



Estudios de Economía Aplicada

ISSN: 1133-3197

secretaria.tecnica@revista-eea.net

Asociación Internacional de Economía
Aplicada
España

WILDI, MARC

Real-Time Signal Extraction: a Shift of Perspective

Estudios de Economía Aplicada, vol. 28, núm. 3, 2010, pp. 497-518

Asociación Internacional de Economía Aplicada

Valladolid, España

Available in: <http://www.redalyc.org/articulo.oa?id=30120334001>

- How to cite
- Complete issue
- More information about this article
- Journal's homepage in redalyc.org

redalyc.org

Scientific Information System

Network of Scientific Journals from Latin America, the Caribbean, Spain and Portugal

Non-profit academic project, developed under the open access initiative

Real-Time Signal Extraction: a Shift of Perspective

MARC WILDI

Institute of Data Analysis and Process Design

ZURICH UNIVERSITY OF APPLIED SCIENCES, SWITZERLAND (SUIZA)

e-mail: wia@zhwin.ch

ABSTRACT

Real-time signal extraction (RTSE) concerns the determination of optimal asymmetric filters towards the end of a time series where —otherwise desirable— symmetric filters cannot be applied anymore. The attractiveness of this particular estimation problem resides in the generality of its scope. For illustrative purposes we here stress real-time monitoring of the US-economy as well as multi-step ahead forecasting. Traditionally, the estimation problem addressed by RTSE is tackled in the methodological framework of the classical maximum likelihood paradigm. We here question the adequacy of this general parametric approach. In particular, we review a statistical apparatus —the DFA— consisting of optimization criteria, diagnostics and tests which accounts for alternative user-relevant aspects of the estimation problem. Interestingly, this customization relates to an uncertainty principle which entails a fundamental shift of perspective. As a result, RTSE emerges as an autonomous discipline with proprietary concepts and statistics. With little suggestive power we may interpret the DFA as a generalization of the traditional model-based approach to more general enquiries about the future than the classical one-step ahead inference.

Keywords: Signal, Real-Time Signal Extraction, Customized Criteria, Filter, Amplitude and Time Shift, Uncertainty Principle, Forecasting Competition, Recession Indicator.

Extracción de señal en tiempo real: un cambio de perspectiva

RESUMEN

La extracción de señal en tiempo real (RTSE, Real-time signal extraction) se refiere al establecimiento de filtros asimétricos óptimos en la parte final de una serie de tiempo donde los filtros simétricos, tan útiles en el resto de la serie, no pueden ser empleados. El atractivo de este problema concreto de estimación reside en la amplitud de su alcance. Como ejemplo ilustrativo, nos centramos aquí en el seguimiento en tiempo real de la economía de los EEUU así como en las predicciones dinámicas (multi-step). Tradicionalmente, el problema de estimación abordado por la RTSE se estudió desde el marco metodológico del paradigma de la máxima verosimilitud clásica. Nosotros discutimos la conveniencia del enfoque paramétrico general, y proponemos, en concreto, un procedimiento estadístico (el DFA) que incluye criterios de optimización, diagnósticos y contrastes que permiten valorar otros aspectos del problema de estimación que son relevantes para el usuario. Es interesante señalar que esta personalización está relacionada con un principio de incertidumbre que trae consigo un cambio sustancial de perspectiva. RTSE, en consecuencia, se plantea como un método de trabajo autónomo con conceptos y herramientas estadísticas propias. Con un poco de imaginación, podríamos interpretar el DFA como una generalización del enfoque tradicional basado en modelos orientada a dar respuestas más amplias sobre el futuro de la serie que las que proporciona la inferencia clásica a un paso (one-step).

Palabras clave: Señal, extracción de señal en tiempo real, criterios personalizados, filtro, amplitud y desplazamiento, principio de incertidumbre, predicciones alternativas, indicadores de recesión.

Clasificación JEL: C30, C32.

Artículo recibido en agosto de 2010 y aceptado en septiembre de 2010.

Artículo disponible en versión electrónica en la página www.revista-eea.net, ref. 28309.

1. INTRODUCTION

From a conceptual perspective, real-time signal extraction (RTSE) concerns the estimation of interesting components—for example trends or cycles—by asymmetric filters towards the end of a time series. It is well-known that the factual asymmetry of these filters entails timeliness and reliability issues: real-time estimates are typically delayed and/or noisy. A popular generic solution, formalized by Cleveland [1972] (stationary case) and Bell [1984] (non-stationary case), then consists in applying a symmetric filter to a time series which has been previously extended by forecasts. Therefore, the resources of time series analysis could be invoked to tackle the problem, in principle. This approach led to a major improvement in official seasonal adjustment procedures, see Dagum [1980]. We here propose a different approach and develop concepts and ideas which position RTSE as an autonomous discipline with 'own rules' and proprietary concepts.

A typical application of RTSE delves into the timing—nowcasting—or the prediction of economic down-turns and recession episodes. A recent review of approaches in the field is proposed in Hamilton [2010] with references to important published indicators. Interestingly, though not surprisingly, conceptually different indicator designs generally rely on common maximum likelihood estimation principles which, schematically, derive overall performances from (mean-square) one-step ahead forecasting performances. Wildi [2008.1] suggests that the intrinsic structure of the RTSE-problem is more complex than could possibly be accounted for by relying on one-step ahead forecasting performances alone. Therefore, general optimization criteria are proposed which embrace the problem structure subject to alternative user-relevant aspects. The so-called Direct Filter Approach (DFA) addresses revisions, timeliness and reliability issues. A recent extension to a multivariate setting (MDFA) is proposed in Wildi [2008.2]. Applications to real-time output-gap and recession indicators are proposed in Sturm/Wildi [2008] and Wildi [2009].

The scope of this article is threefold. First, we provide a convenient summary of the main concepts belaying the DFA by collecting disseminated research results. Second, we contrast the scopes of classical maximum likelihood and (M)DFA in order to motivate the invoked shift of perspective. A hopefully interesting outcome of this confrontation is a refreshed view about the concept of 'parsimony'. Third but not least, we position RTSE as an autonomous discipline with proprietary—customized—optimization criteria and statistics. With little additional effort, this approach may be interpreted as a generalization of the classical one-step ahead model paradigm.

The article is structured as follows: the traditional model-based approach is reviewed in section 0. Section 0 introduces customized variants of DFA and MDFA. Section 0 completes the picture by proposing new filter-diagnostics and tests. An extensive study of a recent real-time business-cycle application illustrates customization-, uncertainty- and parsimony-principles and positions RTSE as an

autonomous discipline. Section surveys recent real-time out-of-sample performances of the new approach, including international forecasting competitions and a multivariate real-time recession indicator monitoring the US economy, the USRI¹. Finally, section 0concludes by a brief summary and an outlook.

2. THE TRADITIONAL ONE-STEP AHEAD PARADIGM

For ease of exposition we here stress a linear univariate framework. Assume, for that purpose, that the interesting signal Y_t and the data 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. RTSE is concerned with the estimation of Y_T given afinite sample X_1, \dots, X_T ². Obviously, a real-time estimate \hat{Y}_T could be obtained by replacing unobserved future realizations X_{T+1}, X_{T+2}, \dots by estimates

$$\hat{Y}_T = \gamma_{T-1}X_1 + \dots + \gamma_0X_T + \gamma_1\hat{X}_{T+1} + \gamma_2\hat{X}_{T+2} + \dots \quad (1)$$

Note that we neglected unobserved past realizations X_0, X_{-1} in this expression because filter coefficients often decay rapidly. Cleveland [1972] (stationary case) and Bell [1984] (non-stationary case) have verified the pertinence of this approach under suitable assumptions—in particular, they assume that the data generating process (DGP) is known—, see also Cleveland/Tiao [1976] and Geweke [1978]. Let us rely on a simple example in order to illustrate the topic and assume, for that purpose, that X_t is a random-walk process. Then $\hat{X}_{T+k} = X_T$, $k > 0$ and (1) becomes

$$\begin{aligned} \hat{Y}_T &= \gamma_{T-1}X_1 + \dots + \gamma_0X_T + \gamma_1\hat{X}_{T+1} + \gamma_2\hat{X}_{T+2} + \dots = \\ &= \gamma_{T-1}X_1 + \dots + \gamma_0X_T + \gamma_1X_T + \gamma_2X_T + \dots = \\ &= \gamma_{T-1}X_1 + \dots + \sum_{k \geq 0} \gamma_k X_T \end{aligned}$$

More generally, if the forecasts \hat{X}_{T+k} are arbitrary linear functions of past and present observations, this expression becomes:

$$\hat{Y}_T = \sum_{j=1}^T \left(\gamma_{T-j} + \sum_{k=1}^{\infty} \gamma_k a_{T+k,j} \right) X_j \quad (2)$$

¹ <http://www.idp.zhaw.ch/usri>.

² We neglect estimates of Y_{T-k} , i.e. smoothing, which could be addressed as well.

where $a_{T+k,j}$ are the coefficients of X_j , $j=1,\dots,T$, in the (linear) forecasting function of X_{T+k} , see Stier/Wildi [2002].

It is likely that the strong intuitive appeal of the above approach has contributed to its widespread diffusion. There are, however, two problems at stake: the solution is optimal in a mean-square sense only—which potentially conflicts with reliability and/or timeliness issues, see section 3—and the DGP is generally unknown. In the latter case, models must be identified, parameters must be estimated and performances must be assessed by relying on suitable diagnostics. Interestingly, the conventional statistical apparatus addresses one-step ahead mean-square performances ‘only’ in all three modeling-phases³. In contrast, the expression (1) addresses a linear combination of one- and multi-step ahead forecasts. This mismatch of the statistical apparatus is at the origin of inefficiencies as revealed and quantified in Wildi [2008.1], chapter 4. McElroy and Wildi [2010.1] propose *model*-diagnostics that emphasize specifically the RTSE-problem and McElroy and Wildi [2010.2] propose generalized model-based estimation principles addressing multi-step ahead forecasting. In contrast, section 0 in this article proposes *concurrent-filter* diagnostics.

3. CUSTOMIZATION AND THE DIRECT FILTER APPROACH

We first review a univariate criterion which emphasizes the revision error variance. Then we propose a generalization which allows for a formal operationalization of the somehow ‘fuzzy’ timeliness and reliability concepts. Finally, a multivariate extension is presented.

The intricate convolutional form of the estimation problem in (2) suggests the pertinence of a frequency-domain approach. Denote by $\Gamma(\cdot)$ and $\hat{\Gamma}(\cdot)$ the transfer functions of symmetric and real-time filters with outputs Y_t and \hat{Y}_t . For stationary processes X_t , the mean-square filter error—the revision error variance—can be expressed as

$$\int_{-\pi}^{\pi} |\Gamma(\omega) - \hat{\Gamma}(\omega)|^2 dH(\omega) = E[(Y_t - \hat{Y}_t)^2] \quad (3)$$

where $H(\omega)$ is the spectral distribution of X_t . For a statistical agency, minimizing the mean-square filter error with respect to the parameters of $\hat{\Gamma}(\cdot)$ would be an appealing criterion. The interpretation of (3) is intuitively appealing as well: the strength of the transfer function fit is modulated by the spectral mass of the

³ Information criteria and estimation emphasize the one-step ahead error variance, diagnostics put forward model residuals i.e. one step ahead errors.

process. Unfortunately, the spectral distribution is generally unknown. Therefore, we may consider an alternative finite sample approximation of the above integral

$$\frac{2\pi}{T} \sum_{k=-T/2}^{T/2} |\Gamma(\omega_k) - \hat{\Gamma}(\omega_k)|^2 I_{TX}(\omega_k) \approx \frac{1}{T} \sum_{t=1}^T (Y_t - \hat{Y}_t)^2 \quad (4)$$

where $\omega_k = k2\pi/T$, is a discrete frequency support in the interval $[-\pi, \pi]$ and

$$I_{TX}(\omega_k) = \frac{1}{2\pi T} \left| \sum_{t=1}^T X_t \exp(-it\omega_k) \right|^2$$

is the periodogram of the input signal X_t . Note that the right-hand side of (4) is not observable directly since Y_t is assumed to be the output of a bi-infinite filter. Wildi [2005], [2008.1] (proposition 10.17) and [2008.2] (proposition 7.4) show that the approximation error in the above expression is of smaller order than the classical $1/\sqrt{T}$ —rule under relatively mild assumptions so that the frequency—domain estimate on the left of (4) is a super consistent estimate of the (unobservable) sample revision error variance. Unlike (3), this estimate is observable. Therefore we may attempt to minimize (4) with respect to the parameters of $\hat{\Gamma}(\cdot)$.

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

Wildi [2008.1] (proposition 10.17), [2008.2] (proposition 7.4) show that the super consistency in (4) applies uniformly under mild regularity assumptions and therefore the solution of (5) minimizes the finite sample revision error variance up to a negligible error term. Since the sample variance can be shown to be an asymptotically efficient estimate of the revision error variance (3), see Wildi [2008.1] proposition 10.16, the interpretation of (5) is straightforward in statistical terms: the solution minimizes a super consistent estimate of an asymptotically efficient estimate of the revision error variance. In the words of Chin/Geweke [2000] the proposed criterion “attempts to enhance accuracy by forecasting directly what is of interest”. Performances of this criterion in the context of business-surveys are documented in Wildi [2008.1]⁴. Applications in the context of recent international forecasting competitions are addressed in section 0. Generalizations to non-stationary (integrated) processes are proposed in Wildi [2008.1], chapter 6.

⁴ Efficiency gains between 20% and 30% are observed for various business survey data when compared to X-12-ARIMA or TRAMO-SEATS.

Criterion (5) addresses specifically applications emphasizing the revision error. In order to address alternative user-requirements, Wildi [1998], [2004] and [2008.1] propose a generalized criterion which relies on a simple trigonometric decomposition of the transfer function error:

$$\begin{aligned} |\Gamma(\omega) - \hat{\Gamma}(\omega)|^2 &= A(\omega)^2 + \hat{A}(\omega)^2 - 2A(\omega)\hat{A}(\omega)\cos(\hat{\Phi}(\omega) - \Phi(\omega)) = \\ &= (A(\omega) - \hat{A}(\omega))^2 + 2A(\omega)\hat{A}(\omega)[1 - \cos(\hat{\Phi}(\omega) - \Phi(\omega))] \end{aligned} \quad (6)$$

Where $\hat{A}(\omega)$ and $\hat{\Phi}(\omega)$ are amplitude and phase functions of $\hat{\Gamma}(\omega)$ i.e. $\hat{\Gamma}(\omega) = \hat{A}(\omega) \exp(-i\hat{\Phi}(\omega))$ (and similarly for A , Φ and Γ). If we assume that Γ is symmetric then $\Phi(\omega) \equiv 0$ (we ignore the practically irrelevant case $\Phi(\omega) \equiv \pi$ which may occur due to a change of sign by the symmetric filter). Inserting (6) into (5) leads to

$$\frac{2\pi}{T} \sum_{k=-T/2}^{T/2} (A(\omega_k) - \hat{A}(\omega_k))^2 I_{TX}(\omega_k) \quad (7)$$

$$+ \frac{2\pi}{T} \sum_{k=-T/2}^{T/2} 2A(\omega_k)\hat{A}(\omega_k)[1 - \cos(\hat{\Phi}(\omega_k))] I_{TX}(\omega_k) \quad (8)$$

As shown in Wildi [2008.1], section 3.1.3, the expression (7) measures that part of the (estimated) mean-square filter error which is attributable to the amplitude function of the real-time filter: an imperfect amplitude-fit will generate an error of the corresponding magnitude assuming an otherwise perfect phase-fit. The expression (8) measures that part of the filter error which corresponds to undesirable phase-artifacts: an imperfect phase-fit will generate an error of the corresponding magnitude assuming an otherwise perfect amplitude-fit. Note that the phase-effect is revealed by a time-shift, ordinarily a delay, of \hat{Y}_t when compared to Y_t . We now refer to this effect by the wording ‘time-shift’. Formally, the time-shift $\phi(\cdot)$ is related to the phase $\Phi(\cdot)$ by $\phi(\omega) = \Phi(\omega)/\omega$, see Wildi [2008.1], chapter 2.

The main advantages of the proposed error-decomposition are twofold: first, the disentanglement of amplitude and time-shift error components enables to ‘act’ on both dimensions separately. Second, the decomposition offers the opportunity to operationalize the somehow fuzzy ‘reliability’ and ‘timeliness’ concepts in a formal way: reliability is achieved by emphasizing the strength of the amplitude matching (7) and timeliness is achieved by emphasizing the time-shift error (8). It should be clear that equivalent error-decompositions and formalizations could not be straightforwardly achieved in the time-domain.

In order to take advantage of the offered opportunities, Wildi [2005] proposes the following generalized criterion:

$$\begin{aligned} & \frac{2\pi}{T} \sum_{k=-T/2}^{T/2} W(\omega_k) (A(\omega_k) - \hat{A}(\omega_k))^2 I_{TX}(\omega_k) \\ & + \lambda \frac{2\pi}{T} \sum_{k=-T/2}^{T/2} 2W(\omega_k) A(\omega_k) \hat{A}(\omega_k) [1 - \cos(\hat{\Phi}(\omega_k))] I_{TX}(\omega_k) \rightarrow \min_{\hat{\cdot}} \end{aligned} \quad (9)$$

Applications to real-time business-cycle monitoring and output-gap measure can be seen in Wildi [2008.1], chapter 5, and Elmer/Wildi [2008]. For $\lambda=1$, $W(\cdot) \equiv 1$ the mean-square revision error criterion (5) results. For $\lambda > 1$ the time-shift error is magnified: everything else being equal, a real-time filter with a larger λ will convey signals, for example turning-points, earlier but the output will be noisier. Note that the weighting-term $A(\omega_k) \hat{A}(\omega_k)$ appearing in the time-shift error component (8) implies that λ is essentially acting on the pass-band of the filter (the product of both amplitude functions can be interpreted as a scaling of the dimensionless consecutive phase-expression). In contrast, the purpose of $W(\omega)$ then consists in emphasizing the fit of the amplitude function in the stop-band: attributing more weight to the stop-band ensures a stronger noise-rejection. Therefore, the reliability of the real-time estimate is improved. Everything else being equal, a stronger noise-rejection results in increased time-delay. The novelty resides in the fact that both aspects, timeliness and reliability, can be addressed simultaneously in a very flexible way: a particular user can set the weighting scheme $\lambda, W(\cdot)$ in order to suit his particular priorities as well as his individual risk profile. It is important, as well, to note that the general criterion (9) addresses a fundamental uncertainty-principle: noise-rejection and time-shift properties are entangled and cannot be improved arbitrarily in either direction. Emphasizing this fundamental uncertainty principle reflects the intrinsic structure of the underlying real-time signal extraction problem and it accounts for important user-requirements in a very flexible way.

The above criteria have been recently generalized to a multivariate setting, see Wildi [2008.2]. It is worth mentioning that the additional cross-sectional dimension of the filtering problem brings new aspects into consideration such as, for example, cointegration. Assume, for now on, that the estimation of $Y_T = \sum_{k=-\infty}^{\infty} \gamma_k X_{T-k}$ can be based on observations X_t as well as on additional time series W_{1t}, \dots, W_{mt} , $t=1, \dots, T$. For illustrative purposes, we here assume non-stationary time series (integration order one) subject to a single cointegration relation⁵

$$C_t = X_t - \alpha_1 W_{1t} - \dots - \alpha_m W_{mt} \sim I(0) \quad (10)$$

⁵ Theorem 7.1 in Wildi [2008.2] proposes solutions for general settings including stationarity as well as arbitrary cointegration ranks.

Under these circumstances, the following estimation criterion generalizes (5) to a multivariate setting:

$$\frac{2\pi}{T} \sum_{k=-(T-1)/2}^{(T-1)/2} |\Xi_{Tr}(\omega_k)|^2 \rightarrow \min_{\hat{\Gamma}_X, \hat{\Gamma}_{W_1}, \dots, \hat{\Gamma}_{W_m}} \quad (11)$$

where

$$\begin{aligned} \Xi_{Tr}(\omega_k) = & \left\{ \Delta\Gamma_X(0) \Xi_{TC}(\omega_k) - \right. \\ & - \exp(i\omega_k) \left[\frac{\Delta\Gamma_X(0) - \exp(-i\omega_k) \Delta\Gamma_X(\omega_k)}{1 - \exp(-i\omega_k)} - \Delta\Gamma_X(0) \right]^+ \Xi_{T\Delta X}(\omega_k) + \\ & + \left[\frac{\Delta\Gamma_X(0) - \exp(-i\omega_k) \Delta\Gamma_X(\omega_k)}{1 - \exp(-i\omega_k)} \right]^- \Xi_{T\Delta X}(\omega_k) + \\ & \left. + \sum_{h=1}^m \exp(i\omega_k) \left[\frac{\hat{\Gamma}_{W_h}(0) - \exp(-i\omega_k) \hat{\Gamma}_{W_h}(\omega_k)}{1 - \exp(-i\omega_k)} - \hat{\Gamma}_{W_h}(0) \right] \Xi_{T\Delta W_h}(\omega_k) \right\} \quad (12) \end{aligned}$$

In this expression $\Delta\Gamma_X(\omega_k) = \Gamma(\omega_k) - \hat{\Gamma}_X(\omega_k)$ corresponds to the difference between the bi-infinite and the real-time filter in criterion (5) and $\hat{\Gamma}_{W_h}(\omega_k)$, $h=1, \dots, m$ are the transfer functions of the real-time filters applied to the additional ‘explaining’ series W_{ht} . $\Xi_{T\Delta X}(\omega_k)$ and $\Xi_{T\Delta W_h}(\omega_k)$ are the discrete Fourier transforms of the (stationary) first differences of the time series. Note that we cannot rely on ‘periodograms’ anymore, as in (5), because the latter do not support the important phase-shift information between adjacent time series. Finally, $\Xi_{TC}(\omega_k)$ is the discrete Fourier transform of the stationary linear combination C_t of the integrated time series (10). The notational convention expressed by ‘plus’ and ‘minus’ exponents on the squared brackets signifies that the corresponding bi-infinite power-series expansions are truncated either to positive or to negative exponents (recall that $\Gamma(\cdot)$ is assumed to be a symmetric bi-infinite filter), see Wildi [2008.2], theorem 7.1 for details. In analogy to (5), the multivariate criterion (11) can be interpreted as the minimization of an efficient estimate of the revision error variance given additional evidence supplied by W_{1t}, \dots, W_{mt} .

The spectral content relayed by the discrete Fourier transforms in (12) relates to stationary processes which complies with standard theory. The presence of the integration operator $1/(1 - \exp(-i\omega_k))$ in the squared brackets, however, suggests the pertinence of so-called pseudo spectral and cross pseudo spectral estimates. Theorem 7.1 in Wildi provides a formal framework for a generalization of classical spectral factorization results to the multivariate RTSE-problem in the case of

integration and cointegration. In the absence of cointegration relations one must impose a set of $m + 1$ filter constraints in the unit-root frequency zero which are

$$\hat{\Gamma}_X(0) = \Gamma(0) \quad (13)$$

$$\hat{\Gamma}_{W_h}(0) = 0, \quad h = 1, \dots, m \quad (14)$$

The first constraint (13) ensures a close track of the trend of Y_t by the output of $\hat{\Gamma}_X(\cdot)$. The latter m restrictions (14) ensure that the uncommon trends in the explaining variables are eliminated (assuming some mild regularity assumptions). Together, these restrictions are necessary to ensure an asymptotically finite revision error variance of the multivariate filter. In the case of the above (single) cointegration constraint, Wildi [2008.2] derives an equivalent transcription of (10) in terms of real-time filter constraints:

$$\hat{\Gamma}_{W_i}(0) = \alpha_i(\Gamma(0) - \hat{\Gamma}_X(0)) \quad (15)$$

for $i = 1, \dots, m$. As can be seen, (15) is satisfied by imposing (13) and (14). But then the cointegration term $\Delta\Gamma_X(0)\Xi_{TC}(\omega_k)$ in (12) would disappear. Conversely, the ‘power’ of the cointegration term as measured by $\Xi_{TC}(\omega_k)$ in (12), determines to which extent $\Delta\Gamma_X(0) = 0$ might appear as a useful constraint. As seen, (15) allows for the relaxation of a constraint by providing an additional degree of freedom. Formally, a cointegration relation augments the rank of the space of admissible filter solutions. The increased flexibility can then be used to improve real-time performances. If all series are stationary then all constraints in frequency zero are relaxed. Note that some mild regularity assumptions proposed in Theorem 7.1 in Wildi [2008.2] are necessary in order to derive the above statements and, in particular, in order to remove singularities on the right-side term of (12) properly. A generalization of (9) to the presented multivariate setting has been proposed recently in Wildi [2010].

4. A PARADIGM SHIFT

The previous section affirmed the originality of the (M)DFA by establishing customized optimization principles. In this section we complete the picture by proposing a set of customized diagnostic —and test—statistics. The resulting statistical apparatus confers an own ‘model-based’ interpretation to the DFA. The pertinence of our approach is illustrated by relying on a recent real-time business-cycle application. Based on the collected evidences we then propose to revisit the famous parsimonious principle. To conclude, we reveal that the RTSE-problem allows for a convenient circumvention of the delicate DGP-identification affecting


model-based approaches. For that purpose we invoke an implicit smoothing argument inherent to the functional (3) and inherited by (5). This distinction achieves the individualism of our approach.

4.1. Filter Diagnostics

The Swiss Institute of Business-Cycle Analysis (<http://www.kof.ethz.ch/>) administers a large-scale business-survey. The following time series, referred to as series '27', belongs to this survey and enters into the design of an important leading indicator for Swiss-GDP. The sample period in our example extends from February 1994 to January 2008 (14 years). In order to measure performances we had to apply symmetric filters and therefore two years of observations were lost at begin and end of the time series: the following plots refer to the shorter truncated sample of length 10 years or 120 months. We rely on this particular example because it lends itself for illustrative purposes and because it has been extensively analyzed in Wildi [2008.1]. The statistical task defined by the economic staff consisted in extracting the low-frequency content of the series by a particular lowpass filter. The data is plotted in figure 1 together with the signal (output of the symmetric filter) and the periodogram. The latter suggests the presence of a cycle component⁶ as well as of a seasonal component. The automatic identification procedure of TRAMO⁷ proposes the following well-known airline model

$$(1 - B)(1 - B^{12})X_t = (1 - 0.3146B)(1 - 0.8565B^{12})\varepsilon_t \quad (16)$$

FIGURE 1
Linearized series, symmetric trend (shaded) and Periodogram.



⁶ The periodogram vanishes in frequency zero because the series is centered.


⁷ <http://www.bde.es/webbde/es/secciones/servicio/software/econom.html>.

Classical model diagnostics (Ljung-Box, ACF and PACF of residuals, not shown here) confirm the adequacy of the model. The DFA-filter is optimized according to (5). It is an ARMA (15,15)-design:

$$\begin{aligned} \hat{y}_t = & 1.1221\hat{y}_{t-1} + 0.6462\hat{y}_{t-2} - 2.0422\hat{y}_{t-3} + 1.2562\hat{y}_{t-4} + 0.6435\hat{y}_{t-5} - \\ & - 1.8344\hat{y}_{t-6} + 1.1496\hat{y}_{t-7} + 0.5889\hat{y}_{t-8} - 1.6788\hat{y}_{t-9} + 1.0521\hat{y}_{t-9} + \\ & + 0.5390\hat{y}_{t-10} - 0.8349\hat{y}_{t-11} + 0.1755\hat{y}_{t-12} + 0.1758\hat{y}_{t-13} + 0.0400\hat{y}_{t-14} + \\ & + 0.0264\hat{y}_{t-15} + 0.3511x_t - 0.1695x_{t-1} - 0.3195x_{t-2} + 0.5272x_{t-3} - \\ & - 0.1755x_{t-4} - 0.3104x_{t-5} + 0.5174x_{t-6} - 0.1755x_{t-7} - 0.3104x_{t-8} + \\ & + 0.5174x_{t-9} - 0.1755x_{t-10} - 0.3104x_{t-11} + 0.1662x_{t-12} - 0.006x_{t-13} + \\ & + 0.0091x_{t-14} - 0.0099x_{t-15} \end{aligned} \quad (17)$$

where x_t is the input series 27 (details of the parameterization are postponed to the next section). Amplitude and time-shift functions of DFA (blue) and TRAMO (red) real-time filters⁸ are plotted in figure 2. For comparative purposes, the transfer function of the symmetric low pass is plotted in green.

FIGURE 2
Amplitude and time shifts DFA (solid) vs. MBA (dotted).



As expected, amplitude as well as time-shift deviate from the idealized target. The amplitude function of the model-based approach (MBA) seems to fit the green line better than the DFA⁹. On the other hand, the time-shift of the DFA seems to

⁸ The model-based real-time filter relies on (1) whereby forecasts were computed according to the above airline model.

⁹ Due to model unit-roots, the amplitude function of TRAMO must satisfy several restrictions: $\hat{A}(0)=1$ and $\hat{A}(k\pi/6)=0$, $k=1,\dots,6$. In contrast, the amplitude function of the DFA is smaller in frequency zero and the seasonal harmonic in $4\pi/6$ is not removed because the periodogram does not suggest evidence of such a component.


outperform TRAMO since the former is smaller in the pass band. Note that the time-shifts of both filters vanish in frequency zero i.e. the corresponding phase functions have a zero of order two. In order to assess filter performances, we rely on the proposed decomposition of the (estimated) mean-square filter error in (7) and (8), see table 1. Note that this decomposition is general and can be applied to any filter, including the model-based design. Results are as expected: the DFA performs slightly worse with respect to its amplitude-fit but it outperforms TRAMO with respect to time-shift characteristics. Since the gain provided by the time-shift exceeds the loss incurred by the amplitude, the DFA wins overall. More refined details about the balancing effect of the uncertainty principle 'at work' can be found below.

TABLE 1
Decomposition of mean-square filter error: DFA vs. MBA.

Series 27	Total	Amplitude	Time Delay
DFA	7.94	6.66	1.28
MBA	9.91	4.47	5.44

A direct comparison of filter errors and of periodograms thereof in fig. 3 reveals further instructive insights. We can infer that the revision errors of the model-based approach suffer from an excessively large low-frequency spectral peak which is due to the larger time-shift: the lagging track of the dominant business-cycle in series 27 by the TRAMO-filter induces a spurious cycle in the filter errors whose frequency thus replicates the business-cycle.

FIGURE 3
Filter errors with periodograms: DFA (solid) vs. MBA (dotted).



The periodogram of the DFA-filter error confirms a slight excess in the stop-band just on the left of the first seasonal frequency $\pi/6$ which is due to a poorer amplitude match there. Remarkably, the DFA-filter seems to outperform the model-based approach towards higher frequencies which results in a smoother filter output. All in all, the DFA is smoother, faster and its revisions are smaller (in a mean-square sense).

The proposed filter diagnostics —amplitude and time-shift of filters and periodograms of filter errors— complement each other meaningfully. From our perspective, they address the RTSE problem more comprehensively than classical (one-step ahead) model diagnostics. In particular, an unsatisfactory filter design could be easily revised by emphasizing locally defective amplitude and/or time-shift fits in the framework of the generalized criterion (9). An instructive example is provided in Wildi [2008.1], section 4.3.2.

4.2. Filter Deconstruction and Parsimony

It is obvious, at first sight, that the imposing ARMA (15,15) filter (17) seems to offend common precautionary principles as entailed by the parsimony concept generally advocated in time series analysis. Therefore, we here propose an in-depth analysis of the rationale behind the proposed design. For that purpose, we ‘atomize’ the DFA-filter into elementary sub-filters whose individual and combined effects are subsequently analyzed: this extended in-depth diagnostic check is called *deconstruction*.

It is convenient to represent the ARMA-filter equation (17) in terms of poles and zeroes. In order to illustrate this alternative coding of the filter we rely on the following ARMA (1,1) toy-example:

$$\hat{y}_t = a\hat{y}_{t-1} + b_0x_t + b_1x_{t-1}$$

The transfer function of the filter is

$$\hat{\Gamma}(\omega) = \frac{b_0 + b_1 \exp(-i\omega)}{1 - a \exp(-i\omega)}$$

The numerator polynomial is $b_0 + b_1x$ with root $Z = -b_0/b_1$. Similarly, the denominator polynomial is $1 - ax$ with root $P = 1/a$. Therefore, an equivalent representation of the transfer function is

$$C \frac{Z - \exp(-i\omega)}{P - \exp(-i\omega)}$$

where $C = -b_1/a$, $P = 1/a$ and $Z = -b_0/b_1$. More generally, the transfer function of an arbitrary ARMA-filter can be alternatively represented as

$$\hat{\Gamma}(\omega) = \frac{\sum_{k=0}^q b_k \exp(-ik\omega)}{1 - \sum_{k=1}^Q a_k \exp(-ik\omega)} = \quad (18)$$

$$= C \frac{\prod_{j=1}^n (Z_{2j-1} - \exp(-i\omega)) (Z_{2j} - \exp(-i\omega))}{\prod_{k=1}^{n'} (P_{2k-1} - \exp(-i\omega)) (P_{2k} - \exp(-i\omega))} \cdot \frac{\prod_{j=2n+1}^q (Z_j - \exp(-i\omega))}{\prod_{k=2n'+1}^Q (P_k - \exp(-i\omega))} \quad (19)$$

where poles can be grouped into complex conjugate pairs $P_{2k-1} = \bar{P}_{2k}$, $k=1, \dots, n'$ and real poles P_k , $k=2n'+1, \dots, Q$ (and similarly for zeroes) and where a_k, b_k denote AR —and MA—parameters of the filter respectively. The purpose of this equivalent coding of the filter is that useful constraints can be defined and imposed much more easily on poles and zeroes than on AR —and MA—parameters. As an example, the important stability of a filter is ensured by requiring $|P_k| > 1$ for all k : in the above ARMA (1,1) toy-example this requirement would ensure $|a| > 1$ which is the classical stationarity assumption of ARMA-models. As a consequence, it is in some sense natural to rely on (19), rather than (18), for optimization purposes. The length and the arguments of the zeroes Z_k and of the poles P_j belonging to the ARMA (15,15)-filter (17) are reported in table 2.


TABLE 2
Series 27: Parameters of ZPC-Filters (without normalizing constant).

	ZPC1	ZPC2	ZPC3	ZPC4	ZPC5	ZPC6	ZPC7	ZPC8
$ Z $:	1.1028	1.9904	1	1	1	2.8526	1	1
$ P $:	1.1077	1.2775	1.03	1.03	1.03	3.9488	1.03	1.03
$\text{Arg}(Z) = \text{Arg}(P)$:	0	$\pm\pi/7.9024$	$\pm\pi/6$	$\pm 2\pi/6$	$\pm 3\pi/6$	$\pm 4\pi/6$	$\pm 5\pi/6$	$\pm\pi$

We see that zeroes and poles are arranged in pairs with common arguments. This particular filter design —ZPC or Zero Pole Combination— was first introduced in Wildi [2005]. The main advantage of ZPC-filters can be seen in the straightforward interpretability of the available —three— degrees of freedom: they can account for the location, the height and the width of spectral peaks, see Wildi [2008.1], section 3.2.2 (note that this restriction could not be easily imposed upon AR- and MA-parameters). ZPC2 to ZPC7 occur in complex conjugate pairs whereas ZPC1 and ZPC8 are real. Since the latter is split, as suggested by the $\pm\pi$ argument, we obtain

effectively an ARMA $(2 \cdot 7 + 1, 2 \cdot 7 + 1)$ -structure. The number of estimated parameters in table 2 is $2 \cdot 8 + 1 = 17$ corresponding to the first two rows and the argument of ZPC2 (for ease of exposition, the normalizing constant C has been concatenated with ZPC1). The other arguments, in particular the seasonal frequencies, are fixed. It may be noteworthy to emphasize that poles number 3, 4, 5, 7 and 8 in the second row of table 2 have attained pre-specified stability constraints: the minimal length 1.03 of the poles has been imposed by a simple rule which depends on the sample length alone, see Wildi [2008.1], section 3.5.4.

FIGURE 4
Amplitude and time shifts: first ZPC-filter.



Moreover, the corresponding zeroes in the first row lie on the unit circle (omitting rounding-errors). Therefore, $5 \cdot 2 = 10$ degrees of freedom are 'unutilized' since the boundaries of the parameter space are attained. As a result, the DFA ARMA (15,15)-filter effectively relies on seven freely determined parameters 'only'. We now attempt to interpret these degrees of freedom by isolating the effect of each elementary ZPC-filter, acknowledging that the overall effect is obtained by serial linkage (thus the individual time-shifts add to the overall shift). Amplitude and time shift functions of ZPC1 are plotted in fig. 4. The contribution of this particular sub-filter to the whole design is the realization of a desirable vanishing time-shift in frequency zero which explains, incidentally, the corresponding restriction of its argument in the above table. We impose this particular time-shift restriction in all our low pass filters because we want to track trends closely. Note, en passant, that the vanishing time-shift of the model-based filter is a consequence of the integration order two imposed by the identified airline model¹⁰. Extensive

¹⁰ The slope of an $I(2)$ -process is asymptotically unbounded —the difference is $I(1)$ —. Therefore, a non-vanishing time-shift in frequency zero would result in an asymptotically unbounded revision error variance, see Wildi [2008.1] chapter 6 for a formal exposition. It should be clear that the


experience in the context of business-surveys suggests that this restriction is useful even in the absence of a marked trend. As expected, the time-shift generated by ZPC1 is negative in order to compensate additively for the delaying effect of the consecutive 'smoothing' filters ZPC2, ..., ZPC8. The amplitude effect of ZPC1 appears to be negligible. The amplitude function of ZPC2 can be seen in fig. 5 (note that the filter has not been normalized yet). The three degrees of freedom of ZPC2 can account for the location, the height and the width of the visible amplitude peak and thereby ZPC2 contributes to a simultaneous match of amplitude and time shift requirements in the pass band. We here develop this argument further in order to highlight the uncertainty principle at work. As seen in figure 2 the 'tip of the nose' of the ARMA (15,15) amplitude function is slightly misaligned and overshooting. This local mismatch was penalized by an excessive amplitude error, see table 1. The periodogram of the filter errors in fig. 3 confirmed a corresponding excess of spectral power just on the left of the fundamental seasonal frequency $\pi/6$. But this apparent misalignment of the amplitude function goes hand in hand with beneficial time-shift properties towards the spectral mass of the dominant cycle whose gain, in (8), overcompensate amplitude losses incurred in (7), see table 1. This subtle balance between amplitude and time-shift error contributions towards the spectral mass of the dominant business-cycle reflects the mechanism of the underlying uncertainty principle. Note that the model-based approach is unable to account for these nuances since the one-step ahead criterion could not resolve the cycle information properly. Instead, a false double unit-root is imposed whose misspecification cannot be identified by traditional model diagnostics¹¹. Obviously, the decreasing amplitude function of ZPC2 in the stop-band contributes to the amplitude-fit of the composed ARMA (15,15) filter. We infer that the task assigned by the criterion (5) to the three degrees of freedom of ZPC2 is more complex than for the previous ZPC1. Amplitude and time shift functions of ZPC3 (and complex conjugate) in fig. 6 are typical for a seasonal adjustment filter. Its two degrees of freedom —lengths of pole and zero— can account for the depth and the narrowness of the amplitude dip of the filter or, equivalently, for the height and the width of the corresponding seasonal component (as noted above: these two degrees of freedom were resorbed by natural boundary constraints). It is a particularity of our approach that the characteristics of the seasonal fundamental and harmonics - height and widths of the peaks - can be fitted individually which amounts to $6 \times 2 = 12$ degrees of freedom in the stop band. The remaining $17 - 12 = 5$ degrees in table 2 attempt to match the pass band. Besides ensuring a more effective adjustment, the precise removal of seasonal peaks alleviates undesirable time-shift effects in the pass band. The consecutive

whole discussion is in some sense completely irrelevant because series 27 is bounded by construction (percentage of affirmative respondents).

¹¹ We tried stationary model alternatives with an additional AR(2)-cycle component to account for these difficulties but possible candidates performed worse than the airline model. This suggests that the problem is inherent to the estimation criterion.

seasonal adjustment sub-filters ZPC4, ..., ZPC8 are omitted from the present analysis because their effects do not provide new insights.


FIGURE 5
Amplitude and time shifts: second ZPC-filter.



A brief summary of the collected evidences is in order here. First, the 30 parameters of the original DFA-ARMA (15,15) filter (17) reduce to only seven effective degrees of freedom. All other degrees are either lost in 'useful' design-constraints—ZPC design—or in a priori constraints—seasonal frequencies—or they are resorbed in natural boundary or stability constraints. Note, however, that the latter dormant parameters could be reactivated by the criterion if needed. The interesting point, confirmed throughout our extensive experience in the domain, is that the proposed customized criteria can make full use of the flexibility of the richly parameterized ARMA-ZPC design without being affected by redundant parameters. A crucial issue in the success of this strategy is that parameters are assigned to precise tasks which are 'meaningful' in a general perspective i.e. out-of-sample: they account for the location, the height and the width of dominant spectral peaks in the time series¹² and they address directly a fundamental uncertainty principle. The resulting task-orientation in conjunction with the precise matching of the problem structure by the proposed criteria could give the parsimony principle a new refreshing interpretation.

¹² In the case of inexistent peaks, like the missing harmonic in $\pi/4$ in series 27, unutilized parameters typically cancel in the ARMA-equation so that the corresponding ZPC-filter degenerates to an identity. Note that ZPC6 did not 'degenerate' because the filter adds a useful touch of smoothing in the stop-band which enforces the effect of ZPC2 towards the higher frequencies.

FIGURE 6
Amplitude and time shifts: third ZPC-filter



4.3. Tests: The DFA in a Model-Based Perspective

The apparent lack of a model for the DGP and the presence of the 'nasty' inconsistent periodogram in the main optimization criteria might suggest that the DFA is a non-parametric approach. There is no formal evidence supporting such a view and, in fact, a closer look at the assumptions of efficiency results in Wildi [2008.1] (theorem 10.14, propositions 10.16 and 10.17) refutes this argument: efficiency claims are anchored in a framework set-up by regularity claims and distributional assumptions which could be interpreted in terms of 'model assumptions', in principle. We here briefly develop this topic and derive a provocative interpretation of the DFA in terms of a generalization of the traditional model-based approach. For this purpose it is convenient to assume a very particular 'signal', namely the (asymmetric) h -step forward-shift operator

$$\Gamma(\omega) = \exp(ih\omega) \quad (20)$$

Under this circumstance, the uniform super consistency-argument applying to (4) implies that (5) would minimize the mean-square h -step ahead forecast error —up to a negligible error term. Assuming $h = 1$ —one step ahead perspective— we infer that efficiency properties of DFA and model-based approaches would coincide, asymptotically. More precisely, asymptotic distributions of DFA filter parameters and of one-step ahead model-parameters would coincide under otherwise similar conditions. Real-time 'filter' and one-step ahead 'model' concepts would melt together. A generalization to RTSE, obtained by relaxing (20) to general bi-infinite signals, is proposed in Wildi [2008.1]. Specifically, the asymptotic distribution of real-time filter parameters is discussed in section 3.4. Common hypothesis-testing (t- and F-tests) could be derived from this result.

Chapter 6 derives a test-statistic for the crucial filter constraint (13) in the unit-root frequency zero, the so-called τ -statistic. The resulting test accounts for one —and multi-step ahead forecast performances as well as for the particular signal specification. Interestingly, the proposed general statistics and distributions dwarf asymptotically into traditional model-based expressions under the assumption (20). In particular the τ -statistic ‘degenerates’ to an ADF-test, see Wildi [2008.1], section 6.6.3. This asymptotic morphism of the DFA to a traditional model-based perspective under the signal definition (20) confirms the pertinence of our initial claim: the DFA unfolds the traditional model-based perspective to more complex forecasting problems, including RTSE.

In a similar vein, we briefly address our choice of the inconsistent spectral estimate, the periodogram, in the DFA-criteria. Besides the aforementioned efficiency proofs —and sufficiency arguments alleged in Bretthorst [1988]— we may invoke an important distinction: our target is to compute real-time filters, not spectral distributions. For that purpose we seek to approximate a functional (3) of the spectral distribution, not the distribution itself. Luckily, the random component of the (pointwise inconsistent) periodogram are smoothed out by the summation-filter (5). This observation puts the whole argument on an implicit meta filter-level: the criterion (5) itself acts as a ‘benevolent’ smoother which relaxes the problem structure it has to solve. These nuances highlight a fundamental methodological difference: the DFA circumvents the delicate DGP-identification by emphasizing directly the relevant functional of the unknown spectral distribution. In principle, a model is assumed implicitly by addressing the convolutional expression (2) in the time-domain; it is not required explicitly, however. The admissible relaxation of the problem structure increases flexibility and refocuses efforts on the core of the estimation problem.

5. PERFORMANCES

In order to witness the pertinence of the previous analysis we rely on recent empirical results involving univariate DFA as well as multivariate MDFA applications. For the former, we decided to participate to recent international forecasting competitions —NN3 (2007) and NN5¹³ (2008)— supported by the International Institute of Forecasters¹⁴ and by SAS. The competitions relied on sets of 111 time series each: monthly data for NN3 ($T \approx 100$, forecasting horizon up to 18-steps ahead) and daily data ($T \approx 750$, forecasting horizon up to 56-steps ahead) for NN5. The observed participation-rate was unusually high —more than 60

¹³ See <http://www.neural-forecasting-competition.com/NN3/> and <http://www.neural-forecasting-competition.com/NN5/>. Despite the suggestive - idiosyncratic —numbering, neither NN1, NN2 or NN4 were ever conducted to our knowledge.

¹⁴ <http://www.forecasters.org/>.

contenders— and the diversity of approaches and methods was on par with this figure. One could enumerate, among others, traditional parsimonious approaches, ranging from various naive benchmarks to classical variants of exponential smoothing and the theta-model (winner of the famous M3 competition), ARIMA-based approaches, including for example X-12-ARIMA, highly customized commercial packages (Autobox, Forecast Pro, runner-up in the M3 competition) and so called ‘computational intelligence’ approaches (including neural nets, of course). The DFA outperformed all contenders on both settings two years in sequence. A common objection against forecasting competitions argues that aggregated performances of competitors cannot be discriminated significantly, due to cross-sectional noise. Given the above double-outcome, we may suggest a simple rank-test, instead.

In the context of multivariate filtering, we proposed a new indicator, the USRI¹⁵, for dating recessions in the US economy in real-time. The design of the indicator relies on a set of five monthly time series, suggested by the National Bureau of Economic Research, NBER. The USRI is on-line since March 2009, see Wildi[2009]. Besides the customized optimization approach, an outstanding particularity of the USRI resides in the publication strategy: the indicator is not revised. Observations prior to 2009 rely on a simulated real-time framework: one-sided filters are applied to data-vintages obtained from the ALFRED-project¹⁶. The singular WYSIWYG-publication strategy addresses users—for example fund managers—who want to calibrate models on the history of the time series. It is important to emphasize, in this context, that revisions of most published indicators preclude a coherent calibration of model parameters: in particular the overwhelming historical evidence conveyed by smoothed indicator performances would erroneously suggest an overly optimistic picture to model-builders. The USRI detected troughs of acceleration-, classical- and growth-cycles in the US as early as December 2008, April 2009 and June 2009 in real-time¹⁷. A longer history, relying on simulated real-time analysis, can be seen on—and downloaded from—the USRI-site.

In addition to these two illustrations, further successful experiences were accumulated over the years in various application fields including business-cycle monitoring (leading indicator for Swiss GDP), output-gap estimation (see Elmer/Wildi[2008] and Sturm/Wildi[2008]), financial trading and health-care forecast (regionally disaggregated multi-step ahead forecasting model for the Swiss National Office of Health-Care).

¹⁵ <http://www.idp.zhaw.ch/usri>.

¹⁶ <http://research.stlouisfed.org/tips/alfred/>.

¹⁷ We here refer to the dating as obtained by the multivariate filter. The adjoined Markov-switching filters relayed these signals by corresponding sharp probability-shifts with a slight delay of approximately one month.

6. CONCLUSION

The DFA unfolds the traditional one-step ahead perspective to more complex forecasting problems, including multi-step ahead projections or RTSE. The shift of perspective conveyed by this approach contrasts with homogenization tendencies enforced by abstract, user- and problem-detached, estimation principles. In particular, the pertinence of the traditional model-based approach in RTSE was challenged by the fact that the delicate identification of the DGP could be circumvented, in principle. In this perspective, the aim of model-based approaches turns out to be an unnecessarily ambitious task which detracts from the true problem. The DFA refocuses the statistical apparatus by emphasizing a fundamental uncertainty principle which reflects important user-priorities. Despite fundamentally different design priorities a connection between DFA and model-based approaches could be established by postulating a particular signal definition corresponding to the one-step ahead forward operator.

An extension to non-linear filtering could be entertained, of course. It is our firm subjective conviction, however, that marginal benefits obtained by exhausting the linear approach are likely to dominate potential gains by non-linear designs—in our particular applications—for some time.

REFERENCIAS BIBLIOGRÁFICAS

- BELL, W. (1984): Signal Extraction for Nonstationary Time Series. *Annals of Statistics*: 12.
- LARRY BRETHORST, G. (1988): Bayesian Spectrum Analysis and Parameter Estimation. *Lectures Notes in Statistics*, 48. Springer.
- CHIN, D.; GEWEKE, J. and MILLER, P. (2000): Predicting turning points. *Research Department Staff Report 267, Federal Reserve Bank of Minneapolis*.
- CLEVELAND, W.P. (1972): Analysis and Forecasting of Seasonal Time series. *PHD-Thesis, University of Wisconsin-Madison*.
- CLEVELAND, W.P. and TIAO, G.C. (1976): Decomposition of seasonal time series. *Journal of the American Statistical Association*, (71).
- DAGUM, E.B. (1980): The X-11-ARIMA Seasonal Adjustment Method. *Statistics Canada, Catalogue 12-564E*.
- GEWEKE, J. (1978): Revisions in seasonality adjusted time series. *SSRI report 78, University of Wisconsin, Madison*.
- HAMILTON, J.D. (2010): Calling Recessions in Real Time. *International Journal of Forecasting: Editor Invited Paper, International Symposium on Forecasting, San Diego*. <http://dss.ucsd.edu/~jhamilto/>
- SIMONE, E. and WILDI, M. (2008): Real-Time Filtering and Turning-Point Detection: Application of Customized Criteria (DFA) to the US Business Cycle. *IDP Working Paper IDP-WP-08Sep-04*. <http://www.idp.zhaw.ch/de/engineering/idp/forschung/finance-risk-management-and-econometrics/signal-extraction-and-forecasting/signal-extraction.html>.
- MCELROY, T. and WILDI, M. (2010.1): Signal Extraction Revision Variances as a Goodness-of-Fit Measure. *Journal of Time Series Econometrics: Vol. 2 : Iss. 1*.
- MCELROY, T. and WILDI, M. (2010.2): Multi-Step Ahead Estimation of Time Series Models (in preparation).

- STIER, W. and WILDI, M. (2002): About Model-Based Times Series Procedures: Some Remarks to TRAMO/SEATS and CENSUS X-12-ARIMA. *Advances in Statistical Analysis*, 86.
- STURM, J-E. and WILDI, M. (2008): Real-Time Filtering: Using the Multivariate DFA to Monitor the US Business Cycle. *IDP working paper*, IDP-WP-08Sep-04. <http://www.idp.zhaw.ch/de/engineering/idp/forschung/finance-risk-management-and-econometrics/signal-extraction-and-forecasting/signal-extraction.html>
- WILDI, M. (1998): Detection of Compatible Turning Points and Signal Extraction for Non-Stationary Time Series. *Operation Research Proceedings*, Springer.
- WILDI, M. (2004): Signal Extraction: How (In)efficient Are Model-Based Approaches? An Empirical Study Based on TRAMO/SEATS and Census X-12-ARIMA. *KOF-Working Paper Nr. 96*, ETH-Zurich.
- WILDI, M. (2005): Signal Extraction: Efficient Estimation, 'Unit-Root'-Tests and Early Detection of Turning-Points. *Lecture Notes in Economic and Mathematical Systems*, 547, Springer-Verlag Berlin Heidelberg.
- WILDI, M. (2008.1): Real-Time Signal-Extraction: Beyond Maximum Likelihood Principles. <http://www.idp.zhaw.ch/de/engineering/idp/forschung/finance-risk-management-and-econometrics/signal-extraction-and-forecasting/signal-extraction.html>.
- WILDI, M. (2008.2): Efficient Multivariate Real-Time Filtering and Cointegration. *IDP-working paper*, IDP-WP-08Sep-01. <http://www.idp.zhaw.ch/de/engineering/idp/forschung/finance-risk-management-and-econometrics/signal-extraction-and-forecasting/signal-extraction.html>.
- WILDI, M. (2009): Real-Time US-Recession Indicator (USRI): A Classical Cycle Perspective with Bounceback. *IDP-working paper*, IDP-WP-09Jun-06. <http://www.idp.zhaw.ch/en/engineering/idp-institute-of-data-analysis-and-process-design/research/finance-risk-management-and-econometrics/economic-indices/us-economic-recession-indicator.html>.
- WILDI, M. (2010): Multivariate Real-Time Signal Extraction: Fast and Reliable Turning-Point Detection. *IDP-working paper*, IDP-WP-10Feb-17.