WORKING PAPERS AND STUDIES

ISSN 1725-4825

2006 EDITION

Real-Time Signal Estimation: an Application of "Customized" Optimization Criteria

Conference on seasonality, seasonal adjustment and their implications for short-term analysis and forecasting

Luxembourg, 10-12 May 2006

EUROPEAN COMMISSION

EURO INDICATORS

eurostat

THEME
General and regional statistics

Europe Direct is a service to help you find answers to your questions about the European Union

Freephone number (*):

00 800 6 7 8 9 10 11

(*) Certain mobile telephone operators do not allow access to 00-800 numbers or these calls may be billed.

A great deal of additional information on the European Union is available on the Internet.
It can be accessed through the Europa server (http://europa.eu).

Luxembourg: Office for Official Publications of the European Communities, 2006

ISBN 92-79-03415-4
ISSN 1725-4825
Catalogue number: KS-DT-06-016-EN-N

Real-Time Signal Extraction: an Application of 'Customized' Optimization Criteria

By Marc Wildi[1]

Institute of Data Analysis and Process Design, University of Applied Sciences, Winterthur and Institute for Business Cycle Research, KOF-ETH, Zurich.

The economic barometer issued by the institute of business cycle research is a typical application of real-time signal extraction: designers (and users) of the indicator are interested in estimating the level of the GDP growth-rate accurately towards the sample end-point $t = T$ and in detecting turning-points fast and reliably. Typically, a strong prospective perspective is associated to this discipline. Traditional model-based approaches (TRAMO/SEATS, Census X-12-ARIMA, STAMP to cite the most renowned) rely either explicitly or implicitly on forecasts of the original time series in order to compute real-time estimates. However, the underlying estimation problems magnify mis-specification issues that are ignored in the methodological framework of model-based approaches. The following paper illustrates the relevant issues and proposes a procedure based on 'customized criteria' which mate the very structure of real-time estimation problems. It is shown that the best mean-square level approximation and the detection of turning-points in real-time are incongruent requirements which deserve specific solutions. Empirical comparisons for two economic indicators - the economic barometer of the KOF and the European Sentiment Indicator (ESI) - illustrate efficiency gains obtained by the new approach when compared to widely-used ARIMA-based approaches as well as logit-models.

KEYWORDS: Real-time signal extraction, economic indicator, amplitude and time shift, customized criteria.

JEL CLASSIFICATION: C13, C53, E32

1 Introduction

Estimating trends or detecting turning-points in real-time are of great importance in various application fields. The present work is based on common research projects of the author with the KOF-ETH (www.kof.ethz.ch) - which is financed by the Swiss National Science Foundation - and with a major Swiss bank.

In order to illustrate issues involved in real-time estimation problems the 'old' economic barometer of the KOF is plotted together with the quarterly (Swiss) GDP growth-rate in fig.1. The indicator is obtained by aggregating time series supposed to anticipate the economic growth such as, for example, selected business survey data. Unfortunately, these series are generally contaminated by noise as well as seasonal components which impair an assessment of the immediate economic 'state'. The latter is commonly reflected by a

[1]Marc Wildi, Institute of Data Analysis and Process Design, ZHW, tel. (32) 345 12 71, e-mail wia@zhwin.ch

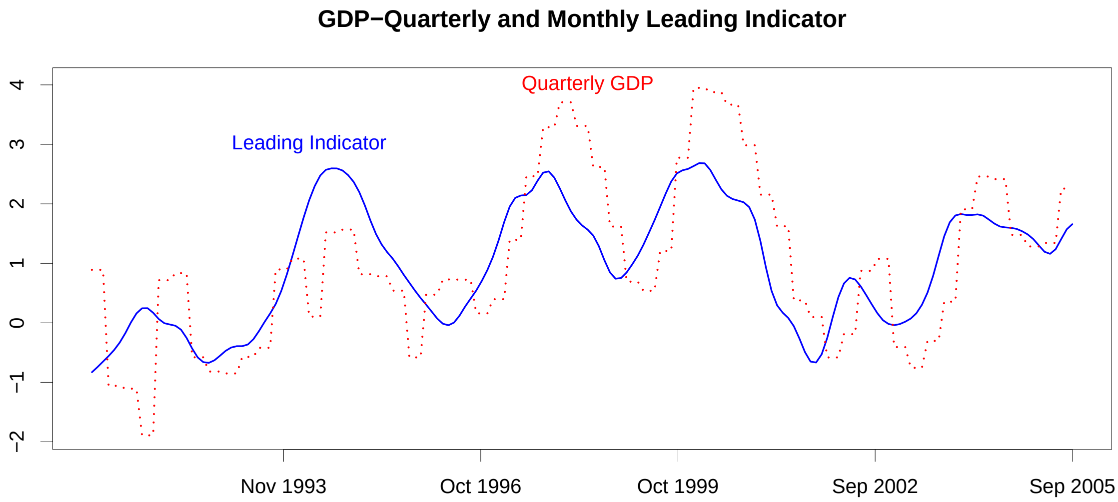

Figure 1: GDP quarterly (dotted line), leading indicator (solid)

particular trend signal which is to be estimated (extracted) at the current boundary $t = T$. The 'raw' indicator is compared to the final (smoothed) indicator in fig.2[2]. In the middle

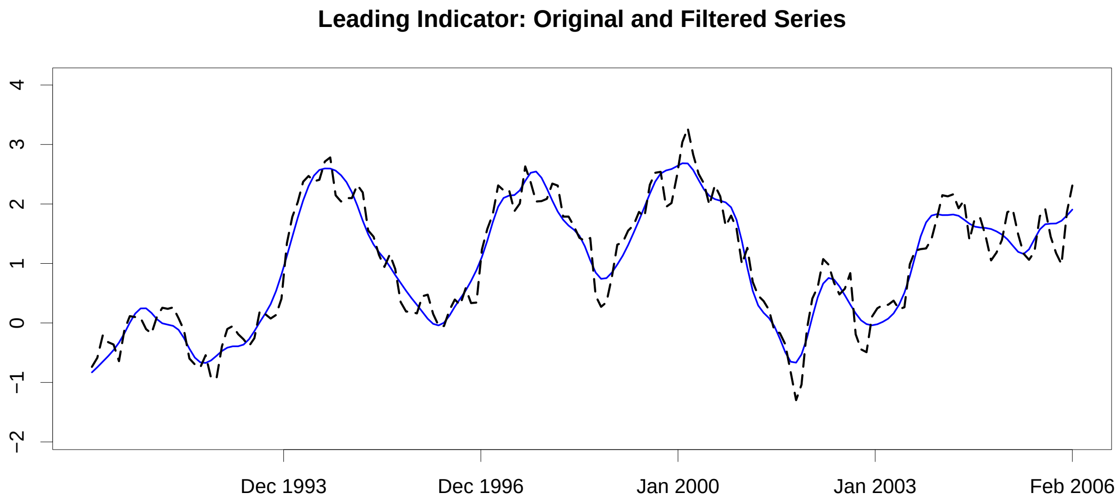

Figure 2: Leading indicator (solid), unfiltered indicator (shaded)

of the sample, i.e. in $t = T/2$, symmetric filters can be used in order to extract the desired trend component[3]. However, towards the current boundary of a finite sample,

[2]The time series in fig.1 is shorter because the GDP-series is not available towards the current boundary. Values for 2005 are still subject to revisions and the last asserted turning-point is in spring 2004.

[3]Symmetric filters are not subject to time shifts so that input series and output signal are synchronized.

symmetric filters have necessarily to be approximated by asymmetric designs. Therefore, current estimates generally differ from the intended final signal. Very often, in practice, current estimates are therefore revised when new data flows in. This effect is illustrated in fig.3 where estimates obtained in real-time (solid lines) are compared to the final X-12-ARIMA trend (dotted) in the vicinity of the last asserted turning-point of the original GDP-series (see footnote 2). One can see that the real-time series differs noticeably from

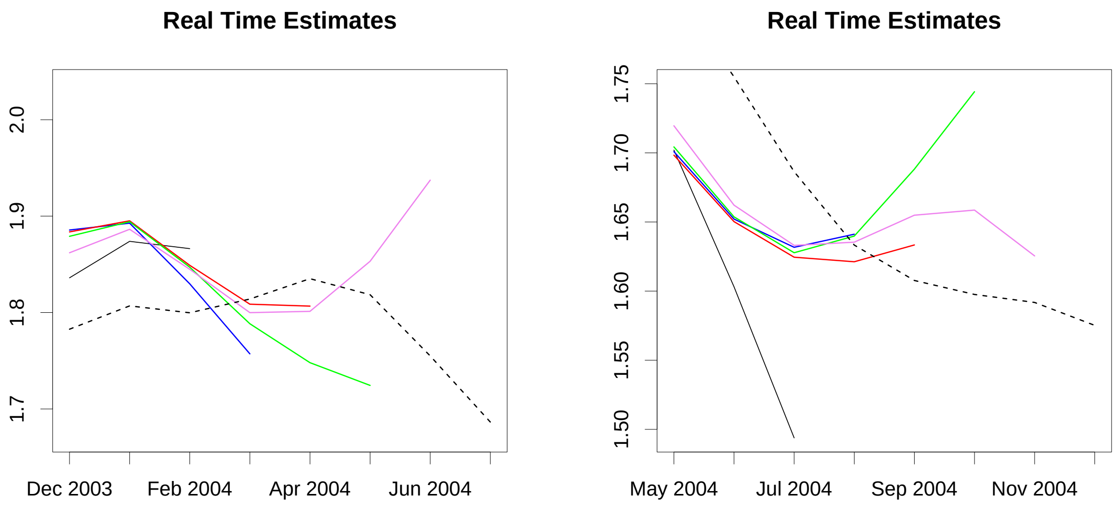

Figure 3: Trend (dotted) and real-time estimates (solid lines)

the ex-post (final) signal. Obviously, identifying turning-points is a challenging exercise because real-time estimates are noisy and delayed. *In this example, one can provide empirical evidence that the ex-post filter leads the GDP growth rate whereas the real-time filter lags it.* This subtle difference is often difficult to communicate to unexperienced users which are easily mystified by graphical representations like fig.1 because the ex-post area overwhelms the tiny but determinant current boundary.

The 'new' KOF economic barometer[4], released in April 2006, is based on a real-time filter which accounts specifically for the problem at hand: the mean-delay vanishes in the pass-band and the undesirable components are strongly damped, see section 4. As a result, turning-points are identified fast and reliably.

[4]See http://www.kof.ethz.ch/deutsch/konjunktur.

2 Model-Based Approaches

2.1 ARIMA-Model Based Approach and Real-Time Estimation Problem

Suppose the stochastic processes Y_t and X_t are related by

$$Y_t = \sum_{|k|<\infty} \gamma_k X_{t-k}$$

where $\gamma_k = \gamma_{-k}$ is a symmetric possibly bi-infinite filter (i.e. there is no m such that $\gamma_k = 0$ for all $|k| > m$) with real transfer function

$$\Gamma(\omega) = \sum_{|k|<\infty} \gamma_k \exp(-ik\omega) = \gamma_0 + 2\sum_{k>1} \gamma_k \cos(k\omega)$$

and suppose we want to approximate Y_T given $X_1, ..., X_T$ (real-time estimation problem). Since the bi-infinite estimate Y_T satisfies

$$Y_T = ... + \gamma_T X_0 + \gamma_{T-1} X_1 + ... + \gamma_0 X_T + \gamma_1 X_{T+1} + ...$$

a straightforward device would be to replace unknown random variables on the right-hand side by estimates

$$\hat{Y}_T := ... + \gamma_T \hat{X}_0 + \gamma_{T-1} X_1 + ... + \gamma_0 X_T + \gamma_1 \hat{X}_{T+1} + ... \tag{1}$$

(note that $\gamma_T = \gamma_{-T}$ by symmetry).

In many applications, the filter coefficients decay rapidly so that $|\gamma_{T+i}|$ is 'small', implying that the estimate (1) can be replaced by

$$\hat{Y}_T = \gamma_{T-1} X_1 + ... + \gamma_0 X_T + \gamma_1 \hat{X}_{T+1} + ... \tag{2}$$

Traditional methods such as implemented in X-11-ARIMA (see Ladiray and Quenneville (2001)), X-12-ARIMA (see Findley, Monsell, Bell, Otto, and Chen (1998)), TRAMO/SEATS (see Maravall and Gomez (1994)) or in the structural model-based approach (see Harvey (1989)) solve the above real-time estimation problem - either explicitly or implicitly[5] - by stretching $X_t, t = 1, ..., T$ by back- and forecasts of the time series based on a time series model for the DGP. Assume, for example, that X_t is a random-walk process. Then $\hat{X}_{T+k} = X_T$, $k > 0$, and 2 becomes

$$\begin{aligned} \hat{Y}_T &= \gamma_{T-1} X_1 + ... + \gamma_0 X_T + \gamma_1 \hat{X}_{T+1} + \gamma_2 \hat{X}_{T+2} ... \\ &= \gamma_{T-1} X_1 + ... + \gamma_0 X_T + \gamma_1 X_T + \gamma_2 X_T ... \\ &= \gamma_{T-1} X_1 + ... + \sum_{k\geq 0} \gamma_k X_T \end{aligned}$$

[5] Kalman-filtering and smoothing in the structural model-based approach does not compute forecasts explicitly.

For a white noise process one would obtain

$$\begin{aligned}\hat{Y}_T &= \gamma_{T-1}X_1 + ... + \gamma_0 X_T + \gamma_1 \hat{X}_{T+1} + ...\\ &= \gamma_{T-1}X_1 + ... + \gamma_0 X_T\end{aligned}$$

and for 'white noise plus constant level' one obtains

$$\begin{aligned}\hat{Y}_T &= \gamma_{T-1}X_1 + ... + \gamma_0 X_T + \gamma_1 \hat{X}_{T+1} + ...\\ &= \gamma_{T-1}X_1 + ... + \gamma_0 X_T + \gamma_1 \overline{X} + \gamma_2 \overline{X} + ...\\ &= (\gamma_{T-1} + 1/T\sum_{k=1}^{T}\gamma_k)X_1 + ... + (\gamma_0 + 1/T\sum_{k=1}^{T}\gamma_k)X_T\end{aligned}$$

where $\overline{X}$ is the mean (level estimate) of the series. More generally, forecasts $\hat{X}_{T+k}$ are functions of past and present observations so that 2 is a particular asymmetric filter, see Stier and Wildi (2002) and Wildi (2004) chapter 2 for details.

2.2 Methodological Limitations

Model-based approaches[6] rely on common methodological principles which might lead to problems in the context of real-time signal extraction. We here briefly summarize the main points.

2.2.1 Estimation

Let $f(X_1, ..., X_T|\theta)$ denote the common density function of $X_1, ..., X_T$. This function is determined (by model assumptions) up to a parameter-vector θ which links the model to the data. The maximum likelihood (ML-) estimate $\hat{\theta}$ of θ is

$$\begin{aligned}&\arg\left(\max_\theta f(X_1, ..., X_T|\theta)\right)\\ =\ &\arg\left(\max_\theta\left(\sum_{t=2}^{T}\ln\left(f(X_t|X_{t-1}, ..., X_1, \theta)\right) + \ln(f(X_1|\theta))\right)\right)\end{aligned} \tag{3}$$

Often, in applications, $\ln(f(X_1|\theta))$ is ignored and the resulting $\hat{\theta}$ is called a conditional ML-estimate. The crucial point is that $f(X_t|X_{t-1}, ..., X_1, \theta)$ is the distribution of the so called innovation, that is the amount of new information contained in X_t which cannot be accounted for by the past $X_{t-1}, ..., X_1$. If we neglect $\ln(f(X_1|\theta)$ in 3, then we deduce that *ML-estimates essentially rely on (the distribution of) one-step ahead forecasting errors, irrespective of the purpose of a particular application.*

Users interested in multi-step ahead forecasts or, more generally, in a particular function of one- and multi-step ahead forecasts, such as 2, are assigned to rely on statistics which are not derived explicitly for their particular purpose.

[6]The treatment here is general i.e. it transcends the preceding ARIMA-based approach.

2.2.2 Identification

Process-identification is guided by a priori knowledge (if available) and by statistics which are tailored to that particular purpose. Information criteria are of the general form

$$\text{Goodness of Fit} + \text{Model Complexity} \rightarrow \min$$

The goodness of fit can be measured by the size of the residuals (estimated innovations) and the complexity of the model can be accounted for by the number of estimated parameters. To illustrate this concept, the widely-used AIC-criterion is based on

$$\ln(\hat{\sigma}^2) + 2\frac{\text{Number of Parameters}}{T}$$

where $\hat{\sigma}^2$ is the mean-square one-step ahead forecasting error.

As for the estimation, the identification generally relies on one-step ahead forecasting errors irrespective of the particular application at hand.

2.2.3 Diagnostics

In order to verify the validity of the imposed model assumptions, typically model-residuals are analyzed. Interesting topics include stationarity (constant mean, constant variance) and independence which, in practice, is often boiled-down to zero-autocorrelation. A widely-used instrument in this context is the so-called Ljung-Box statistic

$$Q_m = T(T+2)\sum_{k=1}^{m}\frac{r_k^2}{T-k}$$

where r_k is the sample autocorrelation of the residuals[7].

As for the preceding two model-building phases, diagnostics generally rely on one-step ahead forecasting errors. New developments emphasizing the real-time estimation problem are provided in McElroy and Wildi (2007).

2.3 The One-Step-Ahead Perspective

Potential weaknesses associated with the above model-based approach are well known. So for example in section 3.7 of the X-12-ARIMA reference manual (version 0.2.10) the authors argue "The point forecasts are minimum mean squared error linear predictions of future y_t's based on the present and past y_t's assuming that the true model is used .... These are standard assumptions, though obviously unrealistic in practical applications.

[7] Often, model-based methods also check additional residual properties such as autocorrelations of squared errors for example.

What is more realistically hoped is that the regARIMA model will be a close enough approximation to the true, unknown model for the results to be approximately valid".

In order to exemplify weaknesses of the model-based approach, we here rely on a particular example related to the KOF economic barometer. For the monthly series (business-survey data) in fig.4, solid line, TRAMO identifies the following airline-model

$$(1 - B)(1 - B^{12})X_t = (1 - 0.662B)(1 - 0.824B^{12})\epsilon_t \qquad (4)$$

after adjustments for outliers and calendar effects. As can be seen from typical diag-

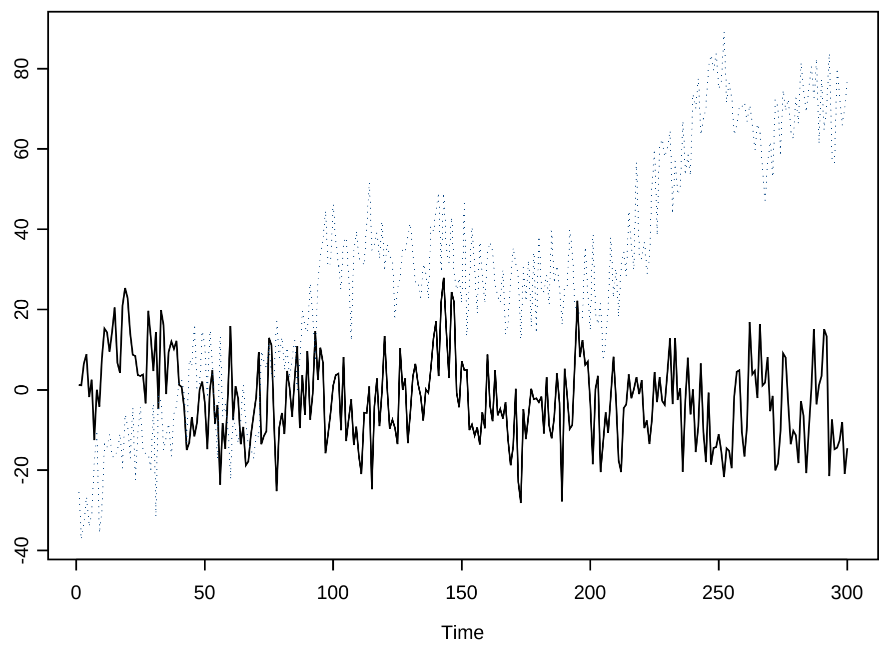

Figure 4: Series 31 (solid) and model simulation (dotted)

nostic plots in fig.5 'standard' model assumptions are met: neither the autocorrelation- nor the partial autocorrelation-function nor the Ljung/Box statistics suggest significant departures from the null-hypothesis[8]. However, a realization - simulation - of the process defined by 4 in fig.4 (dotted line) shows obvious departures from the original path[9]. The artificial series is dominated by a strong trend component which is a 'stylized fact' of I(2)-processes and therefore of the identified airline model. The level of the original time series is much more 'stationary' because the series is *bounded* by construction.

The observed model mis-specifications aren't by far exceptional. In fact, they seem to be the rule, as can be seen in table 2 which compares models identified by three procedures

[8]TRAMO as well as X-12-ARIMA provide additional diagnostic tools such as heteroscedasticity or model stability tests which did not lead to a rejection of the above model neither.

[9]The plotted path represents quite well the main-features of the process-dynamics so that we can safely renounce on simulation experiments.

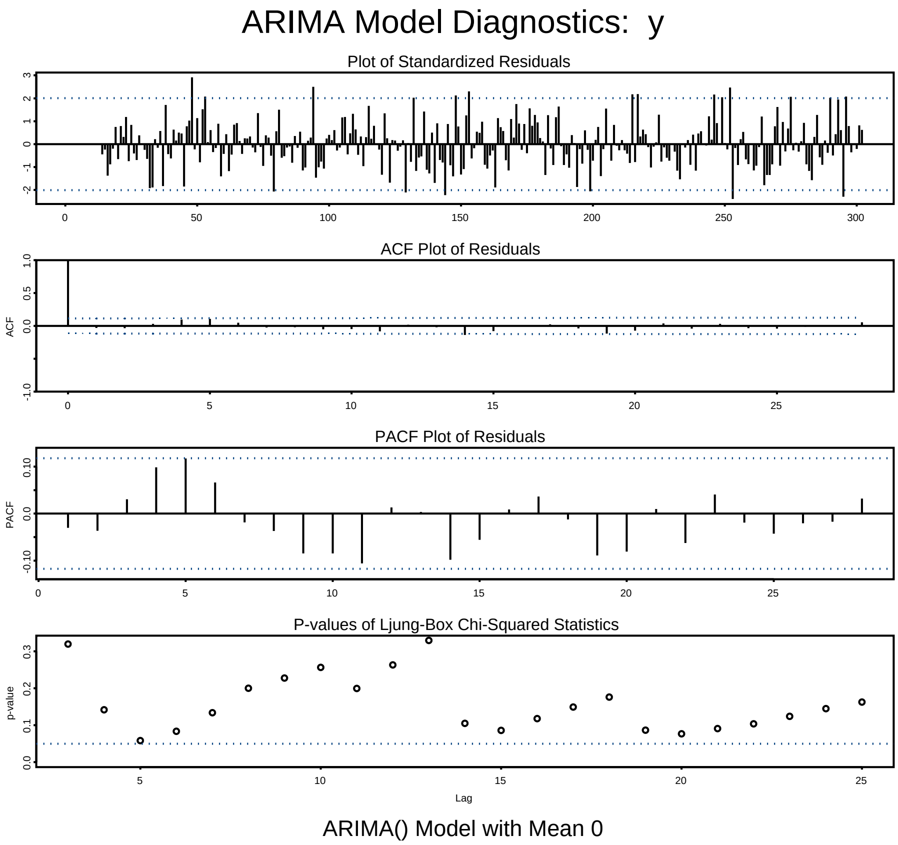

Figure 5: Diagnostics airline model for series 31

for the *bounded* business survey data analyzed in section 3. We gather that approximations of the putative DGP by ARIMA-models can be of rather poor quality in practice because traditional statistics - addressing estimation, identification or diagnostics - are misleadingly conditioned by the rigid short term one-step ahead perspective imposed by maximum-likelihood principles. The main point is that the misspecified model may perform well for the purpose imposed by the optimization criterion, namely one-step ahead forecasting, but for real-time applications, relying also on multi-step ahead forecasts, performances can deteriorate, as shown by our empirical examples below.

Clements and Hendry (2004) show that performances of one- and multi-step ahead forecasts are conflicting in the presence of model mis-specification, which is the rule in practice. *Therefore, we conclude that the traditional model-based approach to real-time signal extraction is subject to a fundamental methodological flaw susceptible to entail efficiency in practice.* The following empirical results suggest that the incurred loss in efficiency can be quite important in applications.

Real-time signal extraction is an important prospective estimation problem whose particular structure is not explicitly addressed by methods relying on maximum likelihood principles. A formal motivation for the latter, namely efficiency, stands with very unrealistic assumptions like 'true' DGP or large (infinite) data sets. Unfortunately, a practically very relevant phenomenon, namely mis-specification, offenses these principles. As stated in Chin, Geweke, and Miller (2000): "*An unwritten rule of forecasting is that accuracy is enhanced by forecasting directly what is of interest*". By contrast, this unwritten rule stands with experience.

3 Customized Criteria: Level Approximation

3.1 Level-Approximation

One-sided real-time filters should not be optimized with respect to potentially misleading one-step ahead forecasting performances. Instead, the mean-square filter error

$$E[(Y_T - \hat{Y}_T)^2] \rightarrow \min$$

should be minimized where Y_T is the signal and $\hat{Y}_T$ is the output of an optimal one-sided filter. The mean-square distance measure emphasizes the level approximation problem. A more general criterion will be used for the identification of turning-points in section 4.

Denote by $\Gamma(\omega)$ and $\hat{\Gamma}(\omega)$ the transfer functions of symmetric and one-sided filters and assume that the time series X_t is stationary. Then Wildi (2004) (chapter 5) and (2007) (chapter 4) shows that parameters of the real-time filter can be optimized by

$$\frac{2\pi}{T} \sum_{k=-T/2}^{T/2} |\Gamma(\omega_k) - \hat{\Gamma}(\omega_k)|^2 I_{TX}(\omega_k) \rightarrow \min \tag{5}$$

where $I_{TX}(\omega_k)$ is the periodogram of X_t and $\omega_k = k2\pi/T$, $k = 0, ..., T/2$. Briefly, the expression in 5 is a finite sample approximation of the following integral

$$\int_{-\pi}^{\pi} |\Gamma(\omega) - \hat{\Gamma}(\omega)|^2 dH(\omega) = E[(Y_t - \hat{Y}_t)^2] \tag{6}$$

where $H(\omega)$ is the spectral distribution of X_t. Generalizations to non-stationary integrated processes are straightforward, by replacing the periodogram in 5 by a so-called pseudo-periodogram, see Wildi (2007), chapter 6[10]. However, these generalizations are often unnecessary in the context of leading indicators because time series are bounded. The resulting real-time estimation method introduced in Wildi (2004) is called *direct filter approach (DFA)*.

One can show that the error linking the finite sample estimate 5 to the mean-square filter

[10] Additionally, suitable filter constraints must be imposed.

error 6 is ‘small’. More precisely, consider the following approximations:

$$\frac{2\pi}{T}\sum_{k=-T/2}^{T/2}|\Gamma(\omega_k)-\hat{\Gamma}(\omega_k)|^2 I_{TX}(\omega_k) \approx \frac{2\pi}{T}\sum_{k=-T/2}^{T/2} I_{T\Delta Y}(\omega_k) \tag{7}$$

$$= \frac{1}{T}\sum_{t=1}^{T}(Y_t-\hat{Y}_t)^2 \tag{8}$$

$$\approx E[(Y_t-\hat{Y}_t)^2] \tag{9}$$

where $I_{T\Delta Y}(\omega_k)$ is the periodogram of the filter error $Y_t - \hat{Y}_t$. In practice, the left-hand side of 7 is observed only. Proposition 10.16 in Wildi (2007) proves that the sample mean-square error 8 is an asymptotically efficient estimate of the mean-square error. Proposition 10.17 implies that the approximation error in 7 - the convolution error - is of smaller magnitude than that in 9 (superconsistency). Therefore, efficiency of the sample mean-square estimate 8 is inherited by the left-hand side of 7. Moreover, uniformity of these results is derived under general regularity assumptions which can be easily implemented in practice, see section 3.5.4 in the cited literature. As a result, *criterion 5 can be interpreted as the minimization of an efficient estimate of the unknown mean-square filter error.*

3.2 Empirical Results

3.2.1 KOF Economic Barometer

Users of the KOF economic barometer are interested in appraising business-cycles. Accordingly, the signal or, more precisely, the trend was defined by

$$\tilde{\Gamma}(\omega) := \begin{cases} 1 & 0 \le |\omega| \le \pi/14 \\ \dfrac{\pi/7-|\omega|}{\pi/7-\pi/14} & \pi/14 \le |\omega| \le \pi/7 \\ 0 & \pi/7 \le |\omega| \le \pi \end{cases} \tag{10}$$

Components with duration longer than 28 months (frequency $\pi/14$) pass the filter unaltered whereas those shorter than 14 months (frequency $\pi/7$) are eliminated[11]. This paper does not discuss the ‘correctness’ of the signal definition which is to be considered as an exogeneous target imposed by users of the above instrument. Basically, any symmetric transfer function can be approximated in real-time by the proposed approach.

The MA-coefficients of the filter can be obtained by inverse Fourier transform:

$$\tilde{\gamma}_k = \begin{cases} -\dfrac{1}{\pi(\pi/7-\pi/14)}\left[\dfrac{\cos(k\pi/7)-\cos(k\pi/14)}{k^2}\right] & k \neq 0 \\ \dfrac{1}{2}\left(\dfrac{1}{7}+\dfrac{1}{14}\right) & \text{k=0} \end{cases} \tag{11}$$

[11]The ‘traditional’ band-pass filter is unnecessary here because the data is informative about growth-rates, not levels. Note that all time series involved are bounded per construction.

It is readily verified that $\tilde{\gamma}_k = \tilde{\gamma}_{-k}$. In our experiments, the signal is generated by a truncated MA(101)-filter $\Gamma(\cdot)$ defined by

$$\gamma_k = \begin{cases} C\tilde{\gamma}_k, & |k| \leq 50 \\ 0, & \text{otherwise} \end{cases} \tag{12}$$

where $C := 1/\sum_{|k|\leq 50} \tilde{\gamma}_k$ ensures that $\Gamma(0) = 1$.

In our comparisons, the above trend definition was fixed for all competing approaches: real-time estimates of ARIMA-based methods are obtained by forecasts which are put into 2, using symmetric filter coefficients defined in 12, and the DFA is based on 5[12]. The following procedures were analyzed: TRAMO as implemented in the latest TSW-package, March 2006 (denoted by TRAMO-TSW in our tables) which can be downloaded from the Bank of Spain (http://www.bde.es/servicio/software/econome.htm) and DEMETRA, version 2.04, the latest release of the EUROSTAT seasonal adjustment software package[13].
Thirty-six time series belong to the analyzed data set, spanning a time period from June 1979 to March 2006[14]. All series were 'linearized' (calendar-effects and/or log-transform and/or calendar effects) prior to processing so that model-assumptions should be satisfied. Automatic model-identification settings were chosen for all model-based approaches (RSA=3 for TRAMO-TSW). Table 1 summarizes mean and median relative performances of the model-based approaches relative to the DFA:

$$\frac{\sum_{t=k_Y+1}^{T-k_Y}(\hat{Y}_t^{MBA} - Y_t)^2}{\sum_{t=k_Y+1}^{T-k_Y}(\hat{Y}_t^{DFA} - Y_t)^2} \tag{13}$$

where $k_Y = 50$ is such that Y_t can be computed in $t = k_Y, ..., T - k_Y$. Out-of-sample periods last two years during which models and filters are kept fixed (no revisions). Re-

	TRAMO-TSW		TRAMO-D		X-12-T		X-12-A	
	in-	out-	in-	out-	in-	out-	in-	out-
mean	136%	146%	134%	143%	134%	144%	124%	131%
median	139%	145%	131%	136%	127%	135%	120%	128%

Table 1: Efficiency loss incurred by MBA in- and out-of-sample'

sults are obtained for samples of length 120[15] (ten years of data). *Numbers larger than*

[12] $\hat{\Gamma}(\omega)$ in 5 is based on a new ARMA-ZPC-filter design, see Wildi (2004), chapter 3, and (2007), chapter 3, for details.

[13] The following versions are used in DEMETRA: TRAMO Nov-99, SEATS Sep-98 (the former is denoted by TRAMO-D in our tables), X-12-ARIMA release 0.2.8. For X-12-ARIMA, we supplemented stationary model alternatives which where unavailable in the original model-file x12a.mdl, namely (2 0 2)(1 0 1),(1 0 1)(1 0 1),(2 0 0)(0 0 0) and (2 0 0)(0 0 1).

[14] The data as well as a detailed description of the survey can be provided by the KOF. Requests are to be addressed to the author at wia@zhwin.ch.

[15] The last four observations of the sample were dropped because of potential (small) data revisions so that the sample length amounts to $T = 318$. In order to compute the symmetric filter, 50 values are lost at the current boundary. Two additional years of data are retained for out-of-sample results so that estimation and identification is restricted to $t = 125, ..., 244 = 318 - 50 - 24$ (time span 01.08.1989-01.07.1999).

one correspond to a gain obtained by the DFA. X-12-T is a hybrid procedure based on identification by TRAMO-D and estimation by X-12-ARIMA. Model-orders identified by

Series	TRAMO-TSW	TRAMO-D	X-12-A
1	(012)(011)	(210)(011)	(210)(011)
2	(210)(001)	(011)(011)	(201)(000)
3	(110)(000)	(011)(011)	(101)(101)
4	(011)(000)	(011)(000)	(100)(000)
5	(211)(000)	(210)(100)	(011)(011)
6	(011)(100)	(110)(000)	(202)(101)
7	(001)(000)	(100)(011)	(212)(011)
8	(011)(000)	(011)(011)	(212)(011)
9	(311)(011)	(011)(011)	(012)(011)
10	(311)(011)	(121)(011)	(212)(011)
11	(311)(011)	(110)(011)	(100)(100)
12	(310)(001)	(011)(000)	(202)(101)
13	(301)(101)	(300)(011)	(210)(011)
14	(011)(011)	(011)(001)	(202)(101)
15	(011)(101)	(011)(011)	(210)(011)
16	(110)(000)	(110)(000)	(202)(101)
17	(311)(000)	(011)(011)	(101)(101)
18	(011)(000)	(011)(011)	(011)(011)
19	(100)(000)	(100)(011)	(101)(101)
20	(011)(000)	(011)(011)	(100)(000)
21	(011)(011)	(011)(011)	(100)(100)
22	(100)(000)	(011)(000)	(100)(000)
23	(100)(100)	(100)(000)	(210)(011)
24	(110)(000)	(010)(001)	(202)(101)
25	(310)(001)	(011)(011)	(011)(011)
26	(013)(000)	(013)(001)	(202)(101)
27	(110)(100)	(011)(011)	(101)(101)
28	(110)(000)	(110)(000)	(201)(000)
29	(311)(011)	(011)(011)	(011)(011)
30	(012)(000)	(012)(000)	(202)(101)
31	(011)(100)	(210)(011)	(210)(011)
32	(011)(000)	(011)(000)	(212)(011)
33	(311)(011)	(010)(011)	(212)(011)
34	(311)(011)	(010)(011)	(011)(011)
35	(311)(011)	(011)(011)	(101)(101)
36	(112)(000)	(112)(000)	(202)(101)

Table 2: Model orders: TRAMO-TSW, TRAMO-D(EMETRA) and X-12-ARIMA, 120 observations (time span 01.08.1989-01.07.1999), Business Survey Data

the above ARIMA-approaches are summarized in table 2. Note that TRAMO-TSW and TRAMO-D disagree in 16 cases with respect to integration orders for example. Differences with respect to X-12-ARIMA are even more striking. The observed model discrepancies illustrate that traditional identification procedures are sensitive to the way they are implemented in practice[16]. This fact may be invoked to justify efforts in merging procedures, see Monsell, Aston, and Koopman (2003).

3.2.2 FED Diffusion Indices

The same experimental design (sample length 120) is applied to a data set comprising 63 time series provided by the federal reserve bank, Philadelphia. Results in- and out-of-sample are summarized in table 3, see Wildi (2007) for details. We infer from these results

	TRAMO-TSW		TRAMO-D		X-12-T		X-12-A	
	in-	out-	in-	out-	in-	out-	in-	out-
mean	135%	161%	130%	157%	137%	177%	132%	141%
median	133%	141%	123%	128%	128%	142%	124%	128%

Table 3: Efficiency loss incurred by MBA in- and out-of-sample

that traditional model-based approaches might be subject to substantial efficiency-losses in real-time applications if the current level of the signal imports.

We are aware that the analyzed series are 'particular' in the sense that they are bounded. However, many important economic time series are bounded too by construction or/and by elementary physical laws or/and by regulating mechanisms. Typical examples are 'growth-rates', which are often more interesting than absolute levels. The proposed DFA reveals weaknesses of ARIMA-based methods in this particular practically relevant framework, characterized by 'small' signal-to-noise ratios.

4 Customized Criteria: Turning-Points

We here assume that turning-points are defined by local extrema of a trend signal Y_t which is defined by the user[17]. In comparison to the preceding level-framework, a customized criterion designed for the detection of turning-points should emphasize filter performances specifically in the vicinity of the local extrema of Y_t. Moreover, in many important real-time applications, anticipations are less harmful than delays. Therefore, a suitable criterion should reflect this error-asymmetry also. We here propose to generalize the optimization criterion 5 in order to account for both problem-specific requirements. To our knowledge, neither is accounted for explicitly by traditional model-based approaches.

[16] Whereas the older TRAMO-D relies more heavily on airline-models, the identification procedure of the newer TSW-package seems to be less biased towards that particular forecasting model.

[17] For the KOF economic barometer, the trend defined in 10 is of interest. However, any signal reflecting the particular purpose of an application and/or the particular intension of a user may be used.

Assume that Y_t is the output of a filter with transfer function $\Gamma(\cdot)$ and that the local extrema of Y_t import. Wildi (2004) proposes a decomposition of the squared filter approximation error into amplitude and phase errors:

$$\begin{aligned}|\Gamma(\omega)-\hat{\Gamma}(\omega)|^2 &= A(\omega)^2+\hat{A}(\omega)^2-2A(\omega)\hat{A}(\omega)\cos\left(\hat{\Phi}(\omega)-\Phi(\omega)\right)\\ &= (A(\omega)-\hat{A}(\omega))^2\\ &\quad +2A(\omega)\hat{A}(\omega)\left[1-\cos\left(\hat{\Phi}(\omega)-\Phi(\omega)\right)\right] \qquad (14)\end{aligned}$$

where $A(\omega)$, $\hat{A}(\omega)$ are the amplitude functions and $\Phi(\omega)$, $\hat{\Phi}(\omega)$ are the phase functions of $\Gamma(\cdot)$ and $\hat{\Gamma}(\cdot)$. If we assume that Γ is symmetric then $\Phi(\omega)\equiv 0$. Inserting 14 into 5 leads to

$$\left(\frac{2\pi}{T}\sum_{k=-T/2}^{T/2}(A(\omega_k)-\hat{A}(\omega_k))^2 I_{TX}(\omega_k)\right.$$
$$\left.+\frac{2\pi}{T}\sum_{k=-T/2}^{T/2}2A(\omega_k)\hat{A}(\omega_k)\left[1-\cos\left(\hat{\Phi}(\omega_k)\right)\right]I_{TX}(\omega_k)\right)\to\min$$

The first summand is that part of the (estimated) mean-square filter error which is attributable to the amplitude function of the real-time filter $\hat{\Gamma}$ and the second summand corresponds to the error attributable to the time shift (phase), see Wildi (2007), chapter 3. Consider a slightly more general criterion

$$\begin{aligned}&\frac{2\pi}{T}\sum_{k=-T/2}^{T/2}(A(\omega_k)-\hat{A}(\omega_k))^2 W(\omega_k)^2 I_{TX}(\omega_k)\\ &+\lambda\frac{2\pi}{T}\sum_{k=-T/2}^{T/2}2A(\omega_k)\hat{A}(\omega_k)\left[1-\cos\left(\hat{\Phi}(\omega_k)-\Phi(\omega_k)\right)\right]\\ &\cdot W(\omega_k)^2 I_{TX}(\omega_k)\to\min \qquad (15)\end{aligned}$$

For $W(\cdot)=1$ and $\lambda=1$, the level-criterion 5 results. If $\lambda>1$, then the time shift in the pass-band[18] is emphasized. The corresponding real-time filter would have smaller time delays i.e. *it would be faster*. If $W(\omega)$ is monotonically increasing in ω, then the amplitude function in the stop-band is emphasized. Stronger damping of undesirable high-frequency components is related to the *reliability* of the real-time filter since 'noise' is damped more heavily.

The key idea behind criterion 15 is that one can account for speed - through λ - and reliability - through $W(\omega)$ - separately. Optimal choices of λ and $W(\omega)$ for the detection of turning-points are proposed in the following section. As we shall see, these clearly differ from the best mean-square level settings ($W(\cdot)=1$ and $\lambda=1$) which confirms that both issues, level approximation and detection of turning-points, are incongruent requirements.

[18]The pass-band is addressed only because the phase error is weighted by $A(\omega)\hat{A}(\omega)$.

4.1 Empirical Results

4.1.1 KOF Economic Barometer

The outcomes of the following three different filters are analyzed in figs.6 and 7: the L-L-filter is based on $\lambda = 1, W(\cdot = 1)$, the L-D-filters are based on $W(\omega) = |1 - \exp(-i\omega)|$ in 15 and either $\lambda = 1$ (no emphasize of time delay) or $\lambda = 6$ (time delay is emphasized). The L-L-filter is the best (mean-square) level approximation of the trend for the *original* data whereas the L-D-filter based on $\lambda = 1$ is the best level approximation for the *differenced* data. The latter is of interest because turning-points are characterized by sign-flips of $Y_t - Y_{t-1}$. In order to compare the various designs, results obtained in first differences are transformed (integrated) to original levels in the figures.

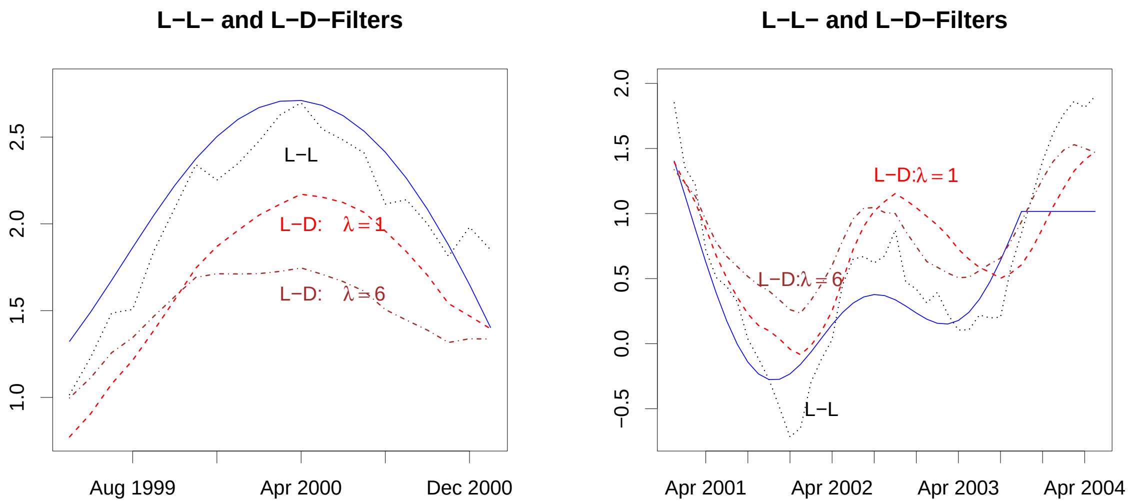

Figure 6: TP-identification: comparison of three different optimization criteria

Fig.6 illustrates that the optimal level filter (L-L-design) is not well suited for the detection of turning-points because it is too noisy, recall fig.3 also. The main problem of this particular design is the insufficient damping in the stop-band as can be seen in fig.7: the amplitude function is larger than 0.5 in $[\pi/2, \pi]$. L-D-filters damp undesirable noise much more strongly, because $W(\cdot)$ emphasizes stop-band issues. Therefore, real-time signals are smoother which is associated to improved reliability or, equivalently, to less false alarms. Unfortunately, the best level approximation in differences ($\lambda = 1$) is subject to substantial delays. However, the filter based on $\lambda = 6$ is faster, as expected. Its speed is comparable to the L-L-filter (similar time shift functions in the pass-band) and damping properties in the stop-band are comparable to the L-D-filter with $\lambda = 1$. These advantageous features are obtained at costs of the amplitude function in the pass-band: the observed negative bias implies worse level performances. *By increasing λ we have traded level-performance against smaller time delays.* Overall, it seems that the fast L-D-filter based on $\lambda = 6$ has inherited advantageous features of both other filter designs. A com-

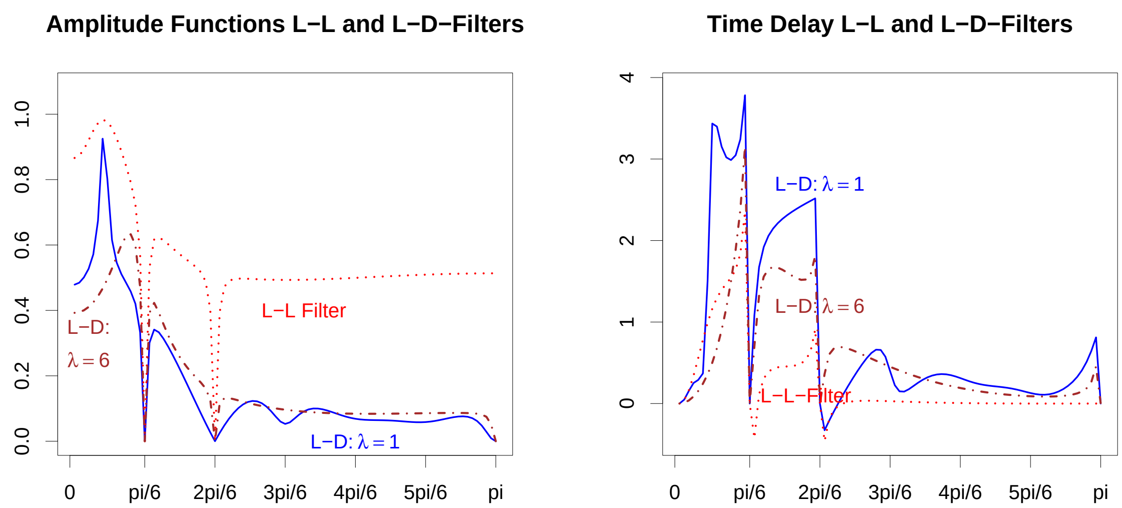

Figure 7: Amplitude and time shift functions of the resulting real-time filters

parison of the rate of false alarms[19] confirms this feeling since we register 'only' 9.1% for the fast L-D-filter, which is substantially smaller than 12.4% (original L-D-) and 14.9% (original L-L-filter) on a time span ranging from January 1984 to March 2006.

At this stage, it is important to analyze the effect of the smaller time delay on filter performances more comprehensively. For that purpose, consider the data and the ideal trend *in first differences* in fig.8[20]. In this representation, the interesting TP's are characterized by sign-flips of the level signal: an ideal real-time filter should perfectly replicate these flips. Therefore, level performances of the filter should be particularly good in the vicinity of TP's.

Two properties characterize TP's in our representation (first differences): the signal crosses the zero-line and its slope is extreme[21]. These characteristics can be accounted for by specialized real-time turning-point filters i.e. by customized optimization criteria. Indeed, the effect of the time delay on the filter error is proportional to the slope of the signal (in the mean): a delay of one time unit induces an error of magnitude $\approx \delta$ (in the mean) if the slope of the signal is δ accordingly. Emphasizing the time delay, as we do by increasing λ, therefore corresponds to a *selective level estimation*. Better performances are obtained specifically in time periods where the slope of the signal is large (in absolute value) which correspond to TP's. *Larger λ implicitly highlight TP's at the cost of the remaining time points.*

In order to verify the above selectivity-effect on the time-axis we plot filter errors $Y_t - \hat{Y}_t$ for both L-D-filters ($\lambda = 1$ and $\lambda = 6$) in fig.9. Vertical lines correspond to TP's (extrema of Y_t). As can be seen, the faster L-D-filter performs better (in the mean) in the

[19] A false alarm is registered in t if the sign of $\hat{Y}_t - \hat{Y}_{t-1}$ does not coincide with that of $Y_t - Y_{t-1}$.

[20] Recall that both L-D-filters are optimized for performances in first differences

[21] The latter property is a 'salient feature' of many economic time series. In our opinion, large absolute second 'derivatives' (second order differences) of the signal in TP's are to some extent the expression of market frictions.

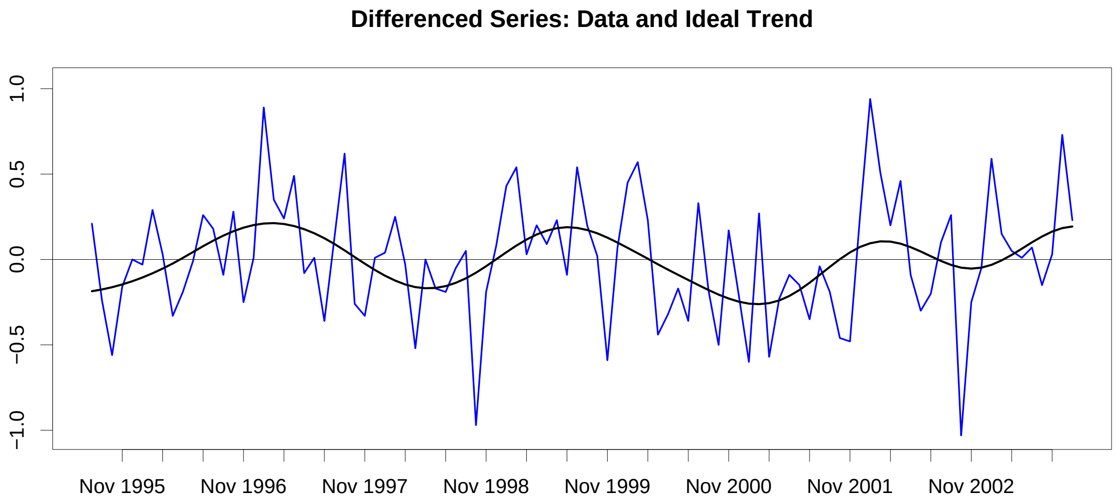

Figure 8: First differences: signal and data

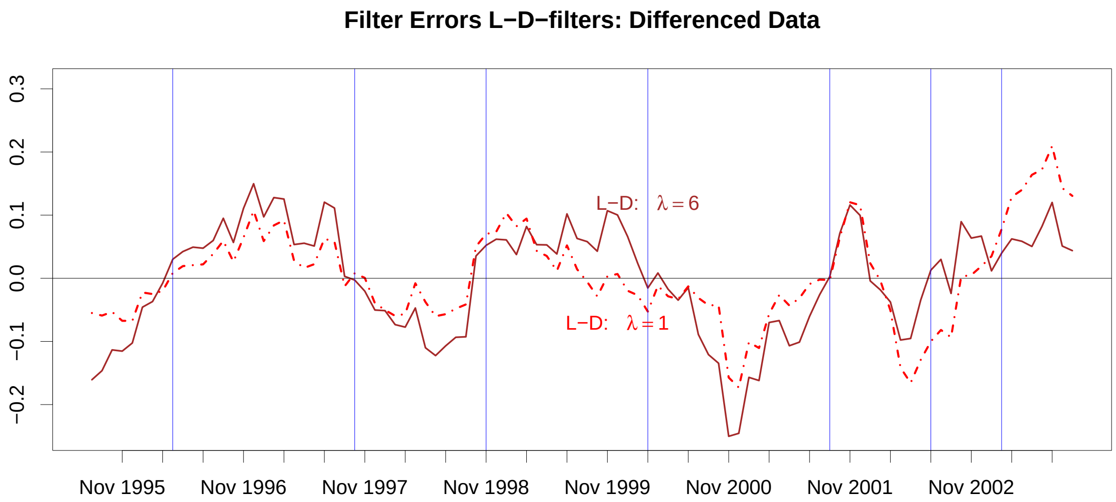

Figure 9: First differences: filter errors L-D-filters $\lambda = 6$ (solid) vs. $\lambda = 1$

TP's but worse 'outside'. The tangible effects are reported in table 4, where mean-square filter errors are computed for both filters on the whole sample and specifically in TP's. As expected, the original L-D-filter ($\lambda = 1$) performs better in the mean on the whole sample because it has been optimized accordingly. However, the faster filter performs better in the TP's[22].

By relying on the decomposition 14, the estimated mean-square filter error can be decomposed into amplitude and time shift error contributions. Both error components are

[22]The difference in performance in the TP's is significant: the t-test of the pairwise comparison is -2.29 (independence of filter errors can be assumed in the test because time intervals between TP's are sufficiently large).

	$\lambda = 6$	$\lambda = 1$
TP's only	0.0008	0.0034
All time points	0.0074	0.0055

Table 4: Filter errors in first differences: whole sample vs. TP's only

reported in table 5. The 'Total'-row is an estimate of the mean-square error computed in

	$\lambda = 6$	$\lambda = 1$
Amplitude	0.0062	0.0023
Time shift	0.00017	0.0019
Total	0.0064	0.0042

Table 5: Filter errors in first differences: amplitude and time shift components

the frequency domain: it corresponds to the time-domain estimates in the second row of table 4. As expected, the amplitude error component has substantially augmented by increasing λ: the resulting disequilibrium between amplitude- and time-shift (mean-square) error components reflects our disposition to trade *mean*-level against *selective*-level or, equivalently, against time delay performances. As a result, *the generalized optimization criterion 15 assigns more weight to delays than to anticipations.* Neither the latter error-asymmetry nor the selective regard on TP's is accounted for explicitly by traditional approaches.

Logit-models are often claimed to mate the real-time turning-point problem better than traditional ARIMA-based approaches. We refer to Birchenhall, Jessen, Osborn, and Simpson (1999), Chin, Geweke, and Miller (2000), Sensier, Artis, Osborn, and Birchenhall (2005), Estrella and Mishkin (1997) and (1998), Lamy (1997) and Filardo (2004) for corresponding applications. The following logit-model has been identified on the basis of a priori knowledge[23] and information criteria

$$P(I_t = 1|X_t, X_{t-1}, ...) = \frac{1}{1 + \exp\left(-c - aX_t - \sum_{k=0}^{4} b_k \tilde{X}_{t-k}\right)} \tag{16}$$

where X_t is the raw (unfiltered) leading indicator, $\tilde{X}_t$ are its first differences and I_t is a dichotomic variable which is zero or one depending on the sign of $Y_t - Y_{t-1}$. The parameter estimates are $c = 2.11$, $a = -1.36$, $b_0 = 5.46$, $b_1 = 4.40$, $b_2 = 3.48$, $b_3 = 4.15$ and $b_4 = 3.90$. The negative sign of a implies, for example, that the probability of an economic expansion (upward TP) tends to increase for smaller values of X_t. The strength of the above logit model is that inferences about TP's explicitly rely on information in levels as well as in first differences which is in accordance with our knowledge about the series dynamics. Clearly, models either in levels only or in differences only perform worse.

[23] The boundedness of the time series implies that the probability of a TP tends to increase with the (absolute) level. Therefore, level information should be important as confirmed by corresponding criteria.

Although the L-D-filter with $W(\omega) = |1 - \exp(-i\omega)|$ and $\lambda = 6$ performs quite well, substantial improvements are still possible by strengthening the weight attributed to undesirable high-frequency components. In figure 10, real-time performances of a turning-point filter (TP-filter) based on $W(\omega) = |\omega|^3$ and $\lambda = 6$ is compared to the above logit-model[24]. In order to accommodate for easy visual inspection, the output of the latter is transformed according to

$$2\hat{I}_t := \begin{cases} 2 & , \quad \hat{P}(I_t = 1|X_t, X_{t-1}, ...) > 0.5 \\ 0 & , \quad \text{otherwise} \end{cases}$$

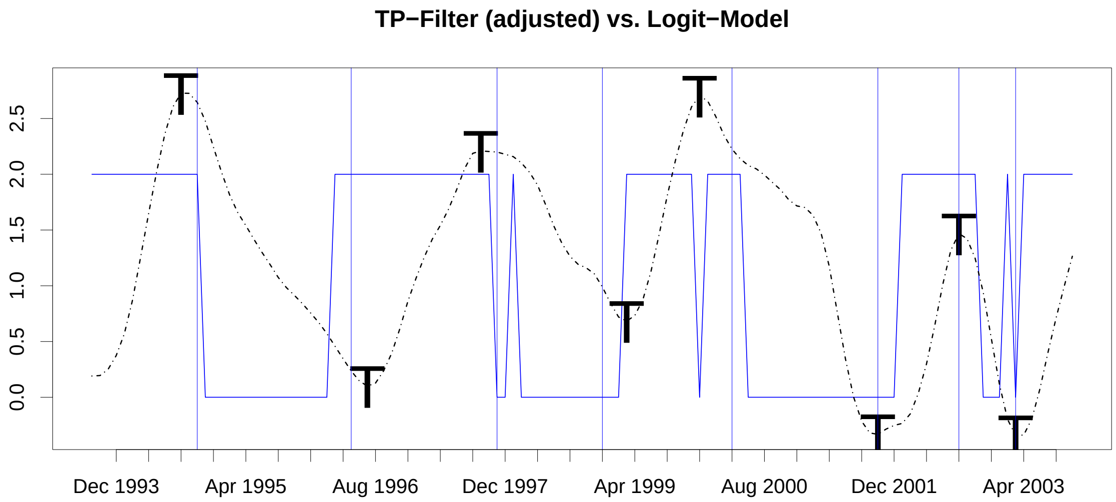

Figure 10: Logit-Model vs. TP-filter adjusted

Sign-flips of $\hat{Y}_t - \hat{Y}_{t-1}$ are marked by a capital 'T'. Except for the particular TP in August 1996, the TP-filter is at least as fast or faster than the logit-model and it is completely free of noisy random alarms. The latter contaminate the logit output after December 1997, before August 2000 and before April 2003. The relatively large lead of the logit-model in August 1996 is due to low time series values which bias the conditional probability of an upswing towards one (because the sign of a is negative in 16). However, the bias induced by extreme observations X_t in 16 is not unproblematic since the additional 'pressure' exerted by X_t can lead to false (random) alarms as observed in fig.10.

Table 6 summarizes the performances of fast L-D- and TP-filters as well as of the logit model. For the concurrent DFA-filters, false alarms are registered if signs of the first increments $\hat{Y}_t - \hat{Y}_{t-1}$ do not coincide with the trend-growth of the signal[25]. The TP-filter anticipates TP's (of the symmetric trend) in April 1995 (by two months), in December

[24] Details related to the choice of W and λ can be found in Wildi (2007).

[25] Such a simple rule is possible because the filter outputs are very smooth. L-L-filters or traditional model-based real-time filters would require more sophisticated rules in order to account for the noisy signal estimates.

	Logit	TP	L-D
Total rate of false signals	11.6%	10.6%	9.1%
Delays	6.6%	4.1%	6.6%
Anticipations	2.5%	6.5%	0.0%
Random alarms	2.5%	0.0%	2.5%

Table 6: Rate of false signals: logit-model vs. TP- and fast L-D-filters

1997 (by two months) and August 2000 (by four months). Therefore, the rate of anticipations is $(2+2+4)/123 = 0.065$[26] as reported in the above table. Rows two and three account for time shift issues and the last row measures the reliability (noise elimination) of each method.

If anticipations are not considered as 'errors', which is the case in our particular leading indicator environment, then the TP-filter performs best with a rate of false alarms (delays exclusively) of 4.1% in *real-time* which is to be compared with 9.1% (random alarms and delays) for both the logit-model and the fast L-D-filter. Random alarms (last row in the above table) of the other two competitors are problematic because they imply that TP's must be confirmed which, in practice, amounts to additional time delays. Finally, note that the apparent equilibrium of delays and anticipations of the TP-filter is a consequence of the vanishing time delay in the whole pass-band of the filter, as can be seen from fig.11. Disequilibria to the disadvantage of anticipations are indicative for delays in the case of

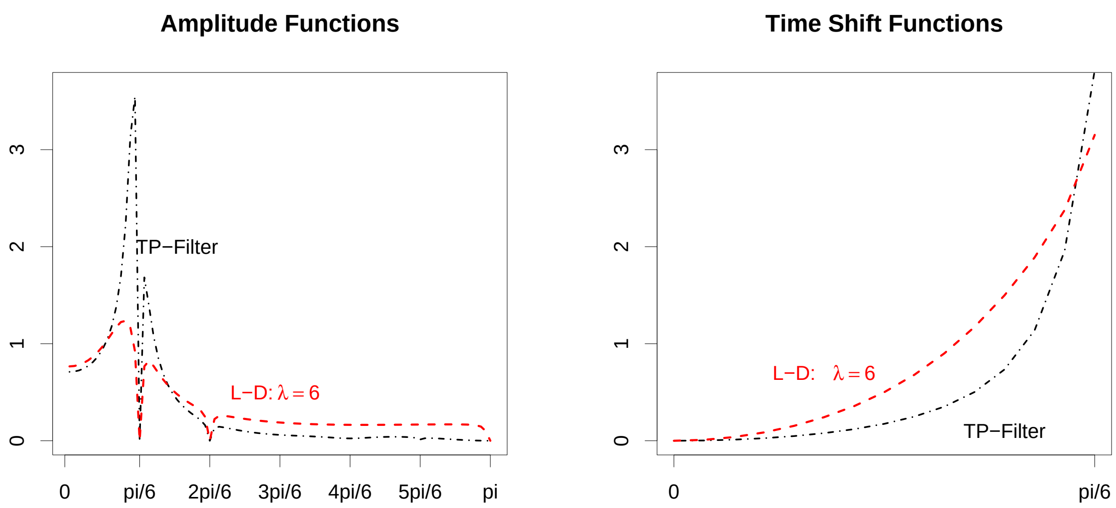

Figure 11: Amplitude and time shift functions: fast L-D vs. TP filter (both adjusted)

the other two approaches.

[26]The symmetric filter can be computed on sample of length 123.

The effect of the strong weights $W(\omega) = |\omega|^3$ and $\lambda = 6$ must have incidences on the un-weighted amplitude function in the pass-band. A distortion of the latter can be seen in fig.11 which emphasizes that (mean-square) level-issues are deliberately sacrificed for better TP-performances. It is worth to emphasize that the smoothing effect of the TP-filter is very pronounced although its time delay vanishes in the pass-band. Note also that various 'risk-aversions' of users can be accounted for explicitly by controlling speed and reliability of the filter through the weights λ and $W(\cdot)$. Neither of these advantageous features is explicitly accounted for by model-based approaches.

The new KOF economic barometer, launched in April 2006, relies on a real-time filter based on the generalized 'customized' criterion 15. The risk-profile has been accounted for by optimal weights λ and $W(\omega)$.

4.2 European Economic Sentiment Indicator (ESI)

The European Economic Sentiment Indicator ('in EU') is a monthly time series issued by the statistical office of the Slovak Republic[27]. The original aggregate is processed by a seasonal adjustment method called DAINTIES, see Franses, Paap, and Fok (2005). A particularity of the latter procedure is that the filter is strictly one-sided too. The last 120 observations of the published (processed) time series are plotted in fig. 12. We were not allowed to plot the raw series (before adjustment). Nevertheless, its periodogram can be seen in the right panel of fig.12.

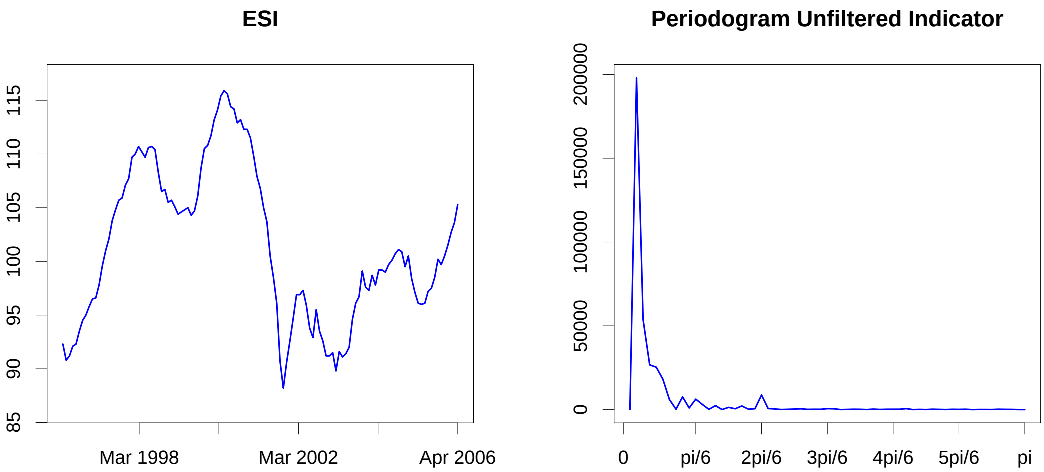

Figure 12: European economic sentiment indicator and periodogram of original time series (the latter is unpublished)

The output of a TP-filter based on $\lambda = 12$ and $W(\omega) = |\omega|^2$ is compared with the pub-

[27]See http://www.statistics.sk/webdata/english/konja/ies_ang.htm

lished indicator in fig. 13. Vertical lines correspond to TP's of the symmetric trend signal 10 (alternative trend definitions could be analyzed of course). As can be seen, the mean

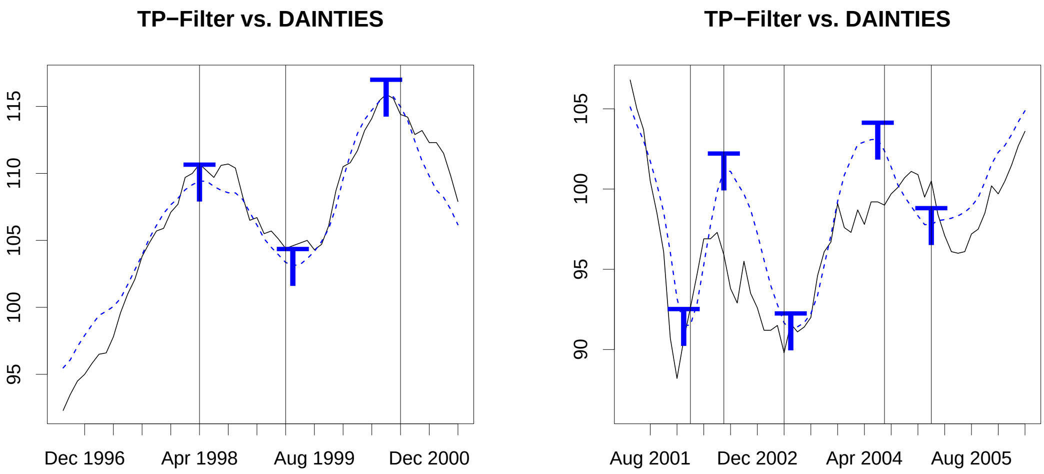

Figure 13: ESI: TP-filter (dashed) vs. DAINTIES

delay of the real-time filter is zero as confirmed by the time shift in the pass-band which is plotted in fig. 14. The TP-filter leads extrema of the published indicator in April 1998 (lead by 4 months), April 1999 (lead by 4 months), June 2004 (lead by 5 months) and February 2005 (lead by 4 months). Lags of one month, with respect to extrema of the published series, occur during the short cycle between February 2002 and April 2003, specifically in February 2002, July 2002 and April 2003. It is worth to emphasize that the apparent lag is falsely suggested by the noisy extrema of the published series in these time points: in fact TP's of the symmetric trend indicate that the TP-filter leads in February 2002, is synchronous in July 2002 and lags in April 2003 by one month, see fig. 13. This example demonstrates that extrema of a noisy signal should not be used as a reference for the identification of TP's. Note also that the DFA-filter does not generate false alarms between consecutive TP's because damping in the stop-band is strong, see fig. 14. In contrast, the published indicator is noisier which implies that one has to wait some time in order to assert the occurrence of a TP with a certain amount of certainty.

Given the advantages of the new filter design it is interesting to compare our analysis with the appraisal by analysts of the EU: "The approach of the European Commission to adjusting the business and consumer surveys (BCS) has always been seasonal adjustment, not trend/cycle extraction, i.e. smoothing. Since the irregular component carries information on respondents' perception of economically relevant special events such as strikes, elections or strong exchange rate or commodity price movements, we believe that retaining the component is important in interpreting the data. We use this approach in a consistent manner for all approx. 50.000 series covered by the comprehensive EU BCS Programme, including sectoral and global composite indicators. To my knowledge, the s.a. approach is also the dominant, if not exclusive, approach used by the national survey institutes, which are conducting the surveys on behalf of the Commission, for their own,

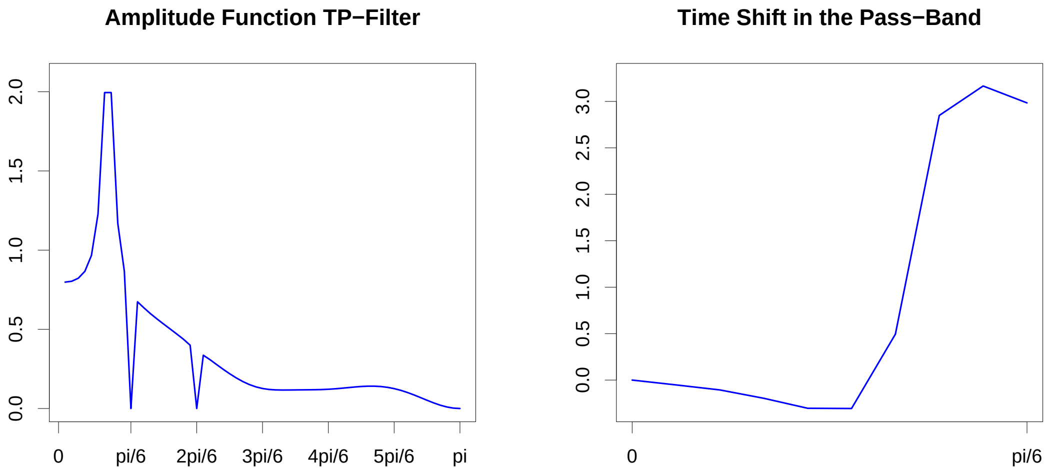

Figure 14: Amplitude und time shift (the latter in the pass-band only)

national indicator releases (some do not adjust the series at all). Of course, the real-time trend/cycle extraction approach of your DFA filter is very relevant in its own right for the identification of the cycle and its turning points. Therefore, your work is very interesting and stimulating. While the calculation and publication of such smooth indicators representing the trend/cycle component is not among our priorities for the nearer future, we would still be interested in learning more about the possible advantageous or complementary features of your DFA approach with respect to our traditional approach of processing the BCS data."

We understand that many experienced users prefer original and/or seasonally adjusted data for their own analysis. However, we feel that there is a growing market for fast and reliable real-time trend estimates of important time series, especially for practitioners interested in business-cycles. Our experience suggests that the new economic barometer of the KOF is appealing to a large audience because its smooth and fast real-time dynamics are less confusing. In this respect the new filter design would be a useful alternative complementing traditional 's.a.'-indicators in the EU.

5 Conclusions

The more or less impressive gains obtained by specialized filter-designs - optimized either for level or for turning-point applications - illustrate the importance of customized optimization criteria that mate the practically relevant estimation problem[28]. In applications with a strong prospective content, many people do not realize or conceive that one-step

[28]Level-filters are compared with respect to mean-square performances in section 3 and TP-filters are compared with respect to speed and reliability of TP-alarms in section 4 which is consistent with the underlying optimization criteria.

ahead forecasting is not directly related to their particular needs. It is true that the holly-grail of statistical principles does not prevent against these odds. We hope that the 'spirit' in this article contributes to question the usefulness of rigid theoretical principles that often do not directly address the practically relevant problem. Our experience so far strongly suggests that 'customized' criteria in conjunction with domain-knowledge can alleviate mis-specification and over-fitting in practice.

Currently, the approach is extended in three directions, namely customized criteria for financial trading[29], multivariate filtering and non-linear filtering. Each extension is specific for a particular application field in which research is conducted in collaboration with corresponding experts so that domain-specific knowledge can be fully integrated. Some people may complain that this interaction can harm to generality. We believe that it benefits to efficiency.

References

Birchenhall, C., H. Jessen, D. Osborn, and P. Simpson (1999): "Predicting US business cycle regimes," *Journal of Business and Economic statistics*, 17.

Chin, D., J. Geweke, and P. Miller (2000): "Predicting Turning Points," Research Department Staff Report 267, Federal Reserve Bank of Minneapolis.

Clements, M. P., and D. Hendry (2004): *Forecasting economic time series*. Cambridge University Press.

Estrella, A., and F. Mishkin (1997): "The predictive power of the term structure of interest rates in europa and the united states: implications for the european central bank," *European Economic Review*.

——— (1998): "Predicting US recessions: financial variables as leading indicators," *Review of Economics and Statistics*.

Filardo, A. (2004): "The 2001 US recession: what did recession prediction models tell us?," Working papers 148, Bank of International Settlements.

Findley, D. F., B. Monsell, W. Bell, M. C. Otto, and B. Chen (1998): "New Capabilities and Methods of the X-12-ARIMA Seasonal-Adjustment Program," *Journal of Business and Economic Statistics*, 16(2).

Franses, P., R. Paap, and D. Fok (2005): "Perfromance of Seasonal Adjustment Procedures: Simulation and Empirical Results," Final report of the study on "seasonal adjustment of business and consumer survey data", ECFIN.

[29]'Turning-points' are an important aspect of the trading-problem but the new criterion on which filters are based addresses specifically earnings, losses, risks and transaction costs. First empirical results are obtained in Oro (2006).

Harvey, A. (1989): *Forecasting Structural Time Series Models and the Kalman Filter.* Cambridge University Press, Cambridge.

Ladiray, D., and B. Quenneville (2001): *Seasonal adjustment with the X-11-method.* Lecture notes in statistics. Springer, Berlin.

Lamy, R. (1997): "Forecasting US recessions: some further results from probit models," Working paper, Finance Canada.

Maravall, A., and V. Gomez (1994): "Program SEATS : User Instructions," .

McElroy, T., and M. Wildi (2007): "Specification Tests and Diagnostics in Real-Time Signal Extraction," In preparation.

Monsell, B., J. Aston, and S. Koopman (2003): "Toward X-13?," in *ASA proceedings, business and economic statistics section, alexandria, VA: ASA.*

Oro, A. (2006): "Trading mit Echtzeitfiltern," Masterthesis, University of Applied Sciences, Winterthur.

Sensier, M., M. Artis, D. Osborn, and C. Birchenhall (2005): "Domestic and international influences on business cycle regimes in Europe," *International Journal of Forecasting, forthcoming.*

Stier, W., and M. Wildi (2002): "About Model-Based Times Series Procedures: Some Remarks to TRAMO/SEATS and CENSUS X-12-ARIMA," *Allgemeines Statistisches Archiv*, (86).

Wildi, M. (2004): *Signal extraction: efficient estimation, Unit-root tests and early detection of turning-points*, no. 547. Lecture notes in economics and mathematical systems. Springer.

——— (2007): *Real-Time Signal Extraction: Beyond Maximum Likelihood Principles.* Springer (in preparation).