Tracking economic activity in the euro area: multivariate direct filter approach*

Ginters Buss†

Abstract

The paper applies the multivariate direct filter approach on selected business and consumer confidence indicators and share price data to construct a real-time indicator tracking the medium-to-long-run component of the quarterly growth of the euro area gross domestic product. Results show that the created indicator behaves similarly to another established indicator - Eurocoin - but slightly leads it after mid 2009. The new indicator is also compared to the Markit euro area composite purchasing managers index and is found to be leading it by about one month and being smoother as well. Overall, the multivariate direct filter approach is found to have merit in tracking business cycle developments; however, the increasing number of free coefficients is an issue for the filter to be applied on rich datasets.

Keywords: real-time signal extraction, business cycles, multivariate time series
JEL code: C13, C32, E32

---

*The author thanks Marc Wildi and anonymous referees for valuable feedback on an earlier draft. All remaining errors are the author's own.

Disclaimer: This report is released to inform interested parties of research and to encourage discussion. The views expressed in this paper are those of the author and do not necessarily reflect the views of the Bank of Latvia.

†Address for correspondence: Latvijas Banka, K. Valdemara 2A, Riga, LV-1050, Latvia; e-mail: ginters.buss@gmail.com.

1 Introduction

The paper applies the recently developed multivariate direct filter approach (henceforth MDFA; see Wildi, 2011) on selected business and consumer confidence indicators and share price data to construct a real-time indicator tracking the medium-to-long-run component of the quarterly growth of the euro area gross domestic product (EA GDP).

The demand for real-time macroeconomic indicators exists mainly because many key macroeconomic variables are released with a considerable lag and are subsequently revised in later releases. For example, the first rough estimate of the main macroeconomic aggregate, GDP, called 'flash GDP', is released only 45 days after the reference period in the European Union (EU), including EA, and happens to be revised substantially. The first official release of GDP in EU is published only 65 days after the reference period, and even this first release is revised in subsequent releases. However, economists are keen to have timely information on the new developments in the economy. Therefore, efforts are made in using timely information to capture the big picture of the overall economy in a more timely manner. For that matter, several indicators are available that try to map timely disaggregate information on the aggregate series, like the GDP. Particularly, there are several established real-time indicators tracking economic situation in the EA, two of which are subjectively selected below.

The (New) Eurocoin (Altissimo et al., 2010) is a replacement of its predecessor, (Old) Eurocoin (Altissimo et al., 2003), since May 2009. This indicator tracks the medium-to-long-run component in the quarterly growth of the EA GDP and is the result of a dynamic factor methodology (Forni et al., 2000, 2005) applied to about 145 series.

The monthly Markit flash EA Purchasing Managers Index[1] (henceforth PMI) is one of the closely watched indices for the EA economy because it is timely, unrevised, and has straightforward methodology. It is based on surveys of companies.

This is the first paper known to the author that implements MDFA to track a trend-cycle, as well as being the first paper that implements this approach to the euro area dataset. Another paper that implements this approach is Wildi and Sturm (2008) but it uses a bandpass specification and applies it to the U.S. dataset. Wildi (2011) develops the filtration methodology implemented in this paper but does not apply it to the data. Recently, Gyomai and Wildi (2012) applied a refined version of the filter (see Wildi, 2012) to track the bandpass of the GDP of several countries but not the euro area.

This paper applies MDFA on selected survey and stock price data to create a real-time indicator for the EA GDP and compares its performance to the Eurocoin and the PMI. The filter output for the EA dataset is found to be about coincident with Eurocoin and to lead the PMI by about one month.

Overall, MDFA is found to have merit in tracking business cycle developments using a few explanatory variables; however, the increasing number of the free coefficients of the filter is an issue for the filter to be applied on rich datasets. See Wildi (2012) for the recent refinement of the methodology attempting to deal with this issue, and Buss (2012) for its application to many variables.

The paper is structured as follows. Section 2 reviews the methodology of the direct filter approach, Section 3 describes the design of the indicator, Section 4 produces the indicator and compares with the alternatives, Section 5 concludes, and the Appendix discusses robustness checks.

2 Background on the direct filter approach

The direct filter approach is concerned with estimating a signal, such as a trend-cycle or a business cycle, in real time. Consider weakly stationary zero-mean input series $\{x_t\}$, $t = 1, \dots, T$, with an absolutely continuous spectral distribution function $H(\omega) = \int_{-\pi}^{\omega} h(\omega')d\omega'$, where $\omega \in [-\pi, \pi]$ denotes the frequency, and $h(\omega)$ is the spectral density of the series, $h(\omega) = \frac{1}{2\pi}\sum_{\tau=-\infty}^{\infty} c(\tau)\exp(-ij\omega)$, where $c(\tau)$ is the autocovariance function, $c(\tau) = \lim_{T\to\infty}\frac{1}{T}\sum_{t=1}^{T-|\tau|} x_t x_{t+|\tau|}$, $|\tau| = 0, \dots, T-1$, $|\cdot|$ denotes the absolute value, and i is the imaginary number defined as $i^2 = -1$. Denote $\{y_t\}$

[1]See http://www.markiteconomics.com/Survey/Page.mvc/PressReleases, accessed on Nov 26, 2012

as the ideal output of a symmetric, possibly bi-infinite filter, $\Gamma(B)=\sum_{j=-\infty}^{\infty}\gamma_j B^j$, applied on an input series $\{x_t\}$:

$$y_t=\sum_{j=-\infty}^{\infty}\gamma_j x_{t-j}, \tag{1}$$

where B is the lag or backshift operator. The filter in (1) is called the ideal filter and the filter output, $\{y_t\}$, is called the ideal filter output. Time series samples have finite length in practice and, therefore, the ideal filter, in general, is infeasible to apply. A practitioner might use a finite length symmetric filter as an approximation to (1) in the middle of the time series but even the symmetric approximation is infeasible to apply at the very end of the sample, i.e., in real time. A real-time estimate of $\{y_T\}$, given a finite input $\{x_1,\ldots,x_T\}$, can be written as

$$\hat{y}_T=\sum_{j=0}^{T-1} b_j x_{T-j}. \tag{2}$$

It is a well-known fact that the filter in (2) generally possesses a nontrivial time shift (defined as the phase (defined below) divided by the frequency), i.e., it's output is lagging in time. In order to define the phase and, in general, the effect of a filter, denote the generally complex transfer functions of filters in (1) and (2) by $\Gamma(\omega)=\sum_{j=-\infty}^{\infty}\gamma_j\exp(-ij\omega)$ and $\hat{\Gamma}(\omega)=\sum_{j=0}^{T-1}b_j\exp(-ij\omega)$, respectively. The number $\Gamma(\omega)$ can be represented in polar coordinates as $\Gamma(\omega)=A(\omega)\exp(-i\Phi(\omega))$, where $A(\omega)=|\Gamma(\omega)|$ is the amplitude, and $\Phi(\omega)=-\arg(\Gamma(\omega))$ is the phase. The effect of a filter - as represented by its transfer function - can be summarized by the effects of the amplitude and phase functions. The amplitude function, $A(\omega)$, can be interpreted as the weight - amplification or damping - attributed by the filter to the input series at frequency ω. The phase function, $\Phi(\omega)$, can be interpreted as a shift function of the amplified or damped signal at frequency ω. This paper is concerned with obtaining a one-sided filter output which would be a good - as good as it can be in real time - approximation to the ideal output. For a weakly stationary input series $\{x_T\}$, the mean squared filter error (MSFE) can be expressed as the mean squared difference between the ideal output and the real-time estimate:

$$\int_{-\pi}^{\pi}|\Gamma(\omega)-\hat{\Gamma}(\omega)|^2 dH(\omega)=E[(y_T-\hat{y}_T)^2]. \tag{3}$$

A finite sample approximation of MSFE, (3), is

$$\frac{2\pi}{T}\sum_{k=-[T/2]}^{[T/2]}\xi_k|\Gamma(\omega_k)-\hat{\Gamma}(\omega_k)|^2 S(\omega_k), \tag{4}$$

where $\omega_k=k2\pi/T$, $[T/2]$ is the greatest integer smaller or equal to $T/2$, and the term ξ_k is defined as

$$\xi_k=\begin{cases}1 & \text{for } |k|\neq T/2\\ 1/2 & \text{otherwise,}\end{cases} \tag{5}$$

(see Brockwell and Davis, 1987, Ch. 10 for the reason for ξ_k; without it the inverse discrete Fourier transform does not replicate the sample autocovariance perfectly). $S(\omega_k)$ in (4) is defined to be an estimate of the unknown spectral density of $\{x_t\}$, which can be any spectral estimate. As discussed in Wildi (2008), consistency of $S(\omega_k)$ is not required because the goal is not to estimate $dH(\omega)$ but the filter mean squared error, (3). Therefore, this paper uses the periodogram, $I_{Tx}(\omega_k)$ - which is sufficient for our purposes, and bears an analogy to a sufficient statistic - defined as:

$$S(\omega_k):=I_{Tx}(\omega_k)=\frac{1}{2\pi T}\left|\sum_{t=1}^{T}x_t\exp(-it\omega_k)\right|^2. \tag{6}$$

Minimizing expression (4) yields the real-time filter output optimally approximated to the ideal output in mean squared error sense. The expression (4) is a generalized problem that encompasses the problem of Baxter and King (1999) (where the feasible filter is assumed to be a symmetric bandpass and the spectral estimate, $S(\omega_k)$ be the spectrum of the white noise), Christiano and Fitzgerald (2003) (where the default specification of the feasible filter is assumed to be a bandpass and the spectral estimate, $S(\omega_k)$ be the spectrum of the random walk), or the Hodrick-Prescott filter (which can be interpreted to be a Butterworth-type filter, assuming the input following an ARIMA(0,2,2) process, and its amplitude function being of a particular truncated bell-shaped form, see King and Rebelo (1993) and Maravall and Rio (2001)). Yet, Wildi (2008) proposes a generalized version of (4) which this paper finds to be useful in creating a real-time indicator, and which is described in the next subsection.

2.1 Univariate direct filter approach

Wildi (2008) proposes a customized version of (4), a part of which is implemented in this paper. Rewrite the discrete version of MSFE, (4), as

$$\frac{2\pi}{T}\sum_{k=-[T/2]}^{[T/2]}\xi_k|\Gamma(\omega_k)-\hat{\Gamma}(\omega_k)|^2 I_{Tx}(\omega_k)W(\omega_k), \tag{7}$$

which is identical to (4) for $W(\omega_k) := 1$. However, a more general version of $W(\omega_k) := W(\omega_k, expw, cut)$ can be written as

$$W(\omega_k, expw, cut) = \begin{cases} 1 & \text{if } |\omega_k| < cut \\ (1+|\omega_k|-cut)^{expw} & \text{otherwise,} \end{cases} \tag{8}$$

which collapses to unity for $expw = 0$, in which case the classical mean squared optimization, (4), is obtained. Parameter cut (for 'cut-off frequency') marks the transition between passband and rightmost stopband, and positive values of $expw$ (for 'exponent weight') emphasize high-frequency components in the rightmost stopband, thus, making the filter output smoother than the one obtained by minimizing (4) for positive $expw$. By choosing $expw > 0$, the additional weighting function $W(\omega_k)$ grows strictly monotonically in the stopband. Thus, it assigns more weight to the fit of $\Gamma(\omega_k)$ by $\hat{\Gamma}(\omega_k)$ towards higher frequencies. Since the ideal trend filter vanishes in the stopband, the additional weighting scheme forces $\hat{\Gamma}(\omega_k)$ to approach $\Gamma(\omega_k)$, which vanishes in the stopband. As a result, high-frequency components are damped more vigorously by the resulting real-time filter: the output, in general, will be smoother.

Univariate filters can be useful in case the target variable of interest, x_t, is timely; however, it is usually not the case for macroeconomic series. For example, one of the key macroeconomic variables, GDP, is released with a delay and is revised in subsequent releases. Therefore, a practitioner might be interested in using other economic variables with shorter delays in creating benchmark indicators for GDP. Therefore, the next subsection turn to the multivariate filter.

2.2 Multivariate direct filter approach

The above univariate customized filter has been generalized to a multivariate filter in Wildi (2011). Rewrite the univariate minimization problem, (7), with the discrete Fourier transform (DFT), $\Xi_{Tx}(\omega_k)$:

$$\begin{aligned} &\frac{2\pi}{T}\sum_{k=-[T/2]}^{[T/2]}\xi_k|\Gamma(\omega_k)-\hat{\Gamma}(\omega_k)|^2 I_{Tx}(\omega_k)W(\omega_k) \\ &= \frac{2\pi}{T}\sum_{k=-[T/2]}^{[T/2]}\xi_k|\Gamma(\omega_k)\Xi_{Tx}(\omega_k)-\hat{\Gamma}(\omega_k)\Xi_{Tx}(\omega_k)|^2 W(\omega_k), \end{aligned} \tag{9}$$

where

$$\Xi_{Tx}(\omega_k) = \sqrt{\frac{1}{2\pi T}}\sum_{t=1}^{T} x_t \exp(-it\omega_k). \tag{10}$$

Assume there are m additional explanatory variables z_n, $n = 1, \ldots, m$ that might help improve the real-time estimate of $\{y_T\}$ obtained with a univariate filter. Then, the second expression in the modulus on the second line of (9), $\hat{\Gamma}_X(\omega_k)\Xi_{Tx}(\omega_k)$, is replaced by

$$\hat{\Gamma}_X(\omega_k)\Xi_{Tx}(\omega_k) + \sum_{n=1}^{m} \hat{\Gamma}_{z_n}(\omega_k)\Xi_{Tz_n}(\omega_k), \tag{11}$$

where

$$\hat{\Gamma}_X(\omega_k) = \left(\sum_{j=0}^{L} b_{xj} \exp(-ij\omega_k) \right) \tag{12}$$

$$\hat{\Gamma}_{z_n}(\omega_k) = \left(\sum_{j=0}^{L} b_{z_n j} \exp(-ij\omega_k) \right) \tag{13}$$

are the one-sided transfer functions applied to the explanatory variables, $\Xi_{Tx}(\omega_k)$, $\Xi_{Tz_n}(\omega_k)$ are the corresponding DFTs, and $L \in \{1, 2, \ldots, T-1\}$ is the filter length. Then, the multivariate version of (9) can be written as

$$\frac{2\pi}{T} \sum_{k=-[T/2]}^{[T/2]} \xi_k \left| \left(\Gamma(\omega_k) - \hat{\Gamma}_x(\omega_k) \right) \Xi_{Tx}(\omega_k) - \sum_{n=1}^{m} \hat{\Gamma}_{z_n}(\omega_k)\Xi_{Tz_n}(\omega_k) \right|^2 W(\omega_k). \tag{14}$$

The paper uses the filter obtained by minimizing (14) (without $\hat{\Gamma}_x(\omega_k)$, since the lagged target series is not included in the set of the explanatory variables, see discussion below) with respect to filter parameters and subject to the amplitude constraint at frequency zero:

$$\sum_{j=0}^{L} b_{z_n j} = 1/m \tag{15}$$

that imposes a constraint at frequency zero according to $\hat{\Gamma}_{z_n}(0) = 1/m$ for all $n = 1, \ldots, m$, such that the sum of the amplitudes over the cross-sectional dimension of the filter, m, yields unity. Such a constraint ensures that the scale of the filter output matches the scale of the target signal. See Wildi (2011) for further details.

3 Indicator design

3.1 The target

The target is the medium-to-long-run fluctuations in the quarterly growth of the seasonally adjusted (SA) EA GDP in chain-linked prices. Specifically, the filter is set to target the ideal lowpass of the quarterly growth of SA EA GDP with the cut-off wave length of 12 months (see Figure 1).

The usually defined minimum length of a business cycle is 18 months (Burns and Mitchell, 1946; Baxter and King, 1999; Christiano and Fitzgerald, 2003). However, passing higher-frequency content through the filter allows for both faster turning point detection and a closer scale fit. The quarterly GDP data are taken from 1995Q1 till 2011Q4. The GDP series is linearly interpolated to monthly frequency (by drawing a straight line between the two nearest observed monthly values in levels), logged, regularly differenced and demeaned before its spectral content enters the filter. Alternatively, the interpolation and differencing can be performed in frequency domain but the results show no obvious improvement with the latter compared to such data transformation in time domain. Linear interpolation of quarterly GDP has little or no effect on the real-time performances of the filter because the interpolation affects the monthly frequencies, but the target of the filter (trend-cycle with a cut-off length of one year) eliminates the information on the monthly frequencies.

Figure 1: Amplitude of the ideal lowpass filter

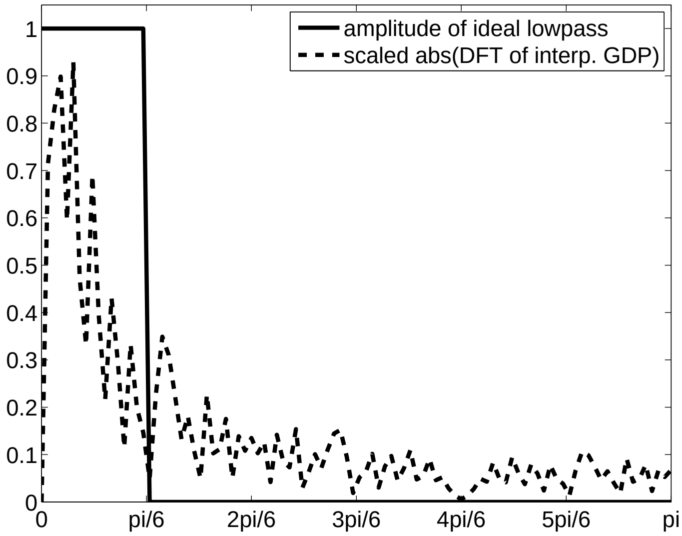

Note: The scaled absolute value of DFT of differenced interpolated GDP is zero in frequency zero because the GDP series is demeaned.

3.2 Explanatory variables

Selected monthly confidence indicators published by the Directorate-General for Economic and Financial Affairs of the European Commission (DG ECFIN) and the United States (US) share price index published by the Eurostat[2] are used as explanatory variables for the EA GDP. DG ECFIN data are usually published at the end of the reference month, except for December for which data are published in early January. DG ECFIN business and consumer surveys data are almost unrevised - this applies both to the seasonally unadjusted and the seasonally adjusted data, as the latter is the product of the seasonal adjustment program 'Dainties' that does not revise history as new data come in[3]. The above-mentioned considerations make DG ECFIN data suitable for real-time filtration. The selected business and consumers confidence data are: production trend observed in recent months (industry), assessment of order-book levels (industry), assessment of stocks of finished products (industry), production expectations for the months ahead (industry), employment expectations for the months ahead (industry), confidence indicator in construction, confidence indicator in retail trade, and consumers confidence indicator.

Other survey data have been tested and found not to add substantial quality to the indicator. SA DG ECFIN data are used (see Figure 2).

Monthly US share price index as published by Eurostat is used as an additional explanatory variable due to the potential economic spillovers from US. It is published about one week after the reference month. Since the real-time indicator - as presented in this paper - is set to be published at the end of reference month, the US share price index enters the filter with a lag of one month. Thus, there is a room for potential improvement in this aspect as financial data generally are available on the go.

The reason for limiting the set of explanatory variables to the above is the restrictive nature of the filter to deal with richer datasets due to the increasing number of free coefficients of the filter that worsens its out-of-sample performance, see discussion below.

Although the GDP series itself could be filtered, it is chosen not to do so as the GDP series is substantially lagged, happens to be revised substantially, and appears not to be particularly useful for the concurrent signal extraction. Therefore, only the spectral content of the GDP series is fed

[2] http://epp.eurostat.ec.europa.eu/portal/page/portal/eurostat/home, accessed on Nov 26, 2012

[3] For details, see 'The joint harmonized EU programme of business and consumer surveys', User Guide, 2007, European Commission Directorate-General for Economic and Financial Affairs, available at http://ec.europa.eu/economy_finance/db_indicators/surveys/documents/userguide_en.pdf, accessed on Nov 26, 2012

to the filter and the monthly variables are set to target that spectral content.

Figure 2: DG ECFIN survey data

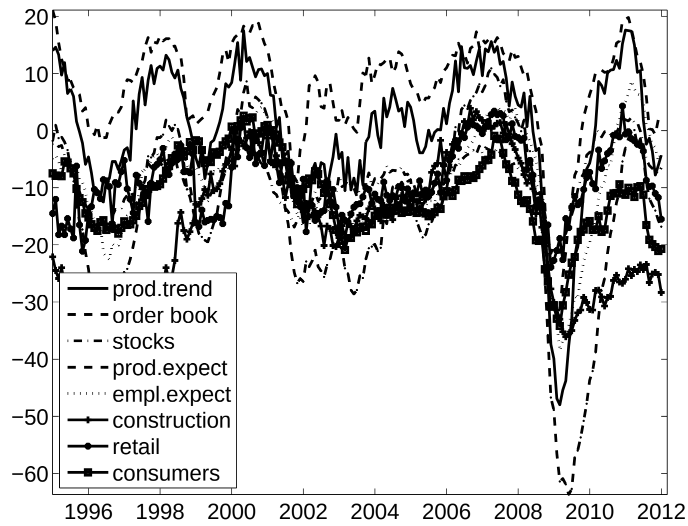

All explanatory variables are taken from January 1995 till January 2012, regularly differenced and standardized to zero mean and unit variance. The list of data and their transformations is shown in Table 1.

Table 1: The list of data used for the construction of the indicator

Variable	Source	Transformation
Gross domestic product, chain-linked, SA	Eurostat	Δ log,lin.interp.
Production trend observed in recent month (industry), SA	DG ECFIN	Δ
Assessment of order-book levels (industry), SA	DG ECFIN	Δ
Assessment of stocks of finished products (industry), SA	DG ECFIN	Δ
Employment expectations for the months ahead (industry), SA	DG ECFIN	Δ
Production expectations for the months ahead (industry), SA	DG ECFIN	Δ
Confidence indicator in construction, SA	DG ECFIN	Δ
Confidence indicator in retail trade, SA	DG ECFIN	Δ
Consumers confidence indicator, SA	DG ECFIN	Δ
The US share price index	Eurostat	Δ log

Notes: 1) Δ denotes the first difference operator, i.e. $\Delta x_t = x_t - x_{t-1}$; 2) lin.interp. denotes 'linear interpolation' defined in the main text.

3.3 Filter dimension

The filter dimension is defined by its length, L, and its cross-sectional dimension, i.e. the number of explanatory variables, m. A higher dimension of the filter leads to a higher number of free coefficients of the filter, which may lead to worse out-of-sample performance. Thus, the filter is tested for various dimension sizes. Unreported results of the author show that, given the available sample length, the cross-sectional dimension of one to four variables show comparable in-sample and out-of-sample performances but using more than four variables manifests the signs of overfitting, i.e. the out-of-sample performance seemingly deteriorates compared to the in-sample results. Therefore, the cross-sectional dimension of the filter is fixed to two variables.

For the same reason, the filter length is set to depend on the number of input variables (negatively) and the sample length (positively). The particular two-variable filter is set to be 38 to 48 observations long, depending on the in-sample length.

3.4 Noise suppression

Extra suppression of high-frequency content in the stopband is implemented with a positive *expw* parameter. Particularly, $expw = 1/2$. As a result, the filter output is smoother than the one of the classical mean squared filter error problem, see discussion in section A.1.

3.5 Cross-sectional aggregation

Given that the considered set of explanatory variables contains more variables than the selected filter cross-sectional dimension, all possible combinations of variables are entered in the filter, such that the final filter output is combined 2-variable filter outputs weighted by weights proportional to inverse estimated mean squared filter error (EMSFE). Alternatively, one could use equal weighting but the resulting output might be more sensitive to the inclusion of irrelevant variables. As another alternative, principal components (Stock and Watson, 2002) could be used to shrink the dimension of the dataset to a few factors but unreported results of the author show no evident improvement in this regard. One explanation for this result might be the small size of the data set. However, the author suspects that there is more than that: the spectral content of the unobserved principal components might be subject to variation over different subsamples which might lead to inferior real-time performance.

The disadvantage of the cross-sectional aggregation procedure is that it does not take into account the full cross-sectional information available. The advantage of this procedure is that the errors of individual filters are partially canceled out. This procedure resembles that of model combination in the time series forecasting literature.

3.6 Output re-scaling and adding the mean back

Filter output obtained thus far is then rescaled to the variance of the output of 31-months long finite symmetric filter for the time span available at the particular real-time estimation moment. Finally, the mean of the output of the finite symmetric filter is added to the filter output. One could use the mean of the GDP series, instead, but since the latter is more volatile than the output of the finite symmetric filter, the latter is preferred.

3.7 Averaging along time dimension

At this point, one can produce the real-time indicator. However, a couple of issues emerge. First, the GDP series happens to be revised. Thus, some of its observations at the end of series - call them 'gap' - are left out of sample before its spectral content is entered into the filter. Denote the whole sample length by 'len'. Then, the maximum in-sample period spans observations $1 : len - gap$. Setting *gap* to nine or 15 months does not make much difference but insures against GDP revisions at the end of series and ensures that the target is about the 'final' GDP as opposed to the 'first-release'. The presented indicator below is the result of setting 'gap' to nine months.

Second, if filter coefficients are updated every time new data become available but the historic values of the indicator are not updated, then the noise generated by the changing estimated latent level of the target signal can suppress the estimated signal. Figure 3 shows an indicator resulting from such filter coefficient re-estimation every quarter. Although the filter output is timely [The peak correlation (0.872) between the filter output and the GDP series is found to be at a one month lag with respect to (w.r.t.) GDP, and the second highest correlation (0.862) located at a zero months lag w.r.t. GDP (see Table 2)], it is not smooth.

In this light, one might choose an in-sample span yielding the most satisfactory out-of-sample performance and fix coefficients. In this case, the real-time estimate is smooth. A real-time recession indicator for the US economy, USRI[4], is based on the latter approach. Yet, another issue emerges - sooner or later the filter coefficients might need an update due to the possible structural changes in the economy.

[4]See http://www.idp.zhaw.ch/usri, accessed on Nov 26, 2012

Figure 3: Filter output resulting from frequent filter coefficient update

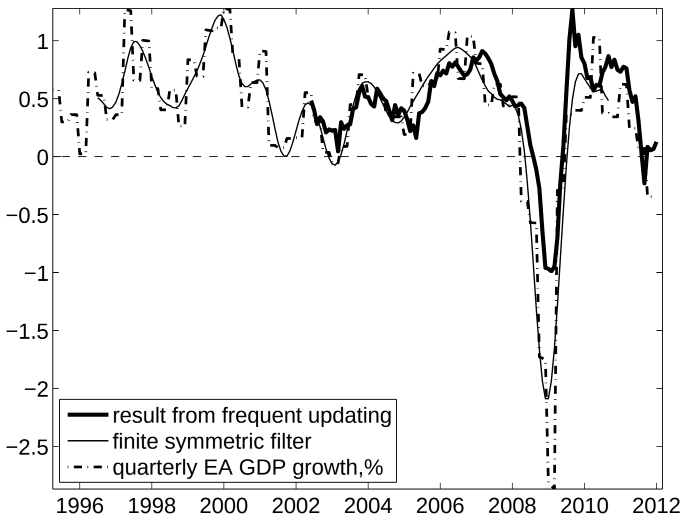

Note: The symmetric filter is 31 months long.

The above issues are solved with the following *ad hoc* filter averaging procedure along the time dimension. Update the filter coefficients every time the new GDP data become available. The final indicator is the average of all filter outputs created up to the estimation period. Thus, iteratively estimate filter coefficients for reduced in-sample length, $1 : len - gap - 3$, $1 : len - gap - 6$, etc. until some minimum in-sample length, but estimate filter outputs for the whole sample span, 'len'. Then, take the arithmetic average[5] of all historically estimated indicators. More indicators are averaged as more data come in but the maximum aggregation time span is set to five years in order to take into account possible structural changes. This updating mechanism overcomes the problem induced by the level of target series and at the same time robustifies signal detection, and incorporates new information as new data become available. As the next section shows, this *ad hoc* procedure does not change the timeliness or the scale of the filter output, compared to the one in Figure 3, considerably, but reduces the short-term fluctuations.

4 Results

The above mechanism is performed for the last ten years. Out-of-sample time span is limited by the filter and sample lengths. The resulting simulated real-time out-of-sample filter output is shown in Figure 4, along with the quarterly growth of the EA GDP and the output of a 31-months long finite symmetric trend-cycle filter.

Figure 4 shows that the new indicator is slightly lagging in a historic perspective, and that its scope is comparable to that of the output of the finite symmetric filter except for the great recession episode. A visual inspection suggests that the indicator gets slightly faster after the great recession episode, when it appears to be coincident with the GDP growth. Methodologically, the indicator is constructed as a coincident indicator. Furthermore, the indicator gets slightly smoother over time. The improvement in the smoothness over time could be explained by the increasing data length available to the filter. Moreover, the smoothness performance is partly due to the filter averaging along the time dimension - the first observation of the indicator in early 2002 is the result of a single cross-sectionally aggregated filter output; the next indicator values are averages of increasing number of filter outputs along the time dimension until five years pass, when only the filter outputs over the last five years are averaged to account for the effect of different phases of a business cycle, as well as possible structural changes.

[5] A weighted average proportional to inverse EMSFE also was tested and found not to yield superior results.

Figure 4: Filter output

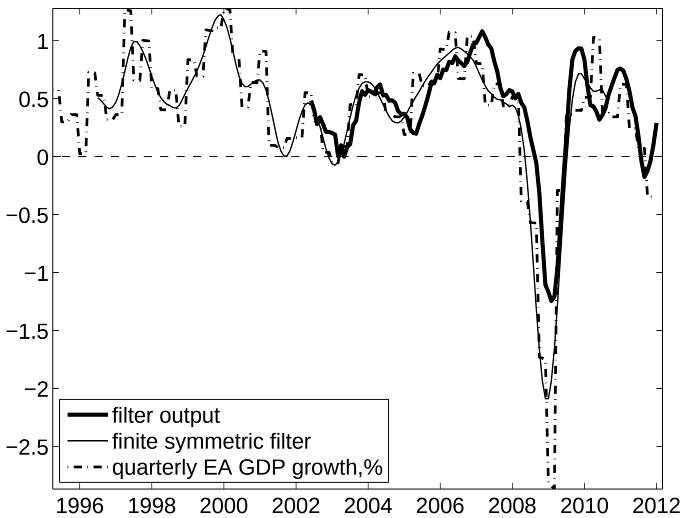

Note: the symmetric filter is 31 months long.

Quantitatively, the filter output cannot be compared to the ideal filter output since the latter is unavailable for finite samples, and its finite approximation might be imprecise for such a short sample span. Therefore, this paper calculates dynamic correlations between the filter output and the GDP series. The dynamic correlations capture well both the scale fit and the time shift. The peak correlation (0.877) between the filter output in Figure 4 and the GDP series is found to be at a one month lag w.r.t. GDP, and the second highest correlation (0.850) located at a zero months lag w.r.t. GDP (see Table 2).

4.1 Comparison with Eurocoin and Markit EA PMI

4.1.1 Eurocoin

Eurocoin (Altissimo et al., 2010) is an established coincident indicator for EA GDP. According to its website[6], it is a smooth real time estimate of the quarterly growth of EA GDP, released at the end of reference month, and is not revised since May 2009. It is the result of the generalized dynamic factor model (Forni et al, 2000, 2005; Forni and Lippi, 2001) applied to - according to Altissimo et al., 2010 - about 145 series. In contrast, the new indicator is a result of a multivariate filter applied on only selected nine explanatory variables. Eurocoin and the new indicator have a few features in common, as well - both target a trend-cycle in the quarterly growth of EA GDP as defined by a lowpass with the cut-off frequency of one year, and both are designed to be coincident indicators. As such, Eurocoin is a decent benchmark for the new indicator.

Figure 5 shows Eurocoin, latest vintage of the quarterly growth of EA GDP, and the filter output for the last ten years.

Figure 5 shows that the scope of both indicators are similar over the considered period. The new indicator is about coincident with Eurocoin during 2002-2004 but less smooth due to the small in-sample period. The new indicator appears to be faster than Eurocoin during 2005, and coincident but with worse level fit during 2007-2008. Eurocoin gets behind slightly after 2009. Given that the true out of sample period for Eurocoin is only since May 2009, and that its pseudo real-time values are calculated on the last vintage data for the period before May 2009[7], it might indicate that the real-time performance of Eurocoin is slightly worse than its pseudo real-time estimates suggest, and that the new indicator slightly leads Eurocoin during the true real time

[6]See http://eurocoin.bancaditalia.it/, accessed on Nov 26, 2012

[7]See the note in http://eurocoin.cepr.org/files/file/Ecoin_realtime_99Feb12.xls, accessed on Nov 26, 2012. An excerpt from the linked text: "Technically we should call this a 'pseudo-realtime exercise' since we use the 'final' data in the estimation and not the first releases since they are not available for all data series."

Figure 5: Comparison to Eurocoin

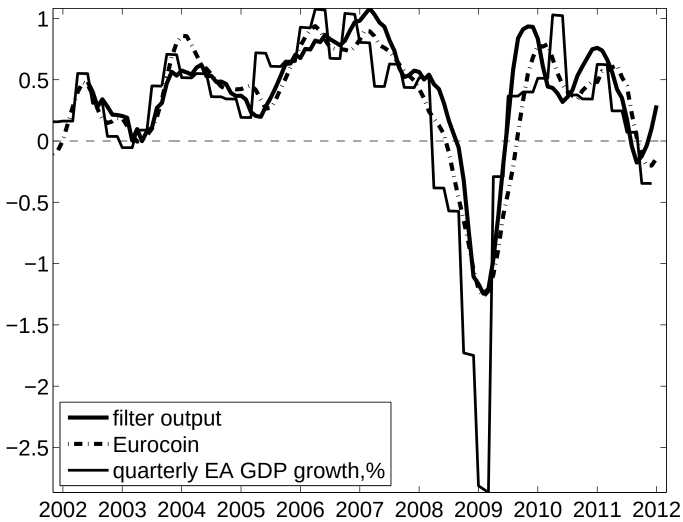

performance of both indicators. Quantitatively, the peak correlation (0.885) between the Eurocoin and the GDP series is found to be at a one month lag with respect to (w.r.t.) GDP, and the second highest correlation (0.880) located at a two months lag w.r.t. GDP (see Table 2). Thus, the new indicator leads the Eurocoin on average by zero to one month.

4.1.2 Markit EA PMI

Markit EA PMI is revealed a few days after the end of reference month and is advertised as being closely correlated with the quarterly growth of EA GDP. PMI is one of the closely watched indices by economists due to its early release, simple design and economic relevance. The PMI data for the last five years are collected from Bloomberg[8] and plotted against the new indicator. Since PMI is of different magnitude than the GDP growth, all variables are normalized to zero mean and unit variance for an easier comparison. Figure 6 shows that the new indicator is about coincident with, but is smoother than PMI. Quantitatively, the peak correlation (0.891) between PMI and GDP is found to be at a two months lag w.r.t. GDP, and the second highest correlation (0.880) located at a one month lag w.r.t. GDP (see Table 2). Thus, the new indicator leads the PMI on average by about one month.

Table 2: Dynamic correlations of indicators with GDP growth rates

Indicators	Dynamic correlation at lag:		
	-2	-1	0
Eurocoin	0.880*	0.885**	0.853
Markit PMI	0.891**	0.880*	0.824
MDFA, expw=0	0.824	0.882**	0.874*
MDFA, expw=0.5	0.841	0.877**	0.850*
MDFA, expw=1	0.846*	0.856**	0.811
MDFA in Fig. 3	0.804	0.872**	0.862*

Notes: 1) The correlations are calculated from the longest common sample such that all the reported numbers are comparable. The sample of the GDP series spans from May 2007 till November 2011.
2) ** marks the peak correlation; * marks the second highest correlation.

[8] http://www.bloomberg.com/, accessed on Nov 26, 2012

Figure 6: Comparison to Markit EA PMI

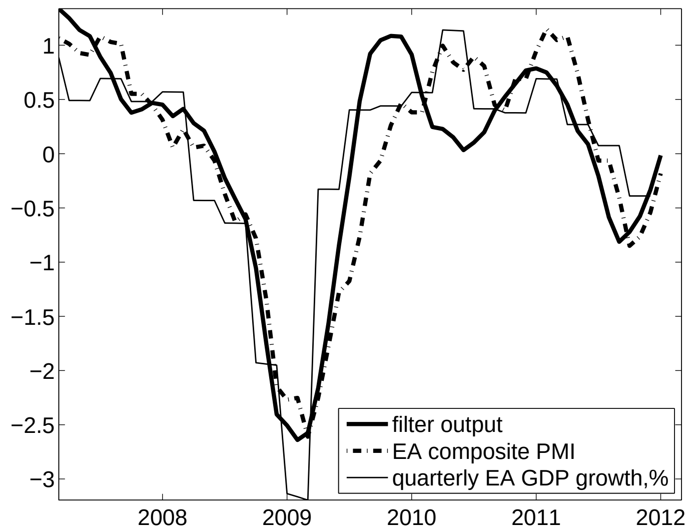

Notes: 1) All series are normalized to zero mean and unit variance.
2) The apparent dip in the filter output right after 2010 is due to the rescaling; see its level in Figure 5.

5 Conclusions

The paper applies the multivariate direct filter approach on selected business and consumer confidence indicators and share price data to construct a real-time indicator tracking the medium-to-long-run component of the quarterly growth of EA GDP.

The results show that the created indicator behaves similarly to another established indicator - Eurocoin - but leads it slightly after mid 2009. The new indicator is also compared to the Markit euro area composite purchasing managers index and is found to be leading it on average by about one month and being smoother as well.

The main conclusion of the paper is that the multivariate direct filter approach has its merits but it is not suitable to process rich datasets due to the increasing number of free coefficients and thus the deteriorating out-of-sample performance. In this regard, either the dimension of the dataset should be reduced prior to applying the filter, or the regularized version of the filter (Wildi, 2012) should be considered.

References

[1] Altissimo, F., Cristadoro, R., Forni, M., Lippi, M. & Veronese, G. 2010. New Eurocoin: Tracking Economic Growth in Real Time. *The Review of Economics and Statistics*, 92(4): 1024-1034. doi:10.1162/REST_a_00045

[2] Altissimo, F., Bassanetti, A., Cristadoro, R., Forni, M., Hallin, M., Lippi, M., Reichlin, L. & Veronese, G. 2003. Eurocoin: a real time coincident indicator for the euro area business cycle. *Computing in Economics and Finance 2003*, 242, Society for Computational Economics.

[3] Arnaud, B., & Hong, E.P. 2001. Comparison of compilation methodologies for the composite leading indicators of Euro area. OECD.

[4] Baxter, M. & King, R.G. 1999. Measuring business cycles: approximate band-pass filters for economic time series. *The Review of Economics and Statistics*, 81(4): 575-593. doi: 10.1162/003465399558454

[5] Bollerslev, T. 1986. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3): 307-327. doi: 10.1016/0304-4076(86)90063-1

[6] Bollerslev, T. & Wooldridge, J.M. 1992. Quasi-maximum likelihood estimation and inference in dynamic models with time varying covariances. *Econometric Reviews*, 11(2): 143-172.

[7] Brockwell, P.J. & Davis, R.A. 1987. *Time Series: Theory and Methods*. New York: Springer Verlag.

[8] Burns, A.F. & Mitchell, W.C. 1946. *Measuring business cycles*. New York: NBER.

[9] Buss, G. 2012. Forecasting and signal extraction with regularised multivariate direct filter approach. Working Papers 2012/06, Latvijas Banka.

[10] Caporello, G., Maravall, A. & Sanchez, F.J. 2001. Program TSW reference manual. Banco de Espana Working Papers 0112, Banco de Espana.

[11] Christiano, L.J & Fitzgerald, T.J. 2003. The band pass filter. *International Economic Review*, 44(2): 435-465. doi: 10.1111/1468-2354.t01-1-00076

[12] Forni, M., Hallin, M., Lippi, M. & Reichlin, L. 2000. The generalized dynamic factor model: identification and estimation. *The Review of Economics and Statistics*, 82(4): 540-554. doi: 10.1162/003465300559037

[13] Forni, M., Hallin, M., Lippi, M. & Reichlin, L. 2005. The Generalized Dynamic Factor Model: One-Sided Estimation and Forecasting, *Journal of the American Statistical Association*, vol. 100, 830-840.

[14] Forni, M. & Lippi, M. 2011. The generalized dynamic factor model: representation theory. *Econometric Theory*, 17(06): 1113-1141.

[15] Gyomai, G. & Wildi, M. 2012. OECD Composite Leading Indicators for G7 Countries: A Comparison of the Hodrick-Prescott Filter and the Multivariate Direct Filter Approach. OECD Statistics Working Papers 2012/5, OECD Publishing.

[16] Hodrick, R. & Prescott, E.C. 1997. Postwar U.S. business cycles: an empirical investigation. *Journal of Money, Credit, and Banking*, 29(1): 1-16.

[17] Jarque, C.M. & Bera, A.K. 1987. A test for normality of observations and regression residuals. *International Statistical Review*, 55(2): 163-173.

[18] King, R.G. & Rebelo, S.T. 1993. Low frequency filtering and real business cycles. *Journal of Economic Dynamics and Control*, 17(1-2): 207-231. doi: 10.1016/S0165-1889(06)80010-2

[19] Kolmogorov, A.N. 1941. Stationary sequences in Hilbert space. (in Russian), Bull. Moscow Univ., 2(6): 1-40.

[20] Massey, F.J. 1951. The Kolmogorov-Smirnov test for goodness of fit. *Journal of American Statistical Association*, 46(253): 68-78.

[21] Maravall, A. & Rio, A. 2001. Time aggregation and the Hodrick-Prescott filter. Banco de Espana Working Papers 0108, Banco de Espana.

[22] Stock, J.H. & Watson, M.W. 2002. Macroeconomic forecasting using diffusion indexes, *Journal of Business & Economic Statistics*, vol. 20(2), 147-162.

[23] Wiener, N. 1949. *Extrapolation, interpolation, and smoothing of stationary time series*, MIT Press.

[24] Wildi, M. 2008. *Real-Time Signal-Extraction: Beyond Maximum Likelihood Principles*, available at http://blog.zhaw.ch/idp/sefblog/uploads/Wildi_Real_Time_SE_0810010.pdf, accessed on Nov 26, 2012.

[25] Wildi, M. 2011. I-DFA and I-MDFA: Companion paper to R-code published on SEFBlog. IDP-Working Paper, available at http://blog.zhaw.ch/idp/sefblog/uploads/working_paper1.pdf, , accessed on Nov 26, 2012.

[26] Wildi, M. 2012. Elements of forecasting and signal extraction, IDP-Working Paper, available at http://blog.zhaw.ch/idp/sefblog/uploads/elements_3007.pdf, accessed on Nov 26, 2012.

[27] Wildi, M. & Sturm, J.E. 2008. Real-time filtering: using the multivariate DFA to monitor the US business cycle. IDP-Working Paper, available at http://www.idp.zhaw.ch/fileadmin/user_upload/engineering/_Institute_und_Zentren/IDP/forschungsschwerpunkte/FRME/sef/signalextraction/papers/IDP-WP-08Sep-02.pdf, accessed on Nov 26, 2012.

Appendix A Robustness checks

A.1 The effect of high-frequency noise suppression

Figure 7 shows the effect of the high-frequency noise suppression in the rightmost stopband by the customized filter.

Figure 7: The effect of high-frequency noise suppression

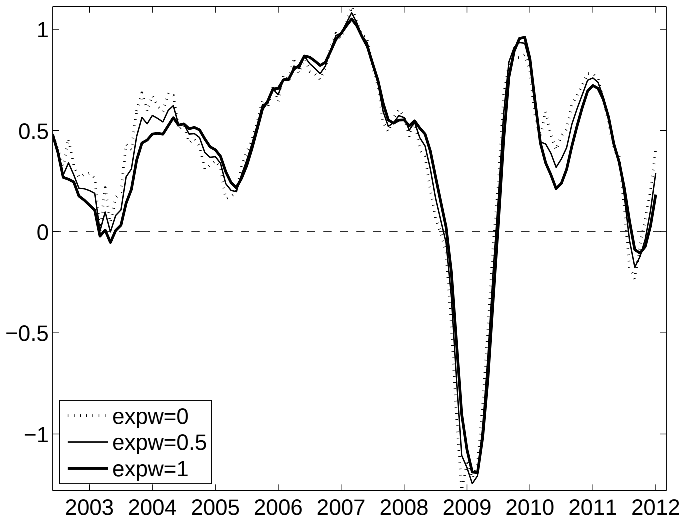

Note: The final indicator presented in this paper is the middle case with $expw = 0.5$.

The dotted line corresponds to the classical mean squared optimization problem with $W(\omega_k) = 1$, i.e., there is no extra noise suppression. The other two lines correspond to the customized filter output with $expw = 0.5$ and $expw = 1$. It can be seen that the indicator without the extra noise suppression is the fastest but also the noisiest of the three. The user of a real-time indicator might want to trade timeliness for extra smoothness such that the real-time signal would be more reliable. In that case, the filter customization helps by suppressing high-frequency content in the stopband. Ideally, that extra noise suppression would not alter the phase function in the passband but in practice, it does to some extent, i.e., the additional noise suppression in the stopband slows down the filter output somewhat. Quantitatively, the peak correlation (0.882) between the filter output with $expw = 0$ and the GDP series is found to be at a one month lag w.r.t. GDP, and the second highest correlation (0.874) located at a zero months lag w.r.t. GDP (see Table 2). As for the filter output with $expw = 0.5$, the corresponding numbers are 0.877 (at a one month lag) and 0.850 (at a zero months lag). The peak correlation (0.856) between the filter output with $expw = 1$ and the GDP series is found to be at a one month lag w.r.t. GDP, and the second highest correlation (0.846) located at a two months lag w.r.t. GDP. The presented final indicator is the middle case with $expw = 0.5$.

A.2 Other data and transformations

Other additional data (e.g., other DG ECFIN data, new registered cars, industrial production, producer and consumer price indices, and more financial variables) were tested and found not to yield conceptually different results - their usage is under consideration but excluded from the paper for simplicity and the lack of proper vintage data.

Both seasonally adjusted and unadjusted survey data were tested as inputs. Preliminary results have shown that the use of seasonally adjusted survey data - by Dainties, as published by DG ECFIN - for filter inputs gives the most satisfactory results. Another nuance is whether to use seasonally adjusted or unadjusted GDP for spectral estimates. If the target is lowpass with cut-off to the left of the one-year frequency, both could be used without much difference in outcome. However, since the cut-off is set to the seasonal frequency, one year, seasonally adjusted GDP is used for the target spectral estimate.

A.3 Data revisions

The seasonal adjustment method used by DG ECFIN - Dainties - to seasonally adjust the survey data does not induce revisions. The last month of the non-adjusted data might be slightly revised, thus inducing revisions in SA data. However, these revisions are negligible and thus are ignored.

Another source of potential data revision is revisions in GDP. However, the solution to the revised GDP series is straightforward - GDP series itself is not filtered; instead, only its spectral content is fed to the filter; therefore, a robust filter output is obtained by feeding the filter with estimated spectra of the final/revised GDP by dropping several observations from the end of the series. That is, the described method does not use uncertain first releases, and dropping the last three or five quarters from the estimation routine gives similar results.

Therefore, the plotted indicator's performance is expected not to deteriorate with new real-time data observations entering the estimation routine.

A.4 Amplitude constraint

The amplitude constraint is imposed by setting the filter amplitude to unity at frequency zero. Strictly speaking, the amplitude constraint is not necessary for stationary data. However, unreported results of the author suggest that an unconstrained filter tends to yield inferior real-time performance with respect to the scale fitting compared to the constrained filter, particularly so for small samples and during the great recession period.

A.5 Filter updating frequency

The default filter updating period is set to every quarter, as new GDP data become available. Filter updating is performed in order to take into account the potential structural changes. Moreover, the indicator has been checked for different - less frequent - filter updating frequency and is found to be robust, which is an expected result since the spectral content of the GDP data in the relevant passband is not expected to change rapidly.