\ \vdots & \vdots & & \vdots & & & b_c & \dots & b_1 & 1 & \vdots \\ 0 & \dots & 0 & \dots & & & b_c & \dots & b_1 & 1 \end{bmatrix}$$

The first c columns of B are set equal to 0, and the a_i 's and b_i 's are determined by $\alpha(L)^c$ and $\beta(L)^c$ in equation (12).

3.3 Putting it all together: The TC filter

Combining the trend and the cycle model in matrix form gives the model of the TC filter:

$$\begin{aligned} X - X^T - X^C &= \varepsilon, & E(\varepsilon) &= 0, & E(\varepsilon'\varepsilon) &= \sigma_\varepsilon^2 \\ \nabla^{d-1}(\nabla X^T - Eb) &= \eta, & E(\eta_t) &= 0 & E(\eta_t^2) &= \sigma_\eta & \forall t = 1 \dots N, & E(\eta\eta') &= \sigma_\eta^2 I_N \\ AX^C &= B\zeta, & E(\zeta_t) &= 0 & E(\zeta_t^2) &= \sigma_\zeta & \forall t = 1 \dots N, & E(\zeta\zeta') &= \sigma_\zeta^2 I_N, \\ & & E(\zeta\varepsilon') &= \mathbf{0}_N, & E(\zeta\eta') &= \mathbf{0}_N. \end{aligned} \tag{13}$$

We assume that ζ and η are white noise error terms. Furthermore, we assume $E(\varepsilon) = 0$, that the variance σ_ε exists and that ε is uncorrelated with the other residuals. ε could follow any stationary ARMA process fulfilling these requirements and is not necessarily a white noise process.

As with the HP filter or the EES, the objective function for this problem is constructed as the sum of the inner products of the residuals $\varepsilon'\varepsilon + \eta'\eta + \zeta'\zeta$. Different from the one-component filters, however, there is no smoothing parameter (such as λ in the HP filter or the EES), and it will be explained below why this is so. This gives the following optimisation problem¹²

$$\begin{aligned} \text{Min}_{X^T, X^C, b} & (X - X^C - X^T)'(X - X^C - X^T) + \\ & (\nabla^{d-1}(\nabla X^T - Ub))'(\nabla^{d-1}(\nabla X^T - Ub)) + \\ & X^{C'} A'(BB')^{-1} AX^C. \end{aligned} \tag{14}$$

The solutions to this problem for the trend and the cyclical processes are obtained by minimising the objective function for X^T , X^C , and also for b if the trend is assumed to

¹² The last expression with X^C in equation (14) can be derived as follows: The objective function involves the minimisation of $\zeta'\zeta$. The minimisation can be carried out in two steps: First, minimise $\zeta'\zeta$ for a *given* X^C under the constraint that the stochastic cycle model $AX^C = B\zeta$ holds. This gives $\zeta = B'\kappa$, with κ as Lagrange multiplier. By replacing ζ in the stochastic cycle model, one obtains $AX^C = BB'\kappa$. From that we derive $\zeta = B'(BB')^{-1}AX^C$ and hence $\zeta'\zeta = X^{C'} A'(BB')^{-1}AX^C$.

follow a first-order random walk with drift ($d = 1$). For the sake of simplicity, let us define the following notations:

$$M_C \equiv (I + A'(BB')^{-1}A)^{-1}$$

$$M_T \equiv \begin{cases} (I + \nabla'(I - UU'(N - 1)^{-1})\nabla)^{-1}, & \text{if } d = 1 \\ (I + \nabla^d \nabla^d)^{-1}, & \text{if } d > 1. \end{cases} \quad (15)$$

We obtain the following system of first order conditions (FOCs):

$$\begin{aligned} X^T &= M_T(X - X^C) \\ X^C &= M_C(X - X^T). \end{aligned} \quad (16)$$

To explain the intuition behind the system of FOCs, observe that M_T is a one-component trend filter which transforms any series X to a trend series. For instance, assuming $d = 2$, we obtain the HP filter with $\lambda = 1$. Similarly, the matrix M_C transforms any (stationary) series to a cycle series. Indeed, it can be shown (Harvey and Trimbur 2003) that the matrix M_C defines a band-pass filter with a gain function spreading around the critical frequency μ . If the order of the stochastic cycle c is increased the cyclical filter approaches a perfect band-pass filter. Thus, the system of FOCs in equation (16) combining the trend and the cyclical band-pass filter can be interpreted as follows: applying the trend filter to a series from which the cyclical process has been removed (i.e. on $X - X^C$) gives the trend X^T , and similarly, if the band pass filter is applied to a series from which the trend has been removed (i.e. on $X - X^T$) the cyclical process follows.

From the FOCs we derive the following solutions for the trend and the cyclical process:

$$\begin{aligned} X^T &= (I - M_T M_C)^{-1} M_T (I - M_C) X \Leftrightarrow X^T = M_{TC} X \\ X^C &= (I - M_C M_T)^{-1} M_C (I - M_T) X \Leftrightarrow X^C = M_{CT} X \end{aligned} \quad (17)$$

Equation (17) defines the TC($d, c, \frac{2\pi}{\mu}, \rho$) filter with a stochastic trend of order d , a stochastic cycle of order c , a critical cycle length of $\frac{2\pi}{\mu}$ and a dampening parameter of ρ .

As equation (17) shows, the two-components TC filter can be regarded as a combination of the one-component trend and the one-component band-pass filter. For instance, using the trend filter to remove the trend in the first step and applying the band-pass filter on the residual yields $\tilde{X}^C = M_C(I - M_T)X$ as the cyclical component. However, this stepwise approach would neglect the simultaneity in the computation of the trend and the cycle and is therefore finally corrected by the correction factor $(I - M_C M_T)^{-1}$. In the special case of $M_C M_T = \mathbf{0}$, there is no simultaneity error so that the stepwise application of the trend and the cyclical filter would not differ from applying the simultaneous TC filter^{13 14}

¹³ Technically, $M_C M_T \rightarrow \mathbf{0}$ means that the intersection of the trend gain with the cycle gain in the frequency domain becomes smaller. This implies that the contribution of the trend to identify the cycle (and vice versa) becomes smaller and that trend and cycle become increasingly independent from each other. Ceteris paribus, the intersection of the gain functions decreases when critical cyclical frequency of the cycle μ becomes higher when the order of the stochastic trend, d , or of the stochastic cycle, c , become smaller.

¹⁴ Equation (17) gives consistent results if one component is missing. For instance, assume that there

As mentioned above, the variance components in the TC filter objective function (14) are not weighted. As the TC filter contains two components which are modelled (the trend and the cycle), two weighting parameters, λ_1 and λ_2 , are necessary to define the objective function with weights as $\varepsilon'\varepsilon + \lambda_1\eta'\eta + \lambda_2\zeta'\zeta$. Under certain assumptions in addition to those in equation (13), minimising the weighted objective would provide the optimal filter for the process defined in equation (13)¹⁵. However, deriving an optimal filter is not our objective. Instead, we want to extend the HP filter with a cyclical model in order to improve certain properties of the HP filter and in order to account for prior assumptions on the cyclical process in more straightforward manner.

With the HP filter, prior assumptions about the cyclical process are in principle reflected in the choice of the smoothing parameter λ . However, as discussed above, the relationship between the assumed cyclical process and the value of lambda is unclear. The TC filter trend can be interpreted as an HP filter trend in which the smoothing parameter λ is replaced by a more complex expression reflecting prior assumptions on the length of the cycle. Rewriting the trend in equation (17) as

$$X^T = (I + (I + (A'(BB')^{-1}A)^{-1})\nabla^d \nabla^d)^{-1}X.$$

reveals that the trend of the TC filter is similar to the HP filter trend in equation (3) with λ replaced by the matrix expression $I + (A'(BB')^{-1}A)^{-1}$. Since this expression depends on μ , the critical frequency of the cycle, it reflects the prior assumption on the average cycle length¹⁶. Thus, by amending the HP filter with a model for the cycle, we have replaced the—to a certain extent arbitrary—smoothing parameter λ with a more general model based expression providing a clear-cut relationship between the cycle length and the filter parameter μ .

3.4 Properties of the TC filter in the time domain

As equation (17) shows, both the stochastic trend and the stochastic cycle model affect the trend and cycle solutions. This is so because the trend and the cycle are determined simultaneously; prior information on the nature of one component is used to identify the other component.

The TC filter reproduces deterministic trends up to order¹⁷ $2d - 1$. This can easily be seen if there is only a trend and no cyclical component, implying $M_C = \mathbf{0}$. It follows that the two-components trend filter collapses to the one-component trend filter: $M_{TC} = M_T$. Respectively, if there is no trend, i.e. if $M_T = \mathbf{0}$, it follows that $M_{CT} = M_C$.

¹⁵ The additional assumptions are that ε , η and ζ are all normally distributed and that the weights are set equal to the respective inverse signal-to-noise variance ratios: $\lambda_1 = \frac{\sigma_\varepsilon^2}{\sigma_\eta^2}$ and $\lambda_2 = \frac{\sigma_\varepsilon^2}{\sigma_\zeta^2}$. However, in equation (13) these variance ratios have been implicitly set to one. This is an important difference to the general Kalman filter approach in (Harvey and Trimbur 2003), in which signal and noise variances are estimated simultaneously with the trend and cycle. Like the HP filter, the TC filter is in general not an optimal filter.

¹⁶For instance, assuming a relevant cycle length of eight years, μ could be set to $\frac{2\pi}{8} \approx 0.8$ with annual data.

¹⁷ A deterministic trend of order k is defined as $\sum_{i=0}^k a_i t^i$ with t denoting the time index.

shown by rewriting the trend in equation (17) as

$$(I + A'(BB')^{-1}A)\nabla^{d'}\nabla^d X^T = A'(BB')^{-1}A(X - X^T).$$

Preserving a deterministic trend implies $X = X^T$, so that the condition $\nabla^{d'}\nabla^d X = 0$ follows. This is equivalent to $L^{d'}\nabla^{2d}X$. As the $2d$ -th difference of any trend of order $2d - 1$ is zero, a trend of order $2d - 1$ fulfills the condition. The TC filter resembles in this respect the HP filter, which preserves deterministic trends of at most third order.

Unlike the HP filter, however, the TC filter preserves deterministic, stationary cycles as well, and its trend is *cyclically neutral* as long as the cycle in the data is consistent with the cyclical model of the filter. This means that applying the TC filter on such a process reproduces the input process completely in the cycle and yields a zero trend. In order to prove this we set $X^C = X$ in equation (17) and derive the condition

$$(I + \nabla^{d'}\nabla^d)A'(BB')^{-1}AX = 0.$$

For this condition to hold it is sufficient that $AX = 0$. This is the case if X is generated by $\alpha(L)^k X = 0$, for $1 \leq k \leq n$ and with $\alpha(L)$ defined as in equation (11). The cyclical neutrality of the trend follows immediately from equation (17) together with the assumption that $AX = 0$.

The cyclical neutrality of the trend is an important improvement compared with a HP filter trend, which is not cyclically neutral: depending on the value of the smoothing parameter λ , the HP filter reproduces harmonic oscillations partly in the trend¹⁸.

The equations of the trend and the cyclical process are completely symmetrical. One can switch from one equation to the other by exchanging ∇^d against A . Indeed, the matrices ∇^d and A can be regarded as containers for arbitrary but distinctive stochastic processes. It is even possible to include exogenous variables in order to identify the trend and the cycle as, for instance, the inflation rate, indicators of capacity utilisation or of consumer sentiments. This is similar to the Multivariate HP filter as proposed by Laxton and Tetlow (1992)¹⁹. Furthermore, as shown in Annex C, structural breaks can be included in a straightforward manner, assuming that the timing of the break is known a priori. Examples are the change from the 1979 to the 1995 European System of National Accounts (ESA) or the German unification in 1991—both are events, which gave rise to jumps in macroeconomic data in specific periods. Although of practical relevance, we do not investigate this issue further but proceed with the frequency domain properties of the standard TC filter.

3.5 Properties of the TC filter in the frequency domain

In this subsection we analyse the properties of the trend-cycle filter in the frequency domain. We derive the polynomial lag forms and subsequently the frequency domain

¹⁸ This is so because the HP filter cannot approximate an ideal filter perfectly, as explained in section 2. The HP filter would give a zero trend only in the limiting case of $\lambda \rightarrow \infty$. The other polar case of $\lambda = 0$ just reproduces the input process. The incorporation of cyclical fluctuations in the HP filter trend reflects the leakage effects of the filter explained above.

¹⁹ For a more recent application of the multivariate HP filter see Gruen et al. (2002) and Boone et al. (2000).

representations—i.e. the *power transfer functions* (PTFs)— of the trend and the cyclical filter in equation (17).

The matrices A and ∇^d in equation (17) are matrix-form translations of the polynomial lags for the stochastic cycle $\gamma(L)^c = (\alpha(L)/\beta(L))^c$ —with $\alpha(L)$ and $\beta(L)$ defined as in equation (11)—and the stochastic trend, $(1 - L)^d$. The transposes of these matrices represent the respective *lead*-polynomials $\gamma(L^{-1})^c$ and $(1 - L^{-1})^d$ in matrix-form. The polynomial lag forms of the trend and the cyclical filter in (17) can therefore easily be derived by replacing ∇ and A with $1 - L$ and $\gamma(L)$ and their transposes with $1 - L^{-1}$ and $\gamma(L^{-1})$, respectively. After simplifying we have:

$$\begin{aligned} x_t^T &= \left(1 + \left(1 + (\gamma(L)\gamma(L^{-1}))^{-c}\right) \left((1 - L)(1 - L^{-1})^d\right)^{-1}\right) x_t \\ &= G_{TC}^T(L, \omega)x_t \\ x_t^C &= \left(1 + \left(1 + ((1 - L)(1 - L^{-1}))^{-d}\right) (\gamma(L)\gamma(L^{-1}))^c\right)^{-1} x_t \\ &= G_{TC}^C(L, \omega)x_t. \end{aligned} \tag{18}$$

The corresponding gain functions, $G_{TC}^T(\omega)$ and $G_{TC}^C(\omega)$, can be obtained by replacing the lag operator L in equation (18) with $U^{-i\omega}$ with ω as the frequency in radians.

As the filters are symmetric, the PTFs are equal to the squared gain functions. The impact of the parameters d , c , μ , and ρ on the behaviour of the TC filter can be best explained by visual inspection of the PTFs as shown in Figure 2 for different parameter settings.

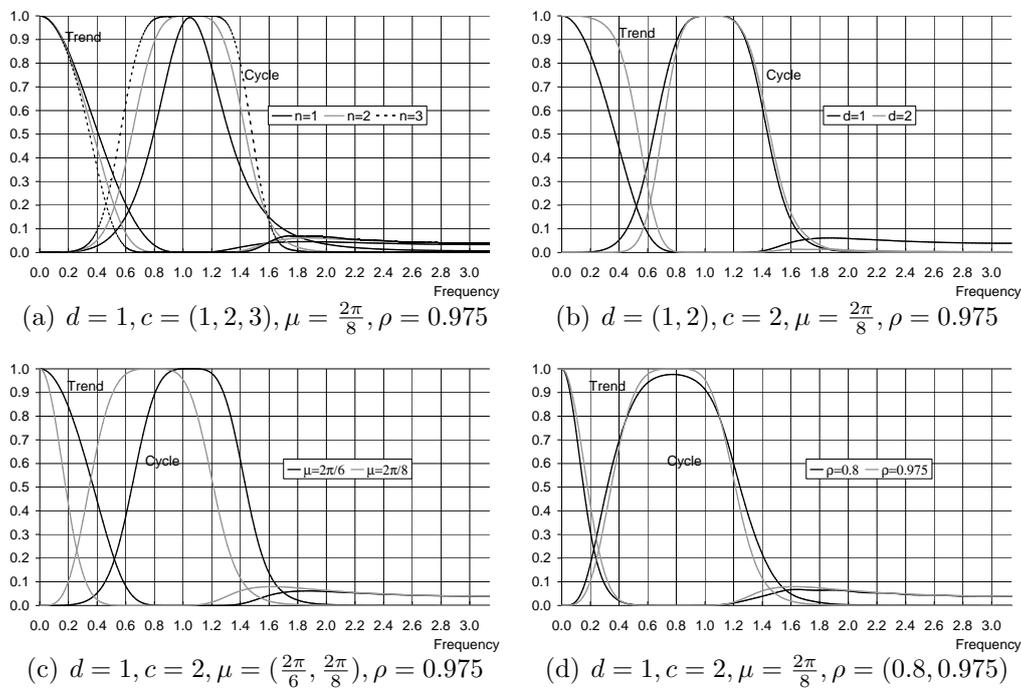


Figure 2: Power-transfer functions of the trend and the cycle of the TC filter for different parameter values

The order of the stochastic cycle, c , determines the band-width of the frequency spectrum contained in the cyclical process. The spectrum expands around the critical frequency μ when c becomes larger. Increasing the order of the stochastic cycle also shifts the trend spectrum to the lower frequency range. This is a consequence of the simultaneous determination of the trend and the cycle. However, the impact of changes in c on the trend-spectrum is minor.

The critical frequency μ determines the centre of gravity in the frequency spectrum of the cycle. Changes in μ give also rise to unidirectional shifts in the position of the trend spectrum, implying that μ does not only affect the volatility of the cycle but to some extent the trend volatility as well. Again, this feature follows from the simultaneous determination of the trend and the cycle.

An increase in the order of the stochastic trend d takes higher frequencies into the trend spectrum, implying that the trend becomes more volatile. The impact of changes in d on the cycle-gain are minor. Thus, by setting the order of the stochastic trend, the trend volatility can be manipulated without affecting the cycle too much, whereas the properties of the cycle are mainly determined through μ and c .

The parameter ρ is necessary to ensure the stationarity of the cycle and should be set close to but less than 1. As Figure 2d shows, the power-transfer functions are quite robust against changes in ρ .

4 An application to real GDP in selected countries

Now we apply variants of TC filter to annual real GDP from 1970-2002 in Germany (DE), Spain (ES), France (FR), Italy (IT), the euro area (EURO), and in the US and compare the results to those obtained with the HP filter and the *Extended Exponential Smoothing* (EES) as suggested by Tödter (2002). The data source is the spring 2004 AMECO database of the European Commission. In order to adjust for the structural jump in the German and the euro area series owing to the German unification, German real GDP was regressed on a constant, a linear trend and a jump dummy which takes a value of 1 from 1991 onwards and of 0 before. The estimated shift parameter value was then added to real GDP before 1991.

We choose a value of 7 for the smoothing parameter for the EES, following Tödter (2002). We fix the λ parameter for the HP filter to 30, as in Bouthevillain et al. (2001). We define an 8 years reference cycle for the TC filters, i.e. $\mu = \frac{2\pi}{8}$, and set the dampening parameter $\rho = 0.975$.

Figure 3 shows the resulting relative cyclical components for the TC(1,2), the TC(2,2), the HP(30) filter and the EES(7). The cyclical components are very similar to each other in the middle of the sample, with the exception of comparatively large TC(1,2) cycles for Spain and the US. More important, however, are the significant differences we observe at the sample fringes: The procession of end-of-sample information seems to constitute the most distinctive feature.

Furthermore, the patterns of trend growth generated with a TC filter are less smooth than the trend growth pattern derived from the one-component filters (Figure 4). In fact, the HP filter has often been criticised for generating an implausibly cyclical—even

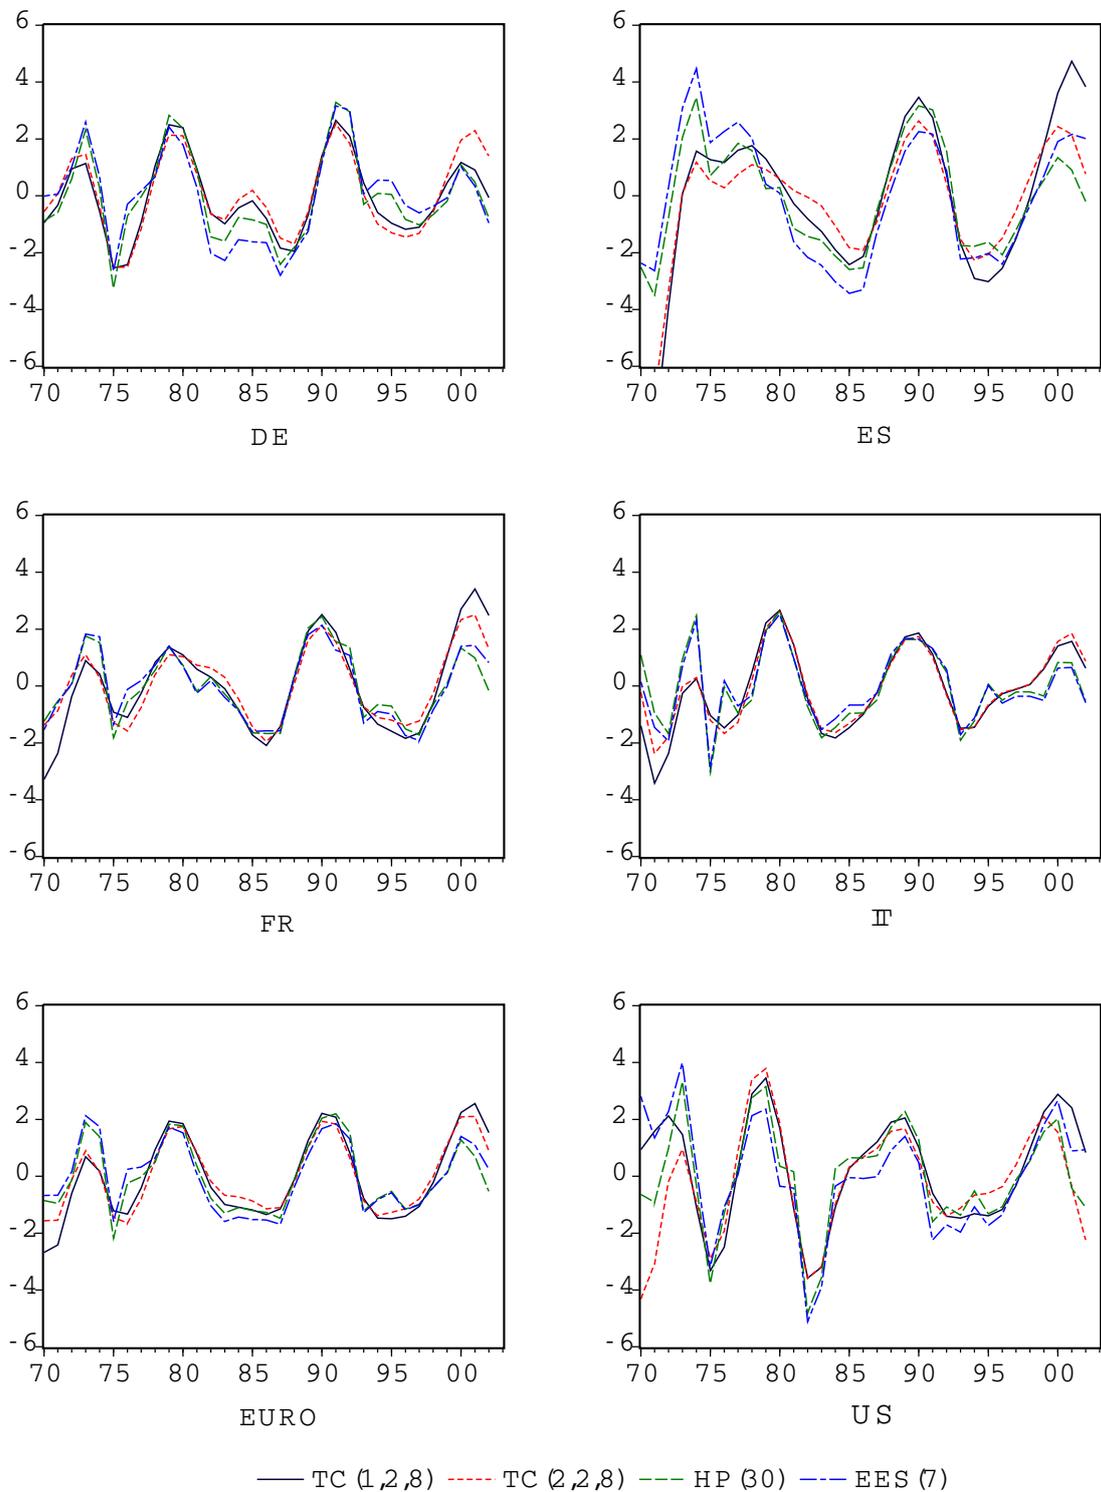


Figure 3: Cyclical components of real GDP (in % of real GDP)

pro-cyclical—pattern in trend growth, which is difficult to reconcile with the common assumption that the long run growth path is mainly affected by irregular supply shocks. At the first sight, it seems as if the zig-zag like movements in the TC filter trend growth rates are more in line with this prior assumption than the patterns of the HP filter or the EES trend growth rates.

In the next sections we analyse the properties of trends and cycles computed with the TC filter more thoroughly and compare them with trends and cycles generated with the HP filter and the EES. In the first subsection, the issue of the so-called endpoint problem is investigated from a more theoretical perspective. It is argued that the forecasting capability of the stochastic model underlying the filter is the main variable triggering the end-of-sample instability. In the second subsection, we explore the forecasting performance of the filters empirically. We find that the stochastic cycle model improves the forecasting performance of filters considerably. Finally, it is shown that some of the assumptions underlying the TC filter can be tested and that the TC filter can to some extent be adjusted to the data.

4.1 The endpoint problem and the predictive capabilities of filters

Many trend-cycle decompositions suffer from the so called end-point problem. The trend in the final period N , x_N^T , is based on information available up to and including period N . It can change significantly if new data for period $N + 1$ become available—irrespective of whether the new data point is driven by cyclical or by structural factors. The real-time allocation of the dynamics to structural and cyclical forces is necessarily uncertain as information on the future path of the economy missing. It is only when new data in future periods become available that the trend-cycle decomposition in period N becomes more certain and stabilises.

While the limited amount of real-time information is a general problem for any trend-cycle decomposition that relies on past and future periods, trend extraction tools differ in the significance of the problem. The problem is less significant, the better the model underlying the filter can forecast the original time series. This can be illustrated by taking the example of the HP filter stochastic model.

The stochastic model of the HP filter can be used to forecast x_{N+1} in period N , once the trend value in N is given. As the trend model is a second order random walk and because the cycle is not modelled, it follows that the optimal forecast for x_{N+1} is equal to $\hat{x}_{N+1} = 2x_{N-1}^T - x_{N-2}^T$. Now extend the original series by \hat{x}_{N+1} to obtain $[x_1 \dots x_N, \hat{x}_{N+1}]$ and apply the HP filter to the extended series. As a result, the trend series up to period N $[x_t^T \dots x_N^T]$ is identical to the one obtained from filtering the non-extended series; the HP filter is consistent with its own forecast (Kaiser and Maravall 1999).

From this we can conclude that there is no endpoint problem if new data that arrive in $N + 1$ comply with the implicit forecast of the HP filter. Stating it the other way round: an end-point problem exists only insofar as the stochastic model underlying the filter is a weak representation of the data generating process.

As a standard remedy to the end-point problem, time series are sometimes extended

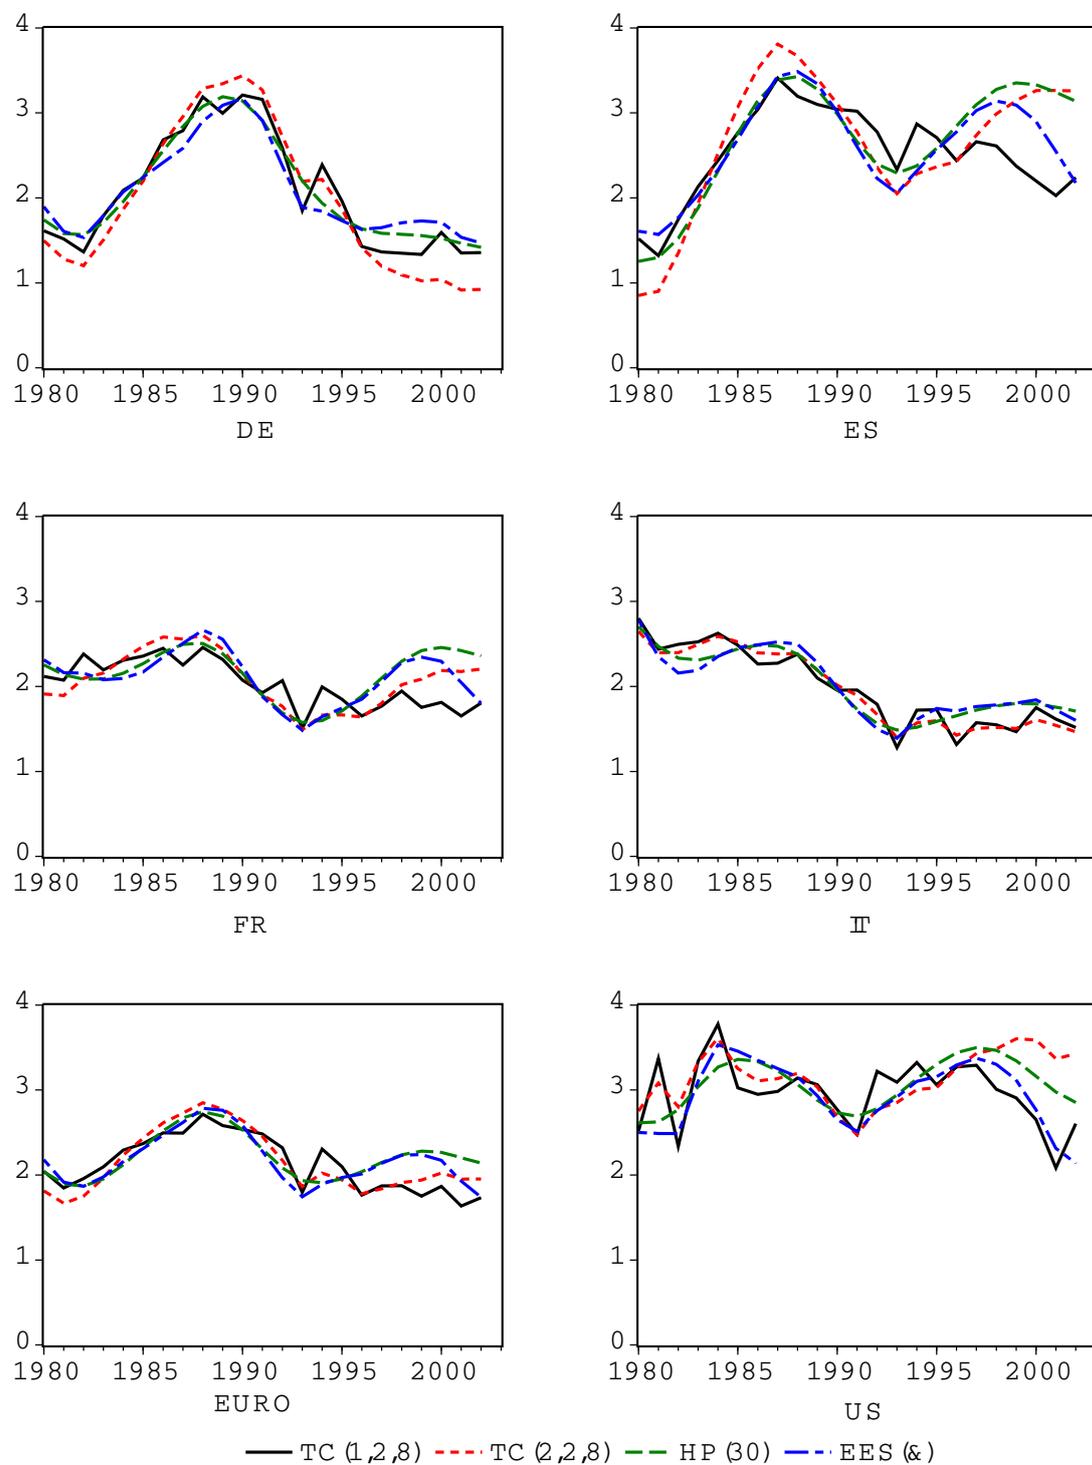


Figure 4: Trend growth rates real GDP in %

by forecasts²⁰, and the filter is applied to the extended series. If the forecast turns out correct ex post, there would not be an end-point bias. However, this approach comes with other problems. It is unclear how the filter processes forecast errors, which translate into errors in the trend estimation. Even if the forecast itself is unbiased and the forecast error is a random white noise process, it is unlikely that the implied errors in the computation of the trend share this feature because the filter model differs from that underlying the forecast.

As we have seen, the HP filter is consistent with forecasts derived from its own time series model. Extending the time series on the basis of a different model means that one does not trust the filter model. However, if there are good reasons to assume that there exists a model with a better forecasting performance than the filter model, the former rather than the latter should be applied for the trend-cycle decomposition.

Thus, rendering the filter model more consistent with the data generating process is a more preferable solution to the endpoint problem than data extensions on the basis of models inconsistent with the filter. It follows that the endpoint problem should be alleviated by improving the forecast performance of the stochastic filter model, i.e. its fit to the actual data.

The forecast performance of the filter and the possibilities to adjust it to the data depend mainly on the complexity of the underlying model. The complexity of the stochastic model of the HP filter, for instance, is low: the second order random walk property of the trend is the only prior piece of information that can be exploited for forecasting. Furthermore, the HP filter provides practically no means to adjust it to the data. Hence, its forecast performance cannot be improved.

The TC filter on the other hand provides a somewhat richer stochastic model as it explicitly accounts for the cycle; but does it give better forecasts and what are the empirical implications for the end-of-sample trend-cycle decomposition?

4.2 The forecasting performance of the HP and the TC filter

We investigate now the iterative one-step-ahead forecasts of the TC and HP filters and the EES. Starting with the sample 1970–1978, we increase the “last year” s of the sample step by step until 2001, apply the filter on each vintage and compute for each of the filters a series from 1979–2002 of one-step-ahead forecasts $\hat{x}_{s+1|s}$ on the basis of the respective stochastic filter model:

$$\hat{x}_{s+1|s} = \begin{cases} x_s^T + b & \text{for the EES} \\ 2x_s^T - x_{s-1}^T & \text{for the HP filter} \end{cases}$$

where $s = 1978 \dots 2002$. The forecasts generated by the TC filter contains two components: the trend forecast \hat{x}_{s+1}^T generated by the stochastic trend model and cycle forecast \hat{x}_{s+1}^C derived from the stochastic cycle model. Note that only the *AR* and not the *MA* part of the stochastic cycle is used to generate the forecast since expected forecast errors

²⁰ The forecasts are often derived from ARIMA models as for instance in Kaiser and Maravall (1999) and in Denis et al. (2002)

are assumed to be equal to zero.

$$\hat{x}_{s+1|s} = \begin{cases} \underbrace{x_s^T + b}_{\hat{x}_{s+1|s}^T} + \underbrace{\sum_{i=1}^{2c} a_i x_{s-i+1}^C}_{\hat{x}_{s+1|s}^C} & \text{for the TC}(1, c) \text{ filter} \\ \underbrace{2x_s^T - x_{s-1}^T}_{\hat{x}_{s+1|s}^T} + \underbrace{\sum_{i=1}^{2c} a_i x_{s-i+1}^C}_{\hat{x}_{s+1|s}^C} & \text{for the TC}(2, c) \text{ filter} \end{cases}$$

where $c = 1, 2$ and $s = 1978 \dots 2001$. The quality of the forecasts can be assessed by testing for $b = 1$ and $\text{const} = 0$ in the regression

$$\Delta x_t = \text{const} + b\Delta\hat{x}_{t|t-1} + u_t. \quad (19)$$

In the case of the TC filter, the additional variance explained by stochastic cycle forecast can be assessed by comparing the explained variance in equation (19) to that in the reduced regression

$$\Delta x_t = \text{const} + b\Delta\hat{x}_{t|t-1}^T + v_t \quad (20)$$

which contains only the trend forecast of the TC filter model.

Table 1—we present only the euro area results of this test because they are similar for the other countries—shows the result of the forecast regressions, together with some indicators of forecast quality, the root mean square error (RMSE), the mean absolute percentage error (MAPE), Theil's inequality coefficient and the coefficient of correlation between the one-step ahead predictions and actual values²¹. The bias and the variance proportion measure the part of the MSE due to differences in the mean and the variation between the predicted and the actual series. The covariance proportion captures remaining unsystematic forecasting errors. The bias, variance and covariance proportion add up to one. Ideally, the bias and variance proportions should be small so that most of the bias concentrates on the covariance proportion.

All filter models predict real GDP growth in the euro area well and are unbiased. The correlation between predicted and actual GDP growth rates increases considerably with the complexity of the underlying filter model; the TC(1,2) and the TC(2,2)-forecast of real GDP growth explain about 80% of actual growth, the EES-forecast only 38%. Furthermore, the stochastic cycle model improves the fit to the data substantially as compared with the forecasts exclusively based on trends. Growth forecasts on the basis of the TC filter variants yield lower RMSE's, lower mean absolute percentage errors and lower Theil inequality statistics than forecasts using the stochastic models HP filter and the EES. The decomposition of the MSE reveals that it is almost fully explained by the non-systematic covariance component in the case of the TC filter, whereas considerable contributions to the mean square error (13.8% in the case of the HP filter and almost 38% with the EES) derive from differences in variation between predicted and actual growth rates when predictions are based on the HP filter and the EES models.

²¹ Theil's inequality coefficient is defined as $\sqrt{\frac{MSE}{\sum \hat{x}^2/n + \sum x^2/n}}$. It takes values between 0 and 1, with values closer to unity indicating worse predictors. The indicators used here are described—for instance—in Maddala (1977).

Filter	Regression [§]			Full forecast			Regression ^{§§}			Trend forecast		
	const	b	Indicators	Forecast error	MSE [†] decomposition	const	b	Indicators	Forecast error	MSE [†] decomposition		
TC(1,1)	Parameter	0.001	0.939	RMSE [†]	0.007	0.000	0.000	0.990	RMSE [†]	0.010	Bias	0.002
	Stdv.	0.004	0.16	MAPE	0.378	0.072	0.008	0.382	MAPE	0.687	Variance	0.340
	F-test [‡]	0.928		Theil	0.156	0.928	0.982		Theil	0.223	Covar.	0.658
			Corr.	0.781				Corr.	0.484			
TC(1,2)	Parameter	0.002	0.896	RMSE	0.005	0.000	0.022	-0.083	RMSE	0.013	Bias	0.008
	Stdv.	0.002	0.091	MAPE	0.208	0.00	0.011	0.497	MAPE	0.900	Variance	0.267
	F-test [‡]	0.530		Theil	0.108	0.999	0.107		Theil	0.279	Covar.	0.726
			Corr.	0.902				Corr.	-0.035			
TC(2,1)	Parameter	0.004	0.818	RMSE	0.006	0.001	0.000	0.980	RMSE	0.006	Bias	0.008
	Stdv.	0.003	0.11	MAPE	0.276	0.004	0.003	0.137	MAPE	0.418	Variance	0.070
	F-test [‡]	0.273		Theil	0.138	0.996	0.908		Theil	0.135	Covar.	0.922
			Corr.	0.846				Corr.	0.837			
TC(2,2)	Parameter	0.003	0.853	RMSE	0.005	0.001	0.012	0.363	RMSE	0.013	Bias	0.014
	Stdv.	0.002	0.09	MAPE	0.232	0.011	0.006	0.231	MAPE	0.987	Variance	0.013
	F-test [‡]	0.283		Theil	0.114	0.987	#0.033		Theil	0.267	Covar.	0.973
			Corr.	0.896				Corr.	0.317			
HP(30)	Parameter						-0.000	0.987	RMSE	0.008	Bias	0.009
	Stdv.						0.004	0.192	MAPE	0.531	Variance	0.138
	F-test [‡]						0.901		Theil	0.168	Covar.	0.853
								Corr.	0.738			
EES(7)	Parameter						-0.005	1.204	RMSE	0.009	Bias	0.003
	Stdv.						0.007	0.331	MAPE	0.649	Variance	0.379
	F-test [‡]						0.800		Theil	0.203	Covar.	0.618
								Corr.	0.613			

[§] $\Delta x_t = \text{const} + b\Delta\hat{x}_{t|t-1} + u_t$ ^{§§} $\Delta x_t = \text{const} + b\Delta\hat{x}_{t|t-1}^T + u_t$

[†] *MSE*: mean square error; *RMSE*: root mean squared error; *MAPE*: mean absolute percentage error;

[‡] *Theil*: Theil inequality measure; *Corr*: Correlation coefficient.

[‡] P-value of Wald-test of H0: const = 0 \wedge b = 1 # H0 rejected at 5% significance level

Table 1: Regression of real GDP growth on one-step-ahead forecasts for real GDP growth for the euro area

To conclude, the endogenous stochastic cycle seems to improve the fit of the stochastic filter model to the actual data²². Therefore, we expect the TC filter to yield more reliable real time trend/cycle estimations than the EES or the HP filter.

4.3 The real time reliability of the TC filter

In order to assess the end-point reliability of trend-cycle decompositions, we generate vintages of trend-cycle estimations by cutting the sample artificially in each year s from 1978-2003 and estimating the trend and the cycle for each sample 1970- s . In this way we obtain for each year s between 1978-2003 one end-point trend/cycle estimation based on the sample 1970- s , the so-called *real-time* estimations $\tilde{x}_s^T, \tilde{x}_s^C$ of the trend and the cycle²³. The regression of the real-time cyclical components \tilde{x}_t^C on the “final” results x_t^C of the 2002 vintage

$$\tilde{x}_t^C = \text{const} + bx_t^C + u_t \quad (21)$$

indicates in how far the real time cyclical components are related to the “true” (the final) ones.

In Rünstler (2002), the “reverse” regression of final on real time results is proposed, which is based on the assumption that deviations of real-time from final results are uncorrelated with real time results. This property of optimal, linear filters is a necessary condition for unbiased, minimum mean square errors of the filter components²⁴, assuming that the underlying stochastic model is correct. Hence, the test in Rünstler (2002) is based on the idea that the filter makes optimal use of real-time information so that subsequent revisions to initial estimates—once additional information comes in—should be orthogonal to the initial estimates. It can therefore be understood as a mis-specification test of the stochastic model underlying the filter. However, as argued above, neither the TC filter, nor the HP filter, nor the EES can be regarded as optimal filters for typical economic time series. Here, we are more interested in the question whether errors are systematically pro- or anti-cyclical when compared to “final” trend deviations and not so much in a specification test for the underlying stochastic model. Under the H0 that errors are not systematically related to “final” results, they should be orthogonal to “final” estimates and the test regression should be specified as in equation (21).

Thus, end point reliability implies that $b = 1$ and $\text{const} = 0$ in equation (21) hold so that real-time cyclical components should be in broadly line with “final” cyclical components. Table 2 presents the results of these regressions, together with the P-value for the Wald test of the joint H0: $\text{const} = 0 \wedge b = 1$.

For the HP filter, the H0 must be rejected in all cases. While the constant is not significantly different from zero, b is consistently below 1: the HP filter cyclical components in real-time underestimate the “true” cycle considerably. In addition, the correlations of

²² It must be kept in mind, though, that an approach with prior parameterisation cannot deliver an optimal fit.

²³ More precisely, these are known as *quasi*-real time vintages, as the s -th vintage does not consist of the data available on period s , but of data available in T . We thus disregard data revisions.

²⁴See Priestley (1981, 775).

Filter	Parameter	DE	ES	FR	IT	EURO	US
TC(1,1)	c	1.379	1.896	0.209	-1.919	4.913	46.935
	std. err. [†]	5.746	2.281	2.982	1.482	12.826	22.257
	beta	0.898	0.999	0.91	1.015	0.914	0.93
	std. err. [†]	0.295	0.253	0.168	0.145	0.17	0.272
	Ftest [‡]	0.937	0.703	0.861	0.359	0.861	0.114
	Correlation	0.683	0.572	0.76	0.849	0.74	0.683
TC(1,2)	c	8.813	5.007	1.582	-1.29	21.831	118.695
	std. err. [†]	9.776	2.728	2.891	1.975	15.047	25.542
	beta	1.424	1.297	1.209	1.317	1.29	1.596
	std. err. [†]	0.406	0.171	0.093	0.132	0.146	0.209
	Ftest [‡]	0.436	#0.024	#0.025	0.074	#0.047	#0.000
	Correlation	0.68	0.807	0.913	0.863	0.849	0.86
TC(2,1)	c	0.01	0.125	0.344	-0.703	0.095	3.923
	std. err. [†]	4.277	1.368	2.567	1.77	10.978	16.373
	beta	0.587	0.637	0.525	0.608	0.538	0.626
	std. err. [†]	0.24	0.24	0.169	0.186	0.192	0.131
	Ftest [‡]	#0.005	0.327	#0.008	0.072	#0.008	#0.029
	Correlation	0.561	0.524	0.531	0.527	0.52	0.615
TC(2,2)	c	1.038	1.463	1.591	0.47	7.37	44.273
	std. err. [†]	9.299	2.845	4.275	2.443	20.188	21.93
	beta	1.354	1.374	1.077	1.355	1.233	1.372
	std. err. [†]	0.346	0.331	0.189	0.187	0.229	0.196
	Ftest [‡]	0.235	0.302	0.717	0.055	0.181	0.061
	Correlation	0.67	0.662	0.75	0.804	0.727	0.874
HP(30)	c	-1.086	1.042	0.715	-0.795	1.738	18.927
	std. err. [†]	5.981	2.001	3.216	1.89	13.613	21.055
	beta	0.422	0.332	0.43	0.503	0.431	0.485
	std. err. [†]	0.177	0.174	0.135	0.11	0.13	0.095
	Ftest [‡]	#0.005	#0.003	#0.001	#0.000	#0.001	#0.000
	Correlation	0.471	0.33	0.511	0.617	0.51	0.589
EES(7)	c	5.208	4.002	0.984	-3.451	13.719	95.813
	std. err. [†]	6.268	2.501	2.808	1.405	12.906	24.857
	beta	0.695	0.741	0.75	0.766	0.717	0.701
	std. err. [†]	0.225	0.189	0.127	0.088	0.145	0.162
	Ftest [‡]	0.303	0.067	0.146	#0.002	0.078	#0.001
	Correlation	0.674	0.598	0.765	0.835	0.722	0.709

Equation: $\tilde{x}_t^C = \text{const} + bx_t^C + u_t$

[†] Newey-West corrected standard errors

[‡] P-value of F-test of $H_0: \text{const} = 0 \wedge b = 1$

H_0 rejected at 5% significance level

Table 2: Regression of the real-time cyclical component on the final cyclical component and correlation between the real-time and the final cyclical component of real GDP

the real-time with “final” cyclical components are low; the “true” cycle explains at most 38% of the variance²⁵ of the cyclical component estimated at real time.

The results are slightly better for the EES. Here, the H_0 $\text{const} = 0 \wedge b = 1$ cannot be rejected except in the cases of Italy and the US²⁶. The slope parameter is closer to 1 than in the case of the HP filter. In two cases (Italy and the US), the real time EES estimates are strongly biased, as the constant is significantly different from zero. The coefficient of correlation between the real time and final cyclical components varies between 0.60 and 0.84, which is higher than for the HP filter.

The TC(2,2) filter turns out best in this exercise. The H_0 is never rejected at the 5% level²⁷. The slope parameter b is close to one, the constant is not significantly different from zero, and the coefficient of correlation varies between 0.57 and 0.91. Decreasing the order of the cycle while maintaining the order of the trend comes at the cost of a considerable decrease in correlation between real-time and final cyclical components. Decreasing the order of the trend gives rise to rejections of the combined H_0 in Spain, France, the euro area and the US. Depending on the time series being filtered, the parameters of the TC filter can to some extent be chosen to adapt the filter to the data generating process.

The underestimation of b gives rise to a pro-cyclical error in the estimation of the trend. This can easily be seen if we approximate the cyclical component by $x_t - x_t^T$. The regression equation $x_t - \tilde{x}_t^T = \text{const} + b(x_t - x_t^T) + u_t$ can be transformed into $x_t^T - \tilde{x}_t^T = \text{const} - (1 - b)(x_t - x_t^T) + u_t$. Values of b between -1 and 1 and different from zero imply that the trend is underestimated in a recession and overestimated in a boom. If $b = 1$ there is no relationship between the cycle and the error in the trend.

Figure D.2 in Annex D compares the errors in the real-time trend with the final cyclical components for the TC(2,2) and the HP(30) filter and the EES(7). As expected, the errors in the real time trend of the TC filter are largely unrelated to the cyclical component. For the HP filter, however, this relationship is strong. The HP filter real-time trend errors approximate very well the final cyclical component. Likewise, the EES induces a pro-cyclical bias in the real-time trend estimations, although the bias is less pronounced than in the case of the HP filter.

An important feature of real-time assessments of the cycle is the behavior around business cycle turning points. Errors in the real-time detection of the “true” turning points might lead to a misdiagnosis of the current situation. The extent the different approaches to trend-cycle decomposition are prone to errors in the detection of turning points can be assessed by the following indices, which rest on the classification shown in Table 3:

- The relative share of wrong signs $(N_{[+-]} + N_{[-+]})/N_{[.]}$.
- The information content defined as $I \equiv N_{[++]}/N_{[+]} + N_{[--]}/N_{[-]} - 1$. This measure takes values between -1 and 1 . Values in the range $0 < I \leq 1$ indicate a positive information content, and $I = 1$ means that the signs of cyclical components in real

²⁵The highest coefficient of correlation amounts to 0.617 (in the case of for IT) so that the explained variance would be $\rho^2 = 0.38$.

²⁶For the euro area and Spain, it would be rejected at the 10% level.

²⁷It would be rejected at the 10% level in Italy and in the US.

		final output gap		
		+	-	sum
real time	+	N_{++}	N_{+-}	$N_{+.}$
output gap	-	N_{-+}	N_{--}	$N_{.-}$
sum		$N_{.+}$	$N_{.-}$	$N_{..}$

Table 3: Reliability of signs of real-time cyclical components

time and final estimates coincide perfectly. If $-1 \leq I < 0$, there is a systematic bias in the signs of cyclical components in real time.

- The cell counts can be compared with the expected ones under the H0 that cell counts are random: $E(N_{[ij]}) = N_{[i.]}N_{[.j]}/N_{[.]}$, $i, j \in \{+, -\}$. The H0 can be tested, using the test statistic $\sum_{i,j \in \{+,-\}} (N_{[ij]} - E(N_{[ij]}))^2 / E(N_{[ij]}) \sim \chi^2(1)$.

Results for these indices for cyclical components of the TC filters with a second order cycle, the HP filter and the EES are shown in table 4. There is no instance with a negative value for I so that the signs of the real-time cyclical components cannot be regarded biased. The relative share of sign misdiagnoses amounts to roughly 10-25% with the TC filter variants. Signs of cyclical components are likewise often wrongly estimated with the EES except in the case of the US, where the EES gives the highest share (38%) of instances with wrong signs. For the other countries and regions, the HP filter yields the highest shares of wrong signs between 35 and 46%. Correspondingly, the HP filter gives the lowest value for the information content measure I , again with the exception of the US, where the EES performs worse. For all regions except Germany and France, I is generally closer to unity for the TC filter variants. In Germany the EES outperforms both trend variants of the TC filters. In France the EES gives a higher value for I than the TC(1,2) filter. The H0 that the cell counts are random can never be rejected at the 5% level with the HP filter. Only HP filtered real GDP in Germany leads to a rejection of the H0 at the 10% level. According to the χ^2 test, the hypothesis of a random distribution of signs can be rejected at least at the 5% significance level for cyclical components computed with the TC Filter and the EES. All in all, the TC filter generally allows for a more consistent determination of signs of cyclical components in real time than the one-component filters. The EES performs remarkable well in this test, while results for the HP filter are less satisfying.

The comparatively weak real time properties of the one-component filters—the HP filter and the EES—derive from the “missing cycle” in these filters. Enhancing these filters with stochastic models for the cycle improves the real-time reliability significantly and removes the pro-cyclical bias in end-of-sample estimates. Obviously, it is not possible to identify the trend at real time in a proper way if a model for the cycle is missing.

Country	Filter	Wrong sign	I	Test statistic	p-value	Significance [†]
DE	TC(1,2)	0.23	0.55	7.80	0.005	***
	TC(2,2)	0.19	0.62	10.40	0.001	***
	EES(7)	0.15	0.69	12.76	0.000	***
	HP(30)	0.35	0.36	3.31	0.069	*
ES	TC(1,2)	0.19	0.62	10.40	0.001	***
	TC(2,2)	0.27	0.45	5.42	0.020	**
	EES(7)	0.31	0.39	3.94	0.047	**
	HP(30)	0.46	0.08	0.18	0.671	
FR	TC(1,2)	0.08	0.87	19.07	0.000	***
	TC(2,2)	0.23	0.55	7.72	0.005	***
	EES(7)	0.15	0.69	12.76	0.000	***
	HP(30)	0.35	0.31	2.48	0.116	
IT	TC(1,2)	0.08	0.85	18.62	0.000	***
	TC(2,2)	0.15	0.69	13.77	0.000	***
	EES(7)	0.23	0.50	7.10	0.008	***
	HP(30)	0.46	0.10	0.25	0.619	
EURO	TC(1,2)	0.15	0.70	12.83	0.000	***
	TC(2,2)	0.19	0.62	10.40	0.001	***
	EES(7)	0.27	0.46	5.57	0.018	**
	HP(30)	0.38	0.24	1.47	0.225	
US	TC(1,2)	0.19	0.55	10.64	0.001	***
	TC(2,2)	0.12	0.77	16.25	0.000	***
	EES(7)	0.38	0.33	4.54	0.033	**
	HP(30)	0.27	0.45	5.42	0.020	**

[†] *, **, ***: Significant at 10%, 5%, 1%

Table 4: Sign tests of real time cyclical components of real GDP

5 Testing the stochastic cycle model of the TC filter and adjusting it to the data

The real-time properties of filters depend to a large extent on how far the stochastic model underlying the filter matches the data generating process. In the previous section, we found that the TC filter seems to provide a comparatively good approximation of the data generating process that drives real GDP in some countries. This is astonishing as the filter was not at all adjusted to the data; rather, all parameters were assigned exogenously selected values.

In this section we show that the stochastic cycle model can to some extent be tested and that the length of the critical cycle in the TC filter can be endogenised and adjusted to the data in a way roughly consistent with the underlying stochastic model, which improves the forecasting capabilities and hence the real-time properties of the filter model. This is an important difference to the one-component filters: it is not possible to endogenise the

exogenous smoothing parameters of the HP filter and the EES.

The parameters μ and ρ of the stochastic cycle can be estimated by iterating through the following steps:

1. Compute the cycle with initial values for μ and ρ .
2. Estimate $\alpha(L)x_t^C = \beta(L)\nu_t$.
3. Compute μ and ρ implied by the estimated parameters of the lag-polynomials $\alpha(L)$ and $\beta(L)$ and use the new values for μ and ρ to compute the cycle with the TC filter again.
4. Go to step 2.

The iteration is to be broken up if changes in μ and ρ between two consecutive iterations fulfill a pre-defined condition.

The parameters ρ and μ are over-identified in step 3 by the parameters of the *ARMA*-regression. One option to deal with this problem is to perform a restricted regression. For instance, if $c = 2$ the four parameters a_i of the autoregressive lag polynomial $\alpha(L)$ and the two parameters b_i of the moving average lag polynomial $\beta(L)$ can be restricted as follows:

$$a_1 = -2w_1; a_2 = (w_1^2 + 2w_2); a_3 = -2w_1w_2; a_4 = w_2^2; b_1 = -w_1; b_2 = \frac{1}{4}w_1^2$$

This gives $\rho = \sqrt{w_2}$ and $\mu = \arccos(w_1/(2\sqrt{w_2}))$.

A similar iteration²⁸ to endogenise the smoothing parameter λ of the HP filter is not feasible. The value of λ would shrink to zero for typical economic time series implying that the inverse signal-to-noise-variance ratio—and hence λ —becomes small (Reeves et al. 1996; Ravn and Marcat 2003).

The results of the iterative estimation of μ and ρ for the TC(1,2) and the TC(2,2) filter are shown in Table 5. For the TC(2,2) filter, the prior critical length of the cycle of 8 years turns out to be largely consistent²⁹: the critical cycle length stabilises at a value close to the prior length. The largest deviations from the prior value occur for Spain with an estimated cycle length of 11.2 and Germany for which the procedure converges to a

²⁸ In the case of the HP filter, the iteration would comprise the steps

1. Set an initial values for λ .
2. Compute the trend.
3. Estimate the noise variance as $\hat{\sigma}_\varepsilon^2 = (X - X^T)'(X - X^T)$ and the variance of the signal as $\hat{\sigma}_\eta^2 = X^T \nabla^2 \nabla^2 X^T$ and set $\lambda = \frac{\hat{\sigma}_\varepsilon^2}{\hat{\sigma}_\eta^2}$.
4. Go to step 2 .

²⁹ The procedure stopped at iteration i if the condition $(\rho_i - \rho_{i-1})^2 + (\mu_i - \mu_{i-1})^2 \leq 0.001$ held. Furthermore, the values for μ obtained from the iterative TC filter do not depend on the initial value of eight years: the same results for the iteratively estimated critical cycle length follows from initial values of 4,5,9,12 or 13 years.

Country	Filter	Final parameters [†]		Regression parameters [‡]		Iterations
		Cycle Length	ρ	w_1	w_2	
DE	TC(1, 2)	10.6	0.88	1.454	0.769	26
	TC(2, 2)	7.1	0.89	1.115	0.78	12
ES	TC(1, 2)	20.5	0.96	1.839	0.930	9
	TC(2, 2)	11.2	0.96	1.618	0.913	10
FR	TC(1, 2)	10.52	0.92	1.523	0.848	7
	TC(2, 2)	8.6	0.91	1.358	0.830	14
IT	TC(1, 2)	8.54	0.90	1.340	0.817	7
	TC(2, 2)	7.8	0.90	1.247	0.808	5
EURO	TC(1, 2)	11.1	0.89	1.502	0.792	15
	TC(2, 2)	8.2	0.90	1.290	0.802	10
US	TC(1, 2)	14.0	0.81	1.457	0.648	9
	TC(2, 2)	7.5	0.88	1.174	0.779	13

[†] Start values: length of crit. cycle = 8 years and $\rho = 0.975$

[‡] Restricted regression:

$$x_t^C = 2w_1x_{t-1}^C - (w_1^2 + 2w_2)x_{t-2}^C + 2w_1w_2x_{t-3}^C - w_2^2x_{t-4}^C + \nu_t - w_1\nu_{t-1} + 0.25w_1^2\nu_{t-2}$$

Table 5: Iterative estimations of the critical cycle length and ρ for the TC(d , 2) filter

seven-years cycle. However, the differences between a cycle generated by a TC(2, 2) filter with a reference cycle of 8 years and one with a reference cycle of 7 or 11.2 years are not large.

The iterative TC(1,2) filter gives rise to implausibly long critical cycles in all countries except Italy. Obviously “cycles” of 10-21 years length seem to be important in the data generating process, and a first order random walk turns out too inflexible to match the spectrum of these long cycles. As a consequence, the assumption of a first order random walk induces the cycle spectrum to shift leftwards thereby incorporating lower frequencies into the cyclical component. In the case of Italy, the assumption of a first order stochastic trend is consistent with the existence of an eight to nine years reference cycle.

The residuals of the AR-regression of the cyclical component indicate the appropriateness of the stochastic cycle model. According to the model specification in equation (13), they should be white noise errors so that we do not expect to find sizeable amounts of autocorrelation in these residuals. Figure D.3 in Annex D shows the Ljung-Box Q-statistics for the residuals of the regressions for all combinations of first and second order trends and cycles, together with the critical values at the 5% significance level. The residuals of the first-order stochastic cycle models contain considerable amounts of autocorrelation which is inconsistent with the model specification. However, the TC filters with stochastic cycles of order 2 perform relatively well: the second order stochastic cycle model gives rise to weakly autocorrelated residuals only in the cases of Germany and the US³⁰.

³⁰ The first- and second-order autocorrelation coefficients for the residuals of the stochastic cycle regression are 0.32 and -0.29 in the case of Germany and 0.27 and -0.40 in the case of the US.

6 The *Trend-Cycle-Season* filter

The TC filter cannot be applied directly to time series with seasonal components in the high-frequency range of the spectrum. The reason is that the gain function of the trend is not perfect: Figure 2 shows that the gain obtains small positive values in the high-frequency range. While this error is negligible when annual time series are filtered, it poses a problem for data with seasonality at higher frequencies. Furthermore, with such data, one often wants to extract the seasonal component in a consistent way together with the cycle and the trend. Therefore, we first discuss some possibilities to account for seasonality. We proceed by deriving a general stochastic seasonal model which we use to extend the TC filter to a *Trend-Cycle-Season* (*TCS*) filter. Finally, the TCS filter is applied to quarterly real GDP data in Germany.

6.1 Integrating seasonal components in the TC filter: general remarks

Several approaches to account for seasonality are possible. First, the season could be modelled in the form of seasonal unit roots. It is straightforward to define the ∇^d matrix appropriately as $\nabla^d \equiv (I_s - L^s)^d$ to model seasonal unit roots at the seasonal frequency s ³¹. Formally, the same solution as in equation (17) would apply. Of course, it is not possible to obtain a distinct seasonal component in this way since the seasonal pattern becomes part of the trend.

Second, deterministic seasonal components can be included in the filter. A straightforward way to do so would be the inclusion of seasonal dummy vectors in the filter model—in the same way as the filter is amended by dummy vectors for structural breaks as explained in Annex C. However, the assumption of deterministic seasonal components is not always appropriate. Rather, the seasonal pattern may be stochastic and thus change over time.

Finally, the seasonal pattern can be modelled through additional stochastic cycles at seasonal frequencies (Harvey 1989). In order to model a seasonal pattern with frequency s , $s/2$ seasonal stochastic cycles would be necessary. For instance, a quarterly seasonal pattern requires two stochastic cycles, one for the two-years and one for the four-years cycle. Together with the stochastic trend and the stochastic cycle at business cycle frequency this would give four components. While solutions for multi-component filters can be developed in a generic way³², the solutions and the application of the filters become increasingly complicated if more and more components are included. This would overburden the filter technique which we want to keep as simple as possible.

6.2 A stochastic model for the seasonal component

Thus, it seems appropriate to restrict the filter approach to three components at most. Therefore, we follow Schlicht and Pauly (1983) in modelling a stochastic seasonal process with just one stochastic ARMA component. Indeed, they suggest a two-component filter

³¹ I_s denotes an identity-matrix in which the first s rows are filled with zeroes (see Annex A).

³² See Annex B.

for seasonal adjustment that is very similar to the TC filter: a second-order trend filter is amended with the seasonal component by minimising the sum of squared residuals of the stochastic trend and season models.

Schlicht and Pauly (1983) assume that a stable seasonal pattern in the seasonal component x_t^S is disturbed by stochastic shocks ν_t :

$$x_t^S = x_{t-s}^S + \nu_t. \quad (22)$$

The stochastic disturbance is modelled as a moving average process:

$$\nu_t = \xi_t - \frac{1}{s} \sum_{\tau=0}^{s-1} \xi_{t-\tau}, E(\xi_t) = 0. \quad (23)$$

Define $z_t \equiv \sum_{\tau=0}^{s-1} x_{t-\tau}^S$ and $\theta_t \equiv \sum_{\tau=0}^{s-2} \frac{s-1-\tau}{s} \xi_{t-\tau}$. Equations (22) and (23) give:

$$z_t - \theta_t = z_{t-1} - \theta_{t-1}.$$

As this has to hold for each t , it follows that

$$z_t - \theta_t \equiv \sum_{\tau=0}^{s-1} x_{t-\tau}^S - \sum_{\tau=0}^{s-2} \frac{s-1-\tau}{s} \xi_{t-\tau} = \text{constant}.$$

Since $z_t - \theta_t$ is non-stochastic, it can be interpreted as the moving seasonal sum of the seasonal component, from which the impact of the stochastic disturbances has been removed. This expression and hence the constant should be set equal to zero. Thus, we can write the stochastic model for the seasonal component as

$$\sum_{\tau=0}^{s-1} x_{t-\tau}^S = \sum_{\tau=0}^{s-2} \frac{s-1-\tau}{s} \xi_{t-\tau}.$$

In matrix form we obtain

$$PX^S = Q\xi, \quad (24)$$

where the $(N - s + 1) \times N$ matrices P and Q are defined as follows:

$$P = \begin{bmatrix} 1 & 1 & \dots & \overset{\text{column } s}{\downarrow} 1 & 0 & \dots & 0 & 0 \\ 0 & 1 & \dots & \dots & 1 & \dots & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & \dots & 1 & \dots & 1 & 1 \end{bmatrix} \quad Q = \frac{1}{s} \begin{bmatrix} 0 & 1 & 2 & \dots & s-1 & 0 & \dots & 0 \\ \vdots & 0 & 1 & 2 & \dots & s-1 & 0 & \dots & 0 \\ \vdots & \vdots \\ 0 & 0 & \dots & 0 & 1 & 2 & \dots & s-1 \end{bmatrix}.$$

We amend equation (13) with the stochastic seasonal model assuming that the covariance of the seasonal disturbance with all other disturbances is zero. This completes the stochastic model of the TCS filter. The trend, the cyclical, and the seasonal component are obtained by minimising the sum of the inner products of the residuals $\varepsilon'\varepsilon + \eta'\eta + \xi'\xi + \zeta'\zeta$ and are derived in Annex B. As the gain of the seasonal component in Figure 5 shows, the seasonal filter lets pass frequencies around the seasonal frequencies for quarterly data of $\frac{\pi}{2}$ and π . As the application of the TCS filter on quarterly real GDP data for Germany shows (see figure D.4 in Annex D), the TCS filter is capable of decomposing an economic time series into a trend, cycle and seasonal component in a plausible way.

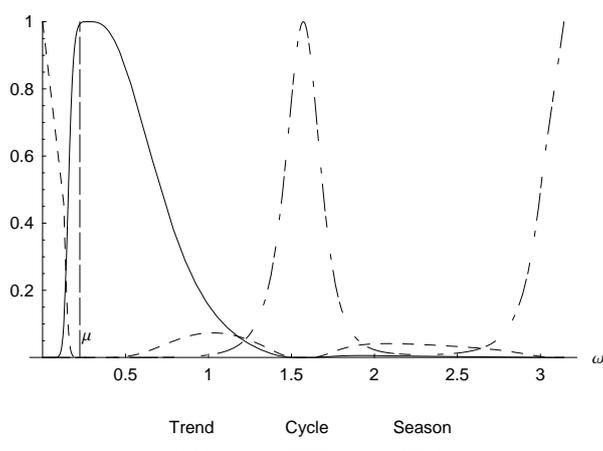


Figure 5: Squared gain functions of the TCS(1,2,28) filter with a quarterly season

7 Conclusion

Univariate trend-cycle decompositions suffer from all-too simple implicit models of the data generating process, while more elaborated approaches—as for instance unobserved components models—are not always easily applicable. This paper develops an intermediate approach by generalising the HP filter and incorporating a cyclical component into the model representation of the filter in the time domain. The resulting trend-cycle filter has better end-of-sample properties than the HP filter or the related Extended Exponential Smoothing (EES) procedure. In particular, the pro-cyclicality in end-of-sample trend/cycle estimations, characterising one-component filters such as the HP filter and the EES, is virtually removed.

The one-component filters are only based on an implicit model for the trend leaving the cycle as a residual from trend-extraction. The incorporation of a cycle model turned out crucial for the favourable properties of the TC filter. Furthermore, structural breaks or exogenous variables to identify the trend and the cyclical component can be easily incorporated in the TC filter. Finally, the Trend-Cycle filter can be expanded towards a Trend-Cycle-Season (TCS) filter in a straightforward way. With the TCS filter, a trend, cyclical and seasonal component can be simultaneously extracted from a time series.

Basic assumptions of the stochastic model underlying the TC filter can be tested, and the model can to some extent be adjusted to the data. As a consequence of the more flexible model-structure, the TC filter produces results, which are more model-consistent than those obtained with the one-component counterparts, the EES and the HP filter.

The TC filter with a second order stochastic trend and a second order stochastic reference cycle of eight years delivers plausible results for all the cases analysed here and can therefore be regarded as an appropriate reference model. However, it is not optimal for all cases. We found that for Spanish GDP the properties of the trend and the cyclical component improve by choosing a nine-years reference cycle. In the case of Italian GDP, the first order stochastic trend gives results more consistent with theoretical model than the second order trend. Different from the TC filter, trend estimates of one-component

filters are largely inconsistent with their underlying stochastic models.

While the TC filter is based on a more complex stochastic model than the EES and the HP filter, its application is comparatively simple. Once the TC filter has been programmed³³, it is straightforward to choose the appropriate stochastic trend and cycle models and to obtain the trend-cycle decomposition. It is not necessary to experiment with prior variance restrictions and start values for unobserved variables as it is sometimes required in unobserved components model estimations.

³³ An implementation in EVIEWS 4.x can be obtained from the author upon request.

Annex A

Lag and difference operators in matrix form

Define the $N \times N$ lag matrix L as $L = \begin{bmatrix} 0 & 0 & \dots & \dots & 0 \\ 1 & 0 & 0 & \dots & 0 \\ 0 & 1 & 0 & 0 & \dots & 0 \\ \vdots & & & & & \vdots \\ 0 & \dots & 1 & \dots & 0 \\ 0 & \dots & 0 & 1 & 0 \end{bmatrix}$. The first row of L is zero as

in finite samples the d -th lag is not defined for the first d data points. This makes some adaptations to the usual lag- and difference operators necessary. Most of their properties, however, carry over to their matrix representations. Lag and difference matrices have the following properties:

Property 1: The d -th lag in matrix form is defined as $L^d = LL^{d-1}$. It holds that $L^d = L^q L^{d-q}$, for any q , $0 \leq q \leq d$. For completeness define $L^0 \equiv I$.

Property 2: The lead operator in matrix form is equal to the transpose of L , L' .

Property 3: Denote an $N \times N$ identity-matrix in which the first d rows are filled with zeroes as I_d . Then, $LL' = I_1$ holds. In general, $L^d L^{d'} = I_d$. Furthermore, it holds that $I'_d = I_d$. For any pair (n, m) , with $n \geq m$, $I_n I_m = I_n$ holds.

Property 4: The matrix of first differences ∇ can be defined as $\nabla \equiv I_1(I - L)$. The I_1 -matrix renders the first row of ∇ zero, accounting for the fact that the lag of the first data point is not defined. In general we define the d -th difference matrix as $\nabla^d \equiv I_d \nabla \nabla^{d-1}$. Again, this is the same as $\nabla \nabla^{d-1}$ with the first d rows set equal to zero as the d -th lag is not defined for the first d data points. It holds that $\nabla^d = I_d \nabla^q \nabla^{d-q}$, for any q , $0 \leq q \leq d$. For completeness we define $\nabla^0 \equiv I$.

Property 5:

$$\nabla^{d'} = \begin{cases} L^{d'} \nabla^d & \text{if } d \text{ is even} \\ -L^{d'} \nabla^d & \text{if } d \text{ is odd.} \end{cases}$$

Proof:

$$\begin{aligned} \nabla^{d'} &= I_d \underbrace{(I_1 (I - L))' \dots (I_1 (I - L))'}_{d \times} = I_d \underbrace{(I_1 - L')' \dots (I_1 - L')'}_{d \times} \\ &= I_d \underbrace{(I'_1 - L') \dots (I'_1 - L')}_{d \times} = I_d \underbrace{(L' (L - I)) \dots (L' (L - I))}_{d \times} \\ &= I_d L^{d'} \underbrace{(L - I) \dots (L - I)}_{d \times} \\ &= \begin{cases} L^{d'} I_d \underbrace{(I - L) \dots (I - L)}_{d \times}, & \text{if } d \text{ is even} \\ -L^{d'} I_d \underbrace{(I - L) \dots (I - L)}_{d \times} & \text{if } d \text{ is odd} \end{cases} \end{aligned}$$

$$= \begin{cases} L^{d'} \nabla^d & \text{if } d \text{ is even} \\ -L^{d'} \nabla^d & \text{if } d \text{ is odd.} \end{cases}$$

Annex B

The generic k -filter

As equation (17) shows, the solution for the two-components TC filter can be expressed in terms of the solution of the one-component trend and cycle filters. This can be generalised: the solution of the system of FOCs for a filter with k -components, the k -filter, can be expressed in terms of the solution for the $(k - 1)$ -filter. Let \mathfrak{K} denote the set of all k components and $\mathfrak{K} \setminus \{\mathbf{i}\}$ denote the set of all k components except component \mathbf{i} . Furthermore, we define the filtermatrix for component $\mathbf{i} \in \mathfrak{K}$ given the set of the other $(k - 1)$ components $\mathfrak{K} \setminus \{\mathbf{i}\}$ as $M_{\mathbf{i}, \mathfrak{K} \setminus \{\mathbf{i}\}}$ so that $X^{\mathbf{i}} = M_{\mathbf{i}, \mathfrak{K} \setminus \{\mathbf{i}\}} X$. Finally, the filter matrix of an 1-filter with component \mathbf{i} is defined as $M_{\mathbf{i}}$. With these definitions, the system of FOCs for a k -filter can be written as

$$X^{\mathbf{i}} = M_{\mathbf{i}, \mathfrak{K} \setminus \{\mathbf{i}, \mathbf{j}\}} (X - X^{\mathbf{j}}), \quad \forall \mathbf{i}, \mathbf{j} \in \mathfrak{K}, \mathbf{i} \neq \mathbf{j} \quad (\text{B.1})$$

which gives the generic solution

$$\begin{aligned} M_{\mathbf{i}, \mathfrak{K} \setminus \{\mathbf{i}\}} &= (I - M_{\mathbf{i}, \mathfrak{K} \setminus \{\mathbf{i}, \mathbf{j}\}} M_{\mathbf{j}, \mathfrak{K} \setminus \{\mathbf{i}, \mathbf{j}\}})^{-1} M_{\mathbf{i}, \mathfrak{K} \setminus \{\mathbf{i}, \mathbf{j}\}} (I - M_{\mathbf{j}, \mathfrak{K} \setminus \{\mathbf{i}, \mathbf{j}\}}) \\ X^{\mathbf{i}} &= M_{\mathbf{i}, \mathfrak{K} \setminus \{\mathbf{i}\}} X, \\ \forall \mathbf{i}, \mathbf{j} &\in \mathfrak{K}, \mathbf{i} \neq \mathbf{j} \end{aligned} \quad (\text{B.2})$$

The necessary and sufficient condition for equation (B.2) to hold is that for each component $\mathbf{i} \in \mathfrak{K}$ the respective 1-filter exists:

$$\det(M_{\mathbf{i}}^{-1}) \neq 0, \quad \forall \mathbf{i} \in \mathfrak{K} \quad (\text{B.3})$$

The generic solution can be applied to obtain the solution for the *Trend-Cycle-Season* filter. To easen notation, define $M_S \equiv (I + P'(QQ')^{-1}P)^{-1}$ where the matrices P and Q are defined as on page 30. Furthermore, we make use of the notation in equation (B.1) to express the trend of the two-components TC filter, for instance, as $X^T = M_{TC} X$. We can now conveniently derive the solutions for the three-components TCS filter from the two-components filter solutions:

$$\begin{aligned} X^T &= (I - M_{TC} M_{SC})^{-1} M_{TC} (I - M_{SC}) X \Leftrightarrow X^T = M_{TCS} X \\ X^C &= (I - M_{CT} M_{ST})^{-1} M_{CT} (I - M_{ST}) X \Leftrightarrow X^C = M_{CTS} X \\ X^S &= (I - M_{ST} M_{CT})^{-1} M_{ST} (I - M_{CT}) X \Leftrightarrow X^S = M_{STC} X \end{aligned} \quad (\text{B.4})$$

In order to explain the rationale behind equation (B.4), observe for instance that the expression $(I - M_{ST}) X$ in the solution for the cyclical component refers to the difference

between the original time series X and the seasonal process we would obtain with a two-components Trend-Season filter applied to X ³⁴. Applying the filter matrix of the cyclical component of the TC filter to the residual from the TS filter gives $M_{CT}(I - M_{ST})X$. Finally, the error, which derives from the fact that the stepwise application of two 2-filters ignores the simultaneity in the determination of the components of X , has to be corrected. Hence, the stepwise filter is multiplied with the correction factor $(I - M_{CT}M_{ST})^{-1}$. In the special case of $M_{CT}M_{ST} = \mathbf{0}$, the TCS filter gives the same solution as the stepwise application of two 2-filters. For X^C , for instance, this condition would hold if the cyclical and the seasonal component of the TC and the TS filter were independent from each other in the sense that their gain functions do not intersect. In the general case, however, this condition does not hold and the simultaneous TCS filter solution differs from the stepwise application of two 2-filters. Finally, note that the necessary condition in equation (B.3) for the existence of a solution holds for the TCS as well as for the TC filter since the matrices $\nabla^{d'}\nabla^d$, $A'(BB')^{-1}A$ and $P'(QQ')^{-1}P$ are non-singular³⁵.

Annex C

Structural breaks

The HP filter, the EES, and the standard TC filter are based on the assumption of a smooth trend without structural breaks. Sometimes, this assumption does not hold. A prominent counter-example is the German unification which gave rise to an upward level shift in German macroeconomic time series. Furthermore, statistical revisions such as the switch from the ESA 79 to the ESA 95 system of national accounts give rise to a structural breaks.

Applying the HP, the EES or the TC filter to a series with a structural break leads to biased trend estimations around the break period. This is because the methods smoothen out the break so that the trend is too high immediately before the break and too low in the periods immediately thereafter.

If the period in which the break occurred is known beforehand, the break can be incorporated in the trend model of the TC filter³⁶ by assuming that X^T follows a purely stochastic trend once the deterministic break has been removed:

$$\nabla^d(X^T - Dv) = \eta \quad (\text{C.1})$$

³⁴ Such a filter, which is conceptually very similar to the TC filter, is suggested by Schlicht and Pauly (1983) as a seasonal adjustment method.

³⁵ Equation (B.4) gives consistent results if one component is missing. For instance, assume that there is a trend and a cyclical but no seasonal component, implying $M_S = \mathbf{0}$. One can derive $M_{ST} = \mathbf{0}$ and $M_{SC} = \mathbf{0}$ from the formulas for the two-components filters. It follows that the three-components filters for the trend and the cyclical component collapse to two-components filters: $M_{TCS} = M_{TS}$ and $M_{CTS} = M_{CT}$. Furthermore, $M_{STC} = \mathbf{0}$. This reasoning holds correspondingly when the trend component or the cyclical component is absent.

³⁶ This appendix draws on Tödter (2002), who expands the EES by structural breaks.

where v denotes a scalar capturing the size of the break and D denotes a dummy vector defining the timing of the break in period τ :

$$D_t = \begin{cases} 0 & \text{if } t < \tau \\ 1 & \text{if } t \geq \tau. \end{cases} \quad (\text{C.2})$$

We obtain the following FOCs for the trend and the cyclical process:

$$\begin{aligned} (I + \nabla^{d'} \nabla^d) X^T &= (X - X^C + \nabla^{d'} \nabla^d D v) \\ (I + A' A) X^C &= (X - X^T) \\ v &= (D' \nabla^{d'} \nabla^d D)^{-1} D' \nabla^{d'} \nabla^d X^T \end{aligned} \quad (\text{C.3})$$

The solution for v is simply the estimated parameter of a regression of $\nabla^d X^T$ on $\nabla^d D$.

To explain the solution for the break parameter v , assume that $d = 2$. It can be shown that the expression $(D' \nabla^{2'} \nabla^2 D)^{-1}$ amounts to $\frac{1}{2}$. Furthermore, $D' \nabla^{2'} \nabla^2$ can be shown to be equal to $-\Delta^3 x_{\tau+1}^T \equiv 2\Delta x_{\tau}^T - (\Delta x_{\tau+1}^T + \Delta x_{\tau-1}^T)$. Thus, v is computed as

$$v = \Delta x_{\tau}^T - \frac{1}{2}(\Delta x_{\tau+1}^T + \Delta x_{\tau-1}^T) \quad (\text{C.4})$$

which is the change in the trend in period τ when the break occurs minus the average trend change immediately before and after the break period. Thus, the parameter v can be understood as locally correcting the bias around the break period τ that the standard TC filter without a structural break in the trend equation would induce³⁷.

The solutions for the trend and the cycle can be written in a convenient way using the fundamental regression matrix W for the regression of $\nabla^d X^T$ on $\nabla^d D$:

$$W \equiv \nabla^d D (D' \nabla^{d'} \nabla^d D)^{-1} D' \nabla^{d'}$$

We obtain $I - W$, the residual projection matrix of the regression on the dummy vector $\nabla^d D$. The solution for the trend and the cycle can be obtained by replacing $\nabla^{d'} \nabla^d$ in equation (13) with $\nabla^{d'} (I - W) \nabla^d$. This leads to the FOCs in which the structural break parameter is eliminated:

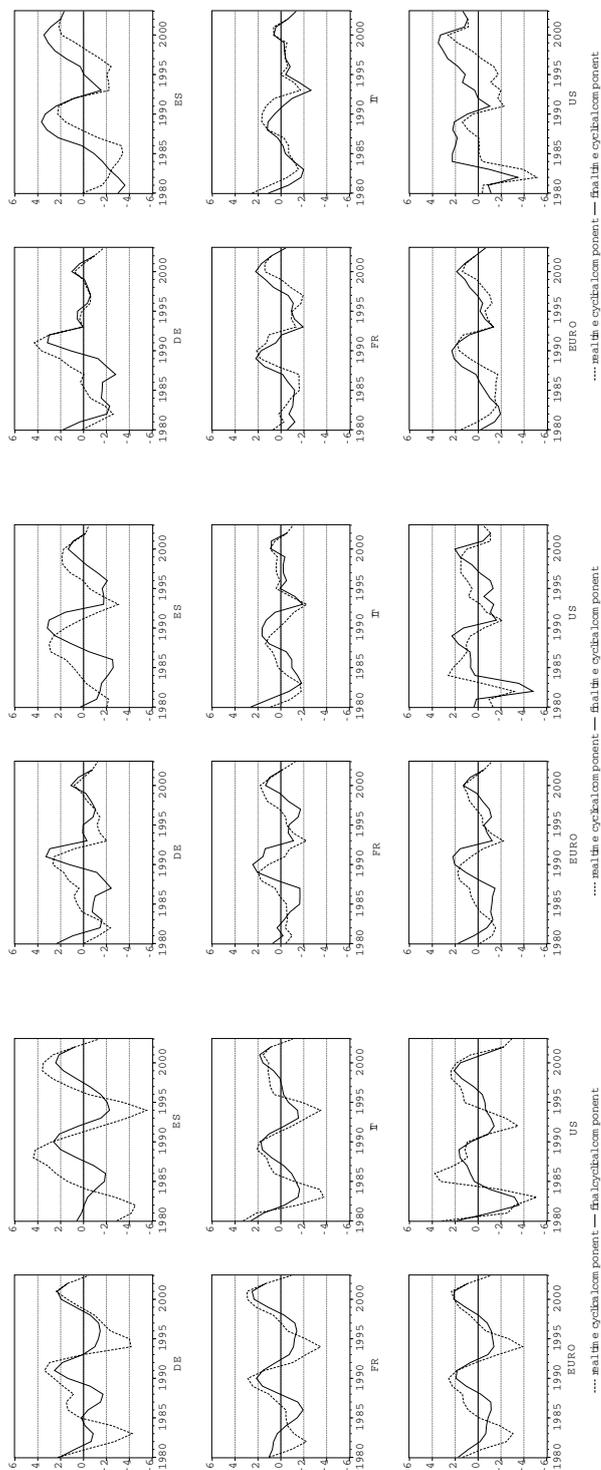
$$\begin{aligned} X^T &= (I + \nabla^{d'} (I - W) \nabla^d)^{-1} (X - X^C) \\ X^C &= (I + A' A)^{-1} (X - X^T). \end{aligned} \quad (\text{C.5})$$

Furthermore, it is straightforward to extend this approach by incorporating more than just one structural break: D could be an $N \times r$ matrix composed of r appropriately defined and linearly independent dummy vectors, and v would then be an $r \times 1$ vector containing the breaks. Besides shift dummies, which introduce permanent level shifts in the trend, the vectors in D could likewise be specified as impulse dummies representing temporary jumps.

³⁷ The solution for a first order random walk with drift ($d = 1$) is slightly more complicated as the constant drift term b must be estimated in addition to the break parameter. It can be shown that $v = x_{\tau}^T - x_{\tau-1}^T - b$ and $b = \frac{x_{\tau}^T - x_{\tau-1}^T + x_{\tau-1}^T - x_1^T}{T-2}$ hold in this case: the shift parameter v is estimated as the trend change in the break period τ , corrected by the drift term whereas the drift term is computed as the global trend change excepting the break period (Tödter 2002).

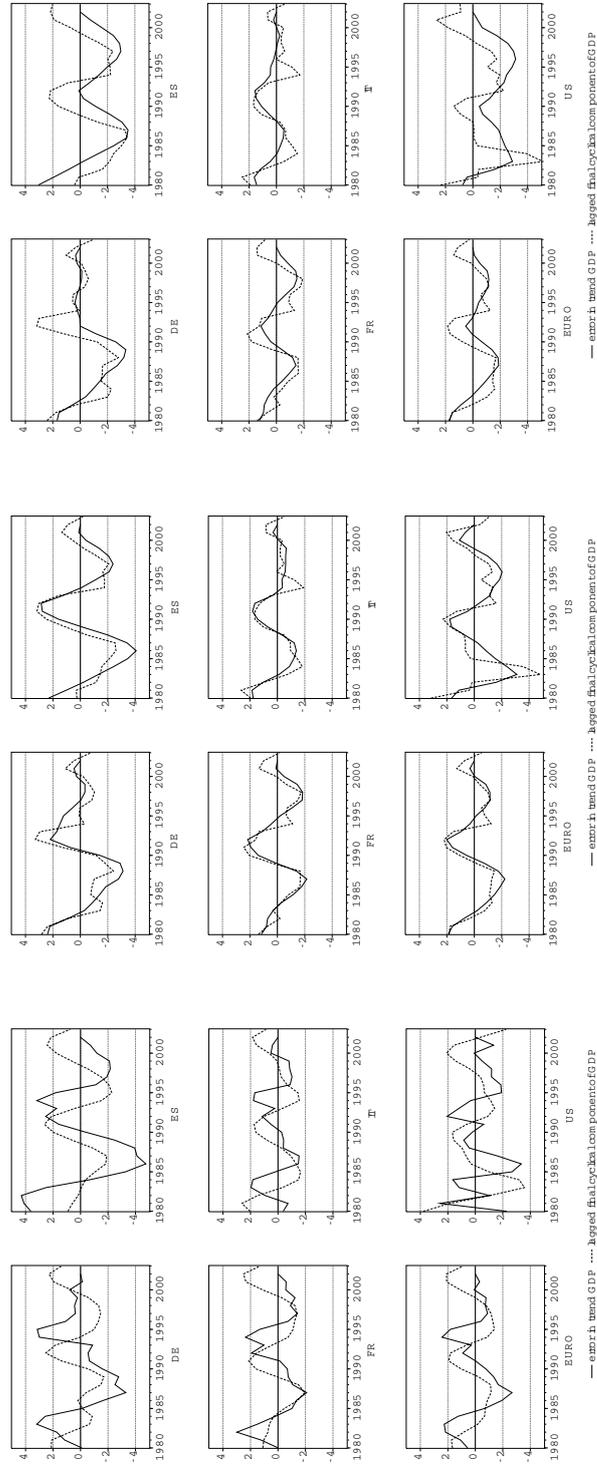
Annex D

Additional figures



(a) TC(2,2,8) filter
 (b) HP(30) filter
 (c) EES(7)

Figure D.1: Cyclical components and real-time cyclical components of real GDP (in % of real GDP)



(a) TC(2,2,8) filter

(b) HP(30) filter

(c) EES(7)

Figure D.2: Real-time minus final trend and final cyclical component (in % of real GDP)

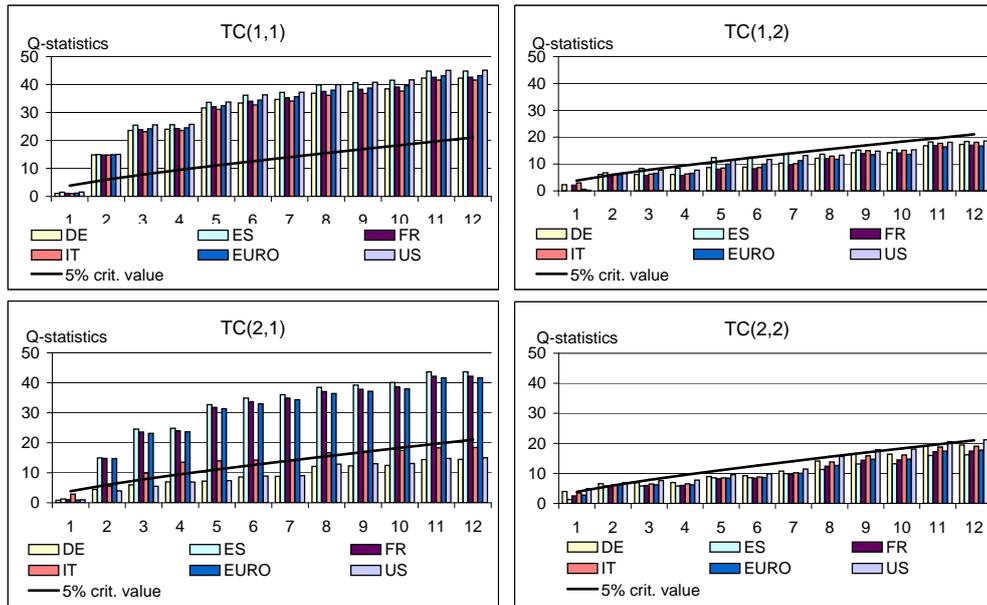


Figure D.3: Q-statistics for the residuals obtained from AR regressions of the cyclical components

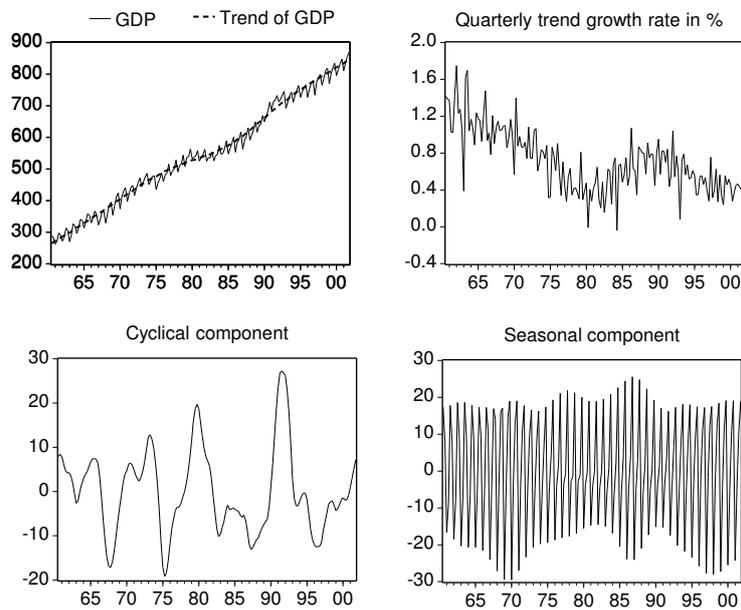


Figure D.4: Application of the TCS filter on quarterly real GDP in Germany

References

- Baxter, M. and R. King (1999). Measuring business cycles: Approximate bandpass filters for economic time series. *Review of Economics and Statistics* 81, 575–593.
- Boone, L., C. Giorno, M. Meacci, D. Rae, P. Richardson, and D. Turner (2000). The concept, policy use and measurement of structural unemployment: Estimating a time varying nairu across 21 OECD countries. OECD Working Paper ECO/WKP(2000)23.
- Bouthevillain, C., P. Cour-Thimann, G. van den Dool, P. H. de Cos, G. Langenus, M. Mohr, S. Momigliano, and M. Tujula (2001). Cyclically adjusted budget balances: An alternative approach. European Central Bank Working Paper 77. <http://www.ecb.int/pub/pdf/scpwps/ecbwp077.pdf>.
- Chagny, O. and J. Döpke (2002). Measures of the output gap in the euro-zone: An empirical assessment of selected methods. *Quarterly Journal of Economic Research – Vierteljahrshefte zur Wirtschaftsforschung* 70, 310–330.
- Cogley, T. and J. N. Nason (1995). Effects of the Hodrick–Prescott filter on trend and difference stationary time series. Implications for business cycle research. *Journal of Economic Dynamics and Control* 19, 253–278.
- Danthine, J.-P. and M. Girardin (1989). Business cycles in Switzerland. A comparative study. *European Economic Review* 33, 31–50.
- Denis, C., K. M. Morrow, and W. Röger (2002). Production function approach to calculating potential growth and output gaps - estimates for the EU member states and the US. European Commission Economic Papers No. 176.
- Döpke, J. (2004). Real-time data and business cycle analysis in Germany: Some preliminary results. Deutsche Bundesbank, Volkswirtschaftliches Forschungszentrum, Working Paper 11/2004.
- Dupasquier, C., A. Guay, and P. St-Amant (1997). A comparison of alternative methodologies for estimating potential output and the output gap. Bank of Canada Working Paper 97-5.
- Gomez, 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.
- Gruen, D., T. Robinson, and A. Stone (2002). Output gaps in real time: Are they reliable enough to use for monetary policy? Reserve Bank of Australia, Economic Research Department Research Discussion Paper 2002-06.
- Harvey, A. (1989). *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press.
- Harvey, A. and A. Jäger (1993). Detrending, stylised facts and the business cycle. *Journal of Applied Econometrics* 8, 231–247.
- Harvey, A. C. and T. M. Trimbur (2003). General model-based filters for extracting cycles and trends in economic time series. *Review of Economics and Statistics* 85, 244–255.

- Hodrick, R. J. and E. C. Prescott (1997). Postwar US business cycles: An empirical investigation. *Journal of Money, Credit, and Banking* 29, 1–16.
- Jäger, A. (1994). Mechanical detrending by Hodrick–Prescott filtering: A note. *Empirical Economics* 19, 493–500.
- Kaiser, R. and A. Maravall (1999). Estimation of the business cycle: A modified Hodrick–Prescott filter. *Spanish Economic Review* 1, 175–200.
- Laxton, D. and R. Tetlow (1992). A simple multivariate filter for the measurement of potential output. Bank of Canada Technical Report No 59.
- Lucas, R. E. (1980). Two illustrations of the quantity theory of money. *American Economic Review* 70, 1005–1370.
- Maddala, G. (1977). *Econometrics*. MacGraw-Hill.
- Mohr, M. (2001). Ein disaggregierter Ansatz zur Berechnung konjunkturbereinigter Budgetsalden für Deutschland: Methoden und Ergebnisse. Deutsche Bundesbank, Volkswirtschaftliches Forschungszentrum, Working Paper 13/2001.
- Pedersen, T. M. (2001). The Hodrick–Prescott filter, the Slutsky effect, and the distortionary effect of filters. *Journal of Economic Dynamics and Control* 25, 1081–1101.
- Planas, C. and A. Rossi (2004). Can inflation data improve the real-time reliability of output gap estimates? *Journal of Applied Econometrics* 19, 121–133.
- Priestley, M. B. (1981). *Spectral Analysis and Time Series*. London: Academic Press.
- Ravn, M. and A. Marcat (2003). The HP-filter in cross-country comparisons. CEPR Working paper 4244.
- Ravn, M. and H. Uhlig (2002). On adjusting the HP-filter for the frequency of observations. *Review of Economics and Statistics* 84, 371–380.
- Reeves, J. J., C. A. Blyth, C. M. Triggs, and J. P. Small (1996). The Hodrick–Prescott filter, a generalisation, and a new procedure for extracting an empirical cycle from the series. University of Auckland, Department of Economics Working Paper 163.
- Rünstler, G. (2002). The information content of real-time output gap estimates: An application to the euro area. European Central Bank Working Paper 182.
- Schlicht, E. and R. Pauly (1983). Descriptive seasonal adjustment by minimizing perturbations. *Empirica* 1, 15–28.
- Tödter, K.-H. (2002). Exponential smoothing as an alternative to the Hodrick–Prescott filter? In I. Klein and S. Mitnik (Eds.), *Contributions to Modern Econometrics: From Data Analysis to Economic Policy. In Honour of Gerd Hansen*, pp. 223–237. Kluwer.
- Whittle, P. (1983). *Prediction and Regulation by Linear Least-Square Methods, 2nd edition*. Minneapolis.

European Central Bank working paper series

For a complete list of Working Papers published by the ECB, please visit the ECB's website (<http://www.ecb.int>)

- 466 "Regulated and services' prices and inflation persistence" by P. Lünemann and T. Y. Mathä, April 2005.
- 467 "Socio-economic development and fiscal policy: lessons from the cohesion countries for the new member states" by A. N. Mehrotra and T. A. Peltonen, April 2005.
- 468 "Endogeneities of optimum currency areas: what brings countries sharing a single currency closer together?" by P. De Grauwe and F. P. Mongelli, April 2005.
- 469 "Money and prices in models of bounded rationality in high inflation economies" by A. Marcet and J. P. Nicolini, April 2005.
- 470 "Structural filters for monetary analysis: the inflationary movements of money in the euro area" by A. Bruggeman, G. Camba-Méndez, B. Fischer and J. Sousa, April 2005.
- 471 "Real wages and local unemployment in the euro area" by A. Sanz de Galdeano and J. Turunen, April 2005.
- 472 "Yield curve prediction for the strategic investor" by C. Bernadell, J. Coche and K. Nyholm, April 2005.
- 473 "Fiscal consolidations in the Central and Eastern European countries" by A. Afonso, C. Nickel and P. Rother, April 2005.
- 474 "Calvo pricing and imperfect common knowledge: a forward looking model of rational inflation inertia" by K. P. Nimark, April 2005.
- 475 "Monetary policy analysis with potentially misspecified models" by M. Del Negro and F. Schorfheide, April 2005.
- 476 "Monetary policy with judgment: forecast targeting" by L. E. O. Svensson, April 2005.
- 477 "Parameter misspecification and robust monetary policy rules" by C. E. Walsh, April 2005.
- 478 "The conquest of U.S. inflation: learning and robustness to model uncertainty" by T. Cogley and T. J. Sargent, April 2005.
- 479 "The performance and robustness of interest-rate rules in models of the euro area" by R. Adalid, G. Coenen, P. McAdam and S. Siviero, April 2005.
- 480 "Insurance policies for monetary policy in the euro area" by K. Küster and V. Wieland, April 2005.
- 481 "Output and inflation responses to credit shocks: are there threshold effects in the euro area?" by A. Calza and J. Sousa, April 2005.
- 482 "Forecasting macroeconomic variables for the new member states of the European Union" by A. Banerjee, M. Marcellino and I. Masten, May 2005.

- 483 “Money supply and the implementation of interest rate targets” by A. Schabert, May 2005.
- 484 “Fiscal federalism and public inputs provision: vertical externalities matter” by D. Martínez-López, May 2005.
- 485 “Corporate investment and cash flow sensitivity: what drives the relationship?” by P. Mizen and P. Vermeulen, May 2005.
- 486 “What drives productivity growth in the new EU member states? The case of Poland” by M. Kolasa, May 2005.
- 487 “Computing second-order-accurate solutions for rational expectation models using linear solution methods” by G. Lombardo and A. Sutherland, May 2005.
- 488 “Communication and decision-making by central bank committees: different strategies, same effectiveness?” by M. Ehrmann and M. Fratzscher, May 2005.
- 489 “Persistence and nominal inertia in a generalized Taylor economy: how longer contracts dominate shorter contracts” by H. Dixon and E. Kara, May 2005.
- 490 “Unions, wage setting and monetary policy uncertainty” by H. P. Grüner, B. Hayo and C. Hefeker, June 2005.
- 491 “On the fit and forecasting performance of New-Keynesian models” by M. Del Negro, F. Schorfheide, F. Smets and R. Wouters, June 2005.
- 492 “Experimental evidence on the persistence of output and inflation” by K. Adam, June 2005.
- 493 “Optimal research in financial markets with heterogeneous private information: a rational expectations model” by K. Tinn, June 2005.
- 494 “Cross-country efficiency of secondary education provision: a semi-parametric analysis with non-discretionary inputs” by A. Afonso and M. St. Aubyn, June 2005.
- 495 “Measuring inflation persistence: a structural time series approach” by M. Dossche and G. Everaert, June 2005.
- 496 “Estimates of the open economy New Keynesian Phillips curve for euro area countries” by F. Rumler, June 2005.
- 497 “Early-warning tools to forecast general government deficit in the euro area: the role of intra-annual fiscal indicators” by J. J. Pérez, June 2005.
- 498 “Financial integration and entrepreneurial activity: evidence from foreign bank entry in emerging markets” by M. Giannetti and S. Ongena, June 2005.
- 499 “A trend-cycle(-season) filter” by M. Mohr, July 2005.

