Real-time signal extraction with regularized multivariate direct filter approach*

Ginters Buss[†]
Bank of Latvia

April 8, 2014

Abstract

The paper studies the regularized direct filter approach (Wildi, 2012) as a tool for real-time signal extraction using high-dimensional datasets. The paper shows that the filter is able to process high-dimensional datasets by controlling for effective degrees of freedom through longitudinal and crossectional regularization. The paper illustrates the merit of the proposed approach by tracking the medium-to-long-run component in the euro area GDP growth. The created real-time indicators outperform Eurocoin (Altissimo et al., 2010) with respect to timeliness.

Keywords: high-dimensional filtering, real-time estimation, business cycle, coincident indicator, leading indicator
JEL code: C13, C32, E32, E37

*I thank Marc Wildi and anonymous referees for valuable feedback. All remaining errors are my own. 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

Nowadays, gathering rich datasets is relatively easy. A more difficult task is to use them effectively for a particular problem at hand. This paper studies the regularized multivariate direct filter approach (Wildi, 2012) as a tool for real-time signal extraction with many variables.

Wildi (2012) derives a regularized multivariate filter as the successor to his multivariate direct filter approach (Wildi, 2011). The regularization, or shrinkage, is performed in three dimensions: 1) longitudinal shrinkage reduces the number of free filter coefficients by forcing more distant filter coefficients to converge to zero; it is, in spirit, similar to the 'lag decay' term in Minnessota prior (e.g. Doan, Litterman and Sims, 1984) in Bayesian econometrics; 2) cross-sectional shrinkage makes the filter coefficients to behave similarly for similar input series, and, in the limit, forces filter coefficients to be the same for all input series; 3) smoothness restriction makes the filter coefficients change smoothly along the longitudinal/time dimension.

Wildi (2012) does not apply the filter to the data. Therefore, the added value of the current paper is to study how the regularized filter can be implemented in real-time signal extraction exercises, and in particularly, in trend-cycle (in this paper, defined by a cycle of one year and longer) estimation using high-dimensional data sets. By high-dimensionality here I mean at least several dozen of variables, since that many variables makes it difficult or impossible to successfully apply traditional, unregularized methods in real-time signal extraction due to the many free estimated coefficients. Trend-cycle estimation is typical in economics and finance, where there is often short-term noise in the data that does not help in understanding the medium-to-long term trends. E.g, Altissimo et al. (2010) use the same definition of the target signal as in this paper.

The key notion I am using to make the multivariate filter perform decently in out of sample is the effective degrees of freedom (the trace of the smoother matrix; henceforth, e.d.f.). My motivation is that, typically, methods with a few degrees of freedom perform well in out of sample. E.g., in forecasting, such would be the autoregressive moving average model, or the factor model with one or a few factors. Therefore, I propose to fix the desired e.d.f. to a sufficiently small value, and use the above shrinkage terms to achieve the desired e.d.f.. This paper finds that two of the Wildi's proposed three shrinkage terms help in real-time signal extraction problems. The most effective is found to be the longitudinal shrinkage but it cannot be applied excessively since a too short filter cannot discriminate well between the frequencies. I find that it is desirable to allow for at least about a half a year of the effective length of the filter (the length of the filter where its coefficients are beyond the neighborhood of zero). The rest of the e.d.f. are suppressed by the the cross-sectional shrinkage. I find that the cross-sectional shrinkage is useful particularly if the dataset is rather homogeneous.

I illustrate the design of the filter and its real-time output by applying the above method to 72 variables in order to track the medium-to-long-run component of the euro area (henceforth, GDP) growth. The results show that the filter output is robust, and and that such a design can mimic and even outperform in terms of timeliness the established Eurocoin indicator (Altissimo et al. 2010) that is based on a dynamic factor methodology (Forni et al. 2000, 2005). A comparison of the proposed method to the dynamic factor methodology is instructive since both methods have much in common but also feature clear-cut differences. While dynamic factor methodology shrinks the dimension of the dataset to a few unobserved factors and thus has a few parameters to estimate, the regularized filter does not shrink the dimension of the dataset but rather imposes restrictions on coefficient behavior. Therefore, the proposed method can be involved in estimating hundreds or even thousands of coefficients. Nonetheless, controlling for effective degrees of freedom by regularization helps achieving the desired out-of-sample behavior. The paper illustrates this point by estimating more than 800 filter coefficients on less than 150 observations long sample. Also, the dynamic factor methodology of Altissimo et al. (2010) (as most other factor methods) is a two-step procedure: in the first step, it extracts factors from the explanatory dataset, and then maps the targeted variable to the extracted factors. This two-step procedure brings a few difficulties in its practical implementation. First, if irrelevant variables dominate the dataset, the extracted generalized principal components might have little in common with the target. Thus, careful preselection of explanatory variables is a prerequisite for a successful application of factor methodology. In contrast, the regularized filter is a one-step method and thus is more robust to irrelevant variables, since the filter would tend to put smaller weights on irrelevant variables, and higher weights - on more relevant ones. Therefore, the regularized filter requires potentially less work in the variable pre-selection step. Indeed, in this paper, I have done little in the variable pre-selection step. Second, the two-step procedure is more difficult to use to decompose the effects of each explanatory variable; in contrast, the partial effect decomposition is straightforward with the Wildi filter. And third, it is not clear how to use the two-step procedure for efficient forecasting, since different subsets of variables may explain the bulk of variation of the target for different targeted lags or leads. Thus, it is no surprise that Altissimo et al. (2010) limit themselves with coincident signal extraction. On the contrary, targeting a specific lead or lag of a signal is straightforward with Wildi filter.

The paper is structured as follows. Section 2 introduces to the regularized direct filter approach. Section 3 proposes the method for real-time trend-cycle filtration using many variables. Section 4 illustrates the proposed method by creating a real-time indicator for quarterly growth rate of euro area GDP. Appendix A lists the tables. Appendix B describes the filter constraints.

2 Background on regularized multivariate direct filter approach

The regularized multivariate direct filter approach is a regularized version of the multivariate direct filter approach (Wildi, 2011). The latter filter contains many free coefficients whose number increases with filter's (either longitudinal or cross-sectional) dimension, posing threat to decent out of sample performance or even the possibility of estimation. Therefore, similar to standard econometric practices in coefficient or variable shrinkage, it is thoughtful to attempt to shrink filter coefficients. Such an attempt is provided by Wildi (2012) who introduces three shrinkage parameters in a multivariate direct filter approach (Wildi, 2011) and that control for cross-sectional shrinkage, shrinkage along time dimension, and that imposes smoothness of filter coefficients along time dimension. Since the details on the filter can be found in Wildi (2011, 2012), this section summarizes its main elements.

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

$$y_t = \sum_{j=-\infty}^{\infty} \gamma_j x_{t-j}, \tag{2.1}$$

where B is the lag or backshift operator.

A real-time estimate of $\{y_T\}$, given a finite input $\{x_1, \dots, x_T\}$, can be written as

$$\hat{y}_T = \sum_{j=0}^{T-1} b_j x_{T-j}. \tag{2.2}$$

Denote the generally complex transfer functions of filters in (2.1) and (2.2) 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.

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:

$$\int_{-\pi}^{\pi} |\Gamma(\omega) - \hat{\Gamma}(\omega)|^2 dH(\omega) = E[(y_T - \hat{y}_T)^2]. \tag{2.3}$$

A finite sample approximation of MSFE, (2.3), is

$$\frac{2\pi}{T}\sum_{k=-[T/2]}^{[T/2]} \xi_k |\Gamma(\omega_k) - \hat{\Gamma}(\omega_k)|^2 S(\omega_k), \tag{2.4}$$

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

$$\xi_k = \begin{cases} 1 & \text{for } |k| \neq T/2 \\ 1/2 & \text{otherwise,} \end{cases} \tag{2.5}$$

see Brockwell and Davis, 1987, Ch. 10 for the reason for w_k. $S(\omega_k)$ in (2.4) 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, (2.3). 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:

$$S(\omega_k) := I_{Tx}(\omega_k) = \frac{1}{2\pi T}\left|\sum_{t=1}^{T} x_t \exp(-it\omega_k)\right|^2. \tag{2.6}$$

Minimizing expression (2.4) yields the real-time filter output optimally approximated to the ideal output in mean squared error sense. The expression (2.4) 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)). Its multivariate extension is elaborated in Wildi (2011), and a regularized version of the latter - in Wildi (2012).

Recalling that Tikhonov regularization problem (e.g. Tikhonov and Arsenin, 1977) can be cast in the form $(Y - Xb)'(Y - Xb) + \lambda b'b \to \min_b$, the regularized direct filter approach problem (Wildi, 2012) can be written in the form:

$$(Y - Xb)'(Y - Xb) + \lambda_c b' Q_c b + \lambda_d b' Q_d b \to \min_b, \tag{2.7}$$

where Y is a $(T/2+1) \times 1$ vector containing information about the spectral content of the target, X is a $(T/2+1) \times (L+1)(m+1)$ matrix containing information about the spectral content of explanatory variables, b is a $(L+1)(m+1) \times 1$ matrix containing the filter's coefficients, T is the even-valued sample size, L is the lag order of a filter such that a one-sided filter's total length is $L+1$, and m is the number of explanatory variables in addition to the target variable; thus the total number of variables is $m+1$. The two additional expressions of bilinear form represent two different regularization dimensions - cross-sectional shrinkage (subsript 'c') and longitudinal shrinkage (subscript 'd'). I omit the third shrinkage dimension based on my unreported results suggesting its limited usefulness in my applications.

The idea behind the cross-sectional shrinkage is that one would expect the filter coefficients to be similar for similar series. This shrinkage is implemented by imposing constraints on b according to

$$\sum_{u=0}^{m} \left(\left(b_0^u - \frac{1}{m+1} \sum_{u'=0}^{m} b_0^{u'} \right)^2 + \left(b_1^u - \frac{1}{m+1} \sum_{u'=0}^{m} b_1^{u'} \right)^2 + \ldots + \left(b_L^u - \frac{1}{m+1} \sum_{u'=0}^{m} b_L^{u'} \right)^2 \right) \tag{2.8}$$

where b_0^0 denotes the first (zero-lagged) coefficient on the target variable. Expression (2.8) yields a symmetric bilinear form with

$$Q_c = \begin{pmatrix} q_{c,1} \\ q_{c,2} \\ \vdots \\ q_{c,(m+1)*(L+1)} \end{pmatrix} \tag{2.9}$$

where

$$
\begin{aligned}
q_{c,1} &= (1-\frac{1}{m+1},0,\ldots,0|-\frac{1}{m+1},0,\ldots,0|-\frac{1}{m+1},0,\ldots,0|\ldots)\\
q_{c,2} &= (0,1-\frac{1}{m+1},0,\ldots,0|0,-\frac{1}{m+1},0,\ldots,0|0-\frac{1}{m+1},0,\ldots,0|\ldots)\\
q_{c,3} &= (0,0,1-\frac{1}{m+1},0,\ldots,0|0,0,-\frac{1}{m+1},0,\ldots,0|0,0,-\frac{1}{m+1},0,\ldots,0|\ldots)\\
&\ldots\\
q_{c,(m+1)*(L+1)} &= (0,0,\ldots,-\frac{1}{m+1}|0,0,\ldots,-\frac{1}{m+1}|0,0,\ldots,-\frac{1}{m+1}|\ldots|0,0,\ldots,1-\frac{1}{m+1})
\end{aligned}
\tag{2.10}
$$

such that each block separated by $|$ is of length $L+1$. Thus there are 1's on the diagonal of Q_c and periodically arranged $-\frac{1}{m+1}$'s which account for the central means in (2.8). A higher λ_c gives preference for more similar filters across series and the limiting case, $\lambda_c \to \infty$ ensures the filter coefficients are identical across series.

The idea behind the shrinkage across time dimension is that a practitioner might give a preference for the filter coefficients that decay to zero progressively as functions of a lag. This shrinkage is implemented by setting Q_d such that

$$
b'Q_d b = \sum_{u=0}^{m}\sum_{l=0}^{L} \tilde{q}_l (b_l^u)^2, \tag{2.11}
$$

where $\tilde{q}_l$ is the l-th element of

$$
\tilde{q} = \left(q^{0\vee h}, q^{|1-0\vee h|}, q^{|2-0\vee h|}, \ldots, q^{|L-0\vee h|}\right), \tag{2.12}
$$

where q is set to $q := 1+\lambda_d$, $\vee$ denotes a maximization operator, and h is the lag at which filter is estimated, i.e., $h=0$ means a concurrent filter that targets $y_{T-h} = y_T$, $h>0$ means the filter is the smoother, and $h<0$ means the filter is targeted to forecast the signal h periods ahead. When estimating y_{T-h} for $h>0$, a practitioner would want to assign the largest filter weight to observations coinciding with y_{T-h}. Thus, (2.12) ensures that minimum regularization is imposed on lag h (since $q^{h-0\vee h} = q$) and a decay is emphasized symmetrically on both sides away from the target lag h. A higher λ_d ensures a faster coefficient decay to zero as a function of a lag.

Since the regularization is cast in bilinear forms, the problem in (2.7) has an analytic solution.

3 Real-time trend-cycle filtering using high-dimensional datasets

Subsequently, this paper narrows down the topic to real-time trend-cycle extraction problem using non-integrated high-dimensional datasets, though the ideas and notions that follow might be used for other signal-extraction problems, as well. I will be using the term 'effective degrees of freedom', therefore its definition follows straightaway.

3.1 Effective degrees of freedom

In an unconstrained ordinary least squares framework the (regression) degrees of freedom is the number of estimated parameters. Given a well-posed ordinary least squares problem,

$$(Y - Xb)'(Y - Xb) \to \min_b,$$

the fitted values of Y can be written in terms of a hat or smoother matrix, S, which is just a projection matrix, P:

$$\hat{Y} = SY = X(X'X)^{-1}X'Y = PY. \tag{3.1}$$

The degrees of freedom is a trace of the projection matrix:

$$d.f. = \mathrm{tr}(P), \tag{3.2}$$

which equals to rank(X). For a regularized problem as in expression (2.7),

$$(Y - Xb)'(Y - Xb) + \lambda_c b' Q_c b + \lambda_d b' Q_d b \to \min_b,$$

the smoother matrix is no longer an orthogonal projection but the same notion applies. Denoting the fitted value of Y by $\hat{Y}$ and the corresponding smoother matrix by $\tilde{S}$:

$$\tilde{S} = \mathrm{Re}(X)\left(X'X + \lambda_c Q_c + \lambda_d Q_d\right)^{-1} \mathrm{Re}(X)', \tag{3.3}$$

where Re denotes the real part of a complex-valued matrix, such that $\hat{Y} = \tilde{S}Y$, the effective degrees of freedom (or, effective number of parameters) is the trace of $\tilde{S}$:

$$e.d.f. = \mathrm{tr}(\tilde{S}), \tag{3.4}$$

see, e.g. Moody (1992), Hodges and Sargent (2001).

3.2 Optimal shrinkage selection

The two most important filter parameters a user has to choose in order to apply the filter to the data is the longitudinal and the cross-sectional shrinkage parameters.[1] One way to choose them would be using cross-validation by selecting, using a grid search, the parameters that yield the best outcome for the training sample. Given the assumed non-integrated data, one way the best outcome can be measured is by maximizing the correlation between the target variable and the filter output. Dynamic correlation can be used if specific lead/lag is targeted. A practical problem to such an empirical cross-validation is the usually short training sample available for economic applications. My suggestion to reduce the extent of the problem is to define the desired effective degrees of freedom (e.d.f.) and then to select the shrinkage parameters that satisfy the e.d.f. restriction. Such an approach can be useful in real-time filtering, since a small number of e.d.f. would ensure that an 'overparameterization', or the problem of too many free

[1] There are more to choose - the target signal, the explanatory variables and their total number, the data transformation, the length of the filter, filter constraints, etc. but these choices are more common across different methods.

coefficients to be estimated - the problem that often is blamed for inferior out-of-sample versus in-sample performance - does not occur. E.g. if one defined the desired e.d.f. to be three (based on the prior real-time estimation experience with other methods employing a small number of estimated coefficients), it would be hard to imagine that the out-of-sample results could deteriorate due to too many free coefficients. Therefore, one would just need to select the two shrinkage parameters satisfying the pre-selected e.d.f. - based on cross-validation or prior knowledge - and the real-time out-of-sample performance of the filter should be comparable to the one during the training sample. Thus, if one is satisfied with the filter's real-time performance during the training-sample, she could use the filter for real-time signal extraction problems with confidence.

The major drawback of this paper is that I do not come up with an automatic shrinkage parameter selection mechanism, although, as mentioned, a simple cross-validation over a grid of parameters could be one way to go. Meanwhile, the next section illustrates that a simple, informal shrinkage selection can suffice for practical purposes. That informal parameter selection is based on my extensive empirical analysis of the behavior of the filter. Due to the space constraint, I will here just summarize my findings reported more extensively in the working paper. The longitudinal shrinkage is found to be the most effective of the two regularization forms; it effectively reduces the e.d.f. and ensures that the more recent information receives the higher weight, mimicking the other common methods placing more weight on recent information, and thus taking into account the potential structural changes along the time dimension. Nevertheless, the longitudinal shrinkage cannot be implemented too strictly, since it would reduce the effective length of the sample to such an extent that the filter could not be able to effectively distinguish between the frequencies. My findings suggest that, for trend-cycle extraction, it is wise to allow for the effective length of the filter (the filter length for which the coefficient values are beyond the neighborhood of zero) to be at least about half a year. The rest of the e.d.f. are reduced by the cross-sectional shrinkage. Increasing the cross-sectional shrinkage parameter to infinity yields filters for all explanatory variables to converge, and e.d.f. to reduce. Therefore, increasing the extent of the cross-sectional shrinkage is less harmful than a too tight longitudinal shrinkage.

The following section illustrates that this simple informal regularization method yields good out-of-sample results.

4 Tracking economic activity in the euro area in real time

4.1 Target

The filter is set to target the ideal lowpass of quarterly growth of euro area GDP with cut-off wave length 12 months. The quarterly GDP data are taken from 1995Q1 till 2011Q4. The GDP series is linearly interpolated to monthly frequency, logged, quarterly differenced and demeaned before its spectral content enters the filter.

4.2 Explanatory variables

Monthly business and consumer confidence indicators published by the Directorate General for Economic and Financial Affairs, European Commission (henceforth, DG Ecfin) and other short-term monthly indicators are used as explanatory variables. In total, 72 monthly variables are used. The choice of the indicators is based on economic relevance and data availability. Appendix A contains a compete list of input data and their transformations. DG Ecfin data are usually published at the end of 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 seasonally unadjusted and seasonally adjusted data, as the latter is the product of a seasonal adjustment program 'Dainties' that does not revise history as new data come in[2]. The above-mentioned considerations make DG Ecfin data useful for real-time signal extraction problems, as far as they contain valuable information about the target signal. Some other explanatory data happen to be revised but the effect of their revision on the filter output is considered to be of minor extent and therefore the final revision data are used.

All explanatory variables are taken from 1995M1 till 2011M4, standardized to zero mean and unit variance. Integrated data are made non-integrated by suitable transformations. Appendix lists the data and their transformations. Note that GDP is not among the explanatory variables, since it is published with a considerable lag and happens to be revised substantially.

4.3 Real-time indicator design

The main explanatory variables are business and consumer survey data, since they are published with almost no delay and have been correlated well with GDP in the past. These data are, strictly speaking, non-integrated and thus are not subjected to differencing. However, as such, their cycle is lagging behind that of quarterly growth of GDP. Therefore, forecasting should be involved (by setting $h < 0$ in expressions (2.12), (B.1), and (B.2), and applying the time shift constraint that restricts the time shift of the filter at the specific h to vanish at zero frequency, see Appendix B) in order to get a timely quarterly growth signal.

The filter can be subjected to also the amplitude constraint that imposes specific values for the amplitude functions at zero frequency and thus can help ensure that the filter output is of the proper scale (see Appendix B). A priori, it is not clear if the amplitude constraint is necessary, therefore both cases - with and without the amplitude constraint - are shown.

4.3.1 Filter with an amplitude constraint.

Amplitude constraint can help contain the filter output on the right scale but it also counteracts with the time shift constraint by partly neutralizing the latter's effect. Therefore, the lead for the time shift constraint is set to six months ($h = -6$). Also, I here differentiate the value of amplitude constraint at zero frequency to be proportional to the

[2]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

in-sample correlation of explanatory series with the target variable.[3] The lag decay parameter is set to $\lambda_d = 0.4$ and the cross-sectional shrinkage parameter is set to $\lambda_c = 1$. This setting gives about three e.d.f.. Filter coefficients and amplitudes are plotted in Figure 4.1.

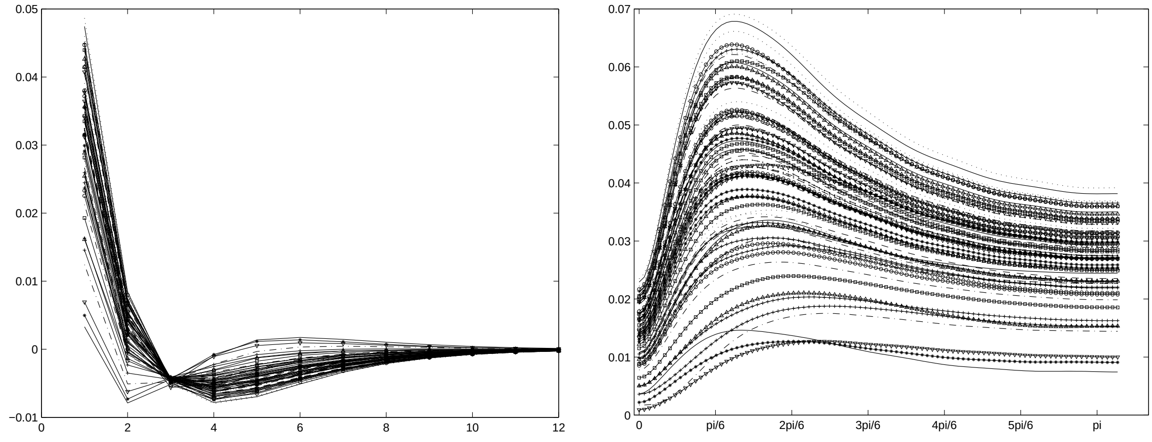

(a) (b)

Figure 4.1: (a) Coefficients for a 72-variables filter with both first- and second-order constraints, $h = -6$, $\lambda_d = 0.4$, $\lambda_c = 1$. (b) Filter amplitudes corresponding to the coefficients in Figure 4.1(a).

The resulting pseudo real-time[4] filter output is plotted in Figure 4.2 along with Eurocoin.

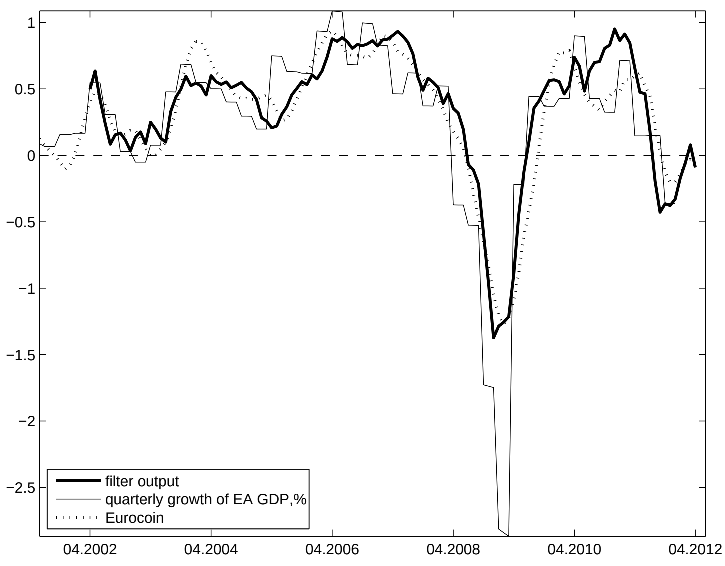

Figure 4.2: Filter output corresponding to filter coefficients in Figure 4.1(a) versus Eurocoin.

[3]The results would be similar if equal weights were used.

[4]Since I have observed the whole sample before designing the filter, I cannot speak about the true real-time performance during the specific sample.

Figure 4.2 shows that the filter output tracks the scale of the target well and precedes Eurocoin on several occasions. Since both indicators target a lowpass of the observed GDP series, traditional mean squared error (MSE) criterion is not suitable for a formal comparison of indicators. Instead, I use the dynamic correlation between the indicator and GDP. The peak correlation of Eurocoin with GDP is located at a two months lag with respect to (henceforth, w.r.t.) GDP, and the second highest correlation being located at a one month lag w.r.t. GDP. For the filter output in Figure 4.2, the peak correlation is located at a one month lag w.r.t. GDP, with the second highest correlation being at a zero months lag w.r.t. GDP (see Table A.1).

Note that the true real-time performance of Eurocoin begins in mid 2009; after that period, the difference between the performances of the two indicators is slightly more evident.

4.3.2 Filter without an amplitude constraint.

With the amplitude constraint absent, it does not interfere with the shift constraint, thus, the targeted lead can be reduced to three months ($h = -3$). Also, absent first-order constraint means more e.d.f., therefore shrinkage is tightened by increasing the cross-sectional shrinkage parameter to $\lambda_c = 5$ (again, the cross-sectional shrinkage is less harmful than the longitudinal shrinkage). This setting gives about eight e.d.f.. Filter coefficients and amplitudes are plotted in Figure 4.3, and the result in 4.4.

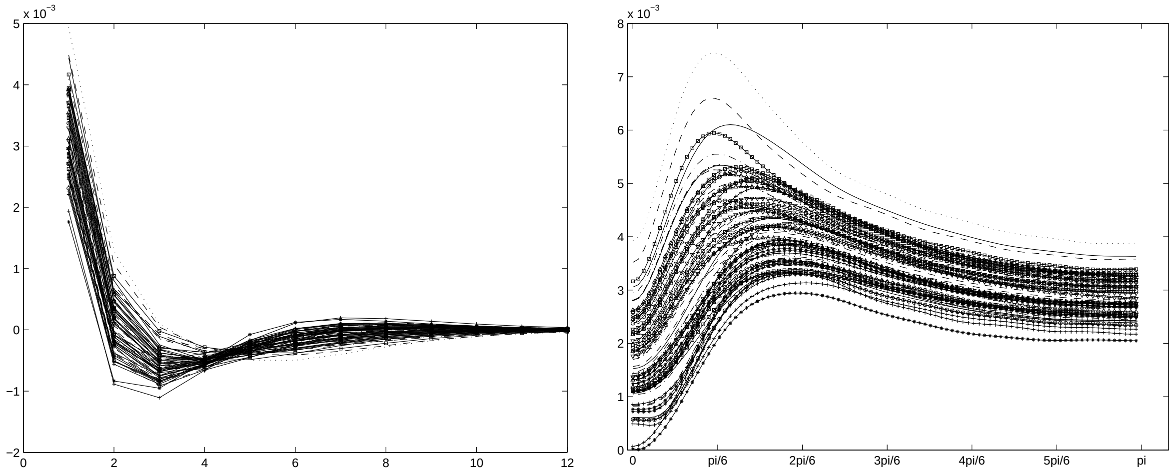

(a) (b)

Figure 4.3: (a) Coefficients for a 72-variables filter with second-order constraint, $h = -3$, $\lambda_d = 0.4$, $\lambda_c = 5$. (b) Filter amplitudes corresponding to the coefficients in Figure 4.3(a).

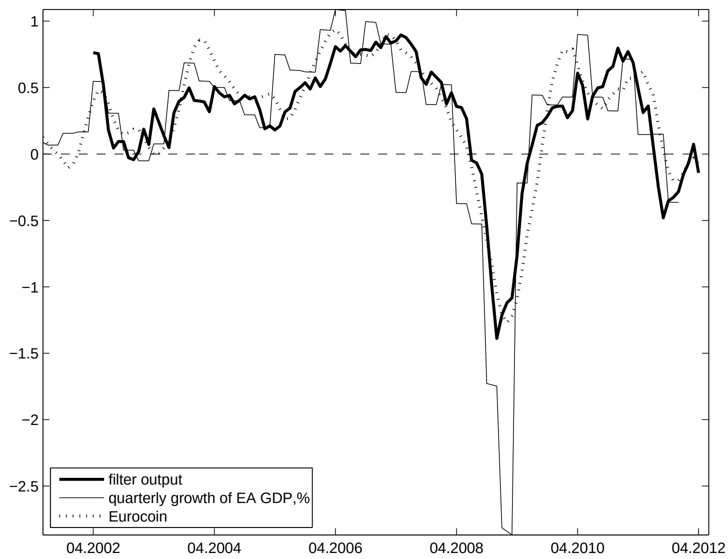

Figure 4.4: Filter output corresponding to filter coefficients in Figure 4.3(a) versus Eurocoin.

The peak correlation of the filter output with GDP is located at a one month lag w.r.t. GDP, and the second highest correlation being located at a zero months lag w.r.t. GDP (see Table A.1).

The two filter designs outperform Eurocoin by about one month in timeliness. The filter design with both constraints appears to be performing slightly better than that without the amplitude constraint. Timeliness is important for real-time indicators, especially if the indicator is lagging with respect to the target variable. Eurocoin is lagging with respect to quarterly GDP growth, on average, by two months. This is a considerable lag in real-time perspective. My created indicators are also lagging the target but the average lag is shorter by one month. To some readers, such an improvement may seem negligible but it is important in real-time setting during uncertain times, i.e., in tracking turning points and the beginning or the end of a recession.

5 Conclusions

This paper considers the regularized multivariate direct filter approach (Wildi, 2012) as a tool for real-time signal extraction using high-dimensional datasets. The key notion I use is the effective degrees of freedom (e.d.f.). The claim in the paper is that if one sets e.d.f. to a sufficiently small number and then search for the shrinkage parameters satisfying the pre-set e.d.f. and maximizing the fit during the training sample (by automatic grid search or, as I have done in this paper - manually), then the out of sample performance of the filter would be comparable to that obtained for the training sample.

Given that the filter's estimation is done in one step, whereas the more traditional factor methods are two-step procedures, the regularized filter has the potential to outcompete the factor methods. To illustrate what one can achieve with the proposed method, I have designed a real-time indicator tracking the medium-to-long-run component of euro area GDP growth using 72 variables. I show that not only the filter can be successfully applied to high-dimensional datasets, but also that the filter output outcompetes the established Eurocoin indicator in timeliness.

Two years have past since I created the above filter design. The indicator has been used since then every month at the Bank of Latvia, and compared with Eurocoin. Although some input data have changed (e.g. due to the new chain-linking of GDP and a few short term indicators) or dropped (due to their changing location or insufficient length), the filter with both constraints have proved time and again to signal more timely the important turning points, compared to Eurocoin. Therefore, for the time being, I have had no thought of discontinuing the indicator.

The greatest drawback of this paper is that I have not come up with an automatic shrinkage selection procedure. To my mind, this is the only noticeable obstacle that refrains widespread usage of the filter. Therefore, future work might be helpful to make the choice of shrinkage parameters automatic for some common signal extraction problems.

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*, MIT Press, vol. 92(4), 1024-1034. DOI:10.1162/REST_a_00045

[2] Baxter M, King RG. 1999. Measuring business cycles: approximate band-pass filters for economic time series, *The Review of Economics and Statistics*, MIT Press, vol. 81(4), 575-593. DOI:10.1162/003465399558454

[3] Brockwell PJ, Davis RA. 1987. *Time Series: Theory and Methods*, New York: Springer Verlag.

[4] Caporello G, Maravall A, Sanchez FJ. 2001. Program TSW reference manual, Banco de Espana Working Papers 0112, Banco de Espana.

[5] Christiano LJ, Fitzgerald TJ. 2003. The band pass filter, *International Economic Review*, vol. 44(2), 435-465. DOI: 10.1111/1468-2354.t01-1-00076

[6] Doan T, Litterman RB, Sims CA. 1984. Forecasting and Conditional Projection Using Realistic Prior Distributions, *Econometric Reviews*, vol. 3(1), 1-100. DOI:10.1080/07474938408800053

[7] Forni M, Hallin M, Lippi M, Reichlin L. 2000. The Generalized Dynamic-Factor Model: Identification And Estimation, *The Review of Economics and Statistics*, vol. 82(4), 540-554. DOI:10.1162/003465300559037

[8] 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.

[9] Hodges JS, Sargent DJ. 2001. Counting degrees of freedom in hierarchical and other richly-parameterized models, *Biometrika*, vol. 88(2), 367-379.

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

[11] King RG, Rebelo ST. 1993. Low frequency filtering and real business cycles, *Journal of Economic Dynamics and Control*, vol. 17(1-2), 207-231. DOI:10.1016/S0165-1889(06)80010-2

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

[13] Moody JE. 1992. The effective number of parameters: an analysis of generalization and regularization in nonlinear learning systems, *Advances in Neural Information Processing Systems 4* (eds J. E. Moody, S. J. Hanson and R. P. Lippmann), 847-854, San Mateo: Morgan Kaufmann.

[14] Tikhonov AN, Arsenin VY. 1977. *Solutions of Ill-Posed Problems*, Washington: V. H. Winston & Sons.

[15] 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`, last accessed in April 3, 2014.

[16] 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`, last accessed in April 3, 2014.

[17] Wildi M. 2012. "Elements of forecasting and signal extraction", IDP-Working Paper, available at `http://blog.zhaw.ch/idp/sefblog/uploads/elements_3007.pdf`, last accessed in April 3, 2014.

Appendix A Tables

Indicators	Dynamic correlation at lag:					
	-3	-2	-1	0	1	2
Eurocoin	0.852	0.882**	0.879*	0.849	0.792	0.706
Proposed filter in Fig. 4.2	0.774	0.846	0.883**	0.870*	0.819	0.736
Proposed filter in Fig. 4.4	0.751	0.824	0.869**	0.862*	0.817	0.740

Note: ** marks the peak correlation; * marks the second highest correlation.

Table A.1: Dynamic correlations of indicators with GDP growth rates.

Variable	Source	transf.yoy	transf.qoq
Real gross domestic product, chain-linked, EA, SA	Eurostat	Δ_{12} log,lin.interp.	Δ log,lin.interp.
Production trend observed in recent month (industry), EA, SA	DG Ecfin	-	-
Assessment of order-book levels (industry), EA, SA	DG Ecfin	-	-
Assessment of export order-book levels (industry), EA, SA	DG Ecfin	-	-
Assessment of stocks of finished products (industry), EA, SA	DG Ecfin	-	-
Production expectations for the months ahead (industry), EA, SA	DG Ecfin	-	-
Selling price expectations for the months ahead (industry), EA, SA	DG Ecfin	-	-
Employment expectations for the months ahead (industry), EA, SA	DG Ecfin	-	-
Confidence Indicator in services, EA, SA	DG Ecfin	-	-
Business situation development over the past 3 months (services), EA, SA	DG Ecfin	-	-
Evolution of the demand over the past 3 months (services), EA, SA	DG Ecfin	-	-
Expectation of the demand over the next 3 months (services), EA, SA	DG Ecfin	-	-
Evolution of the employment over the past 3 months (services), EA, SA	DG Ecfin	-	-
Consumer confidence indicator, EA, SA	DG Ecfin	-	-
Financial situation over last 12 months (consumers), EA, SA	DG Ecfin	-	-
Financial situation over next 12 months (consumers), EA, SA	DG Ecfin	-	-
General economic situation over last 12 months (consumers), EA, SA	DG Ecfin	-	-
General economic situation over next 12 months (consumers), EA, SA	DG Ecfin	-	-
Price trends over next 12 months (consumers), EA, SA	DG Ecfin	-	-
Unemployment expectations over next 12 months (consumers), EA, SA	DG Ecfin	-	-
Major purchases at present (consumers), EA, SA	DG Ecfin	-	-
Savings over next 12 months (consumers), EA, SA	DG Ecfin	-	-
Confidence indicator in retail, EA, SA	DG Ecfin	-	-
Business activity (sales) development over the past 3 months (retail), EA, SA	DG Ecfin	-	-
Volume of stock currently hold (retail), EA, SA	DG Ecfin	-	-
Orders expectations over the next 3 months (retail), EA, SA	DG Ecfin	-	-
Business activity expectations over the next 3 months (retail), EA, SA	DG Ecfin	-	-
Employment expectations over the next 3 months (retail), EA, SA	DG Ecfin	-	-
Confidence indicator in construction, EA, SA	DG Ecfin	-	-
Building activity development over the past 3 months (construction), EA, SA	DG Ecfin	-	-
Employment expectations over the next 3 months (construction), EA, SA	DG Ecfin	-	-
Prices expectations over the next 3 months (construction), EA, SA	DG Ecfin	-	-
Production trend observed in recent month (industry), DE, SA	DG Ecfin	-	-
Assessment of order-book levels (industry), DE, SA	DG Ecfin	-	-
Assessment of stocks of finished products (industry), DE, SA	DG Ecfin	-	-
Production expectations for the months ahead (industry), DE, SA	DG Ecfin	-	-
Employment expectations for the months ahead (industry), DE, SA	DG Ecfin	-	-
Confidence indicator in construction, DE, SA	DG Ecfin	-	-
Confidence indicator in retail, DE, SA	DG Ecfin	-	-
Consumer confidence indicator, DE, SA	DG Ecfin	-	-
Confidence Indicator in services, DE, SA	DG Ecfin	-	-
Production trend observed in recent month (industry), FR, SA	DG Ecfin	-	-
Assessment of order-book levels (industry), FR, SA	DG Ecfin	-	-
Assessment of stocks of finished products (industry), FR, SA	DG Ecfin	-	-
Production expectations for the months ahead (industry), FR, SA	DG Ecfin	-	-
Employment expectations for the months ahead (industry), FR, SA	DG Ecfin	-	-
Confidence indicator in construction, FR, SA	DG Ecfin	-	-
Confidence indicator in retail, FR, SA	DG Ecfin	-	-
Consumer confidence indicator, FR, SA	DG Ecfin	-	-
Confidence Indicator in services, FR, SA	DG Ecfin	-	-
Production trend observed in recent month (industry), IT, SA	DG Ecfin	-	-
Assessment of order-book levels (industry), IT, SA	DG Ecfin	-	-
Assessment of stocks of finished products (industry), IT, SA	DG Ecfin	-	-
Production expectations for the months ahead (industry), IT, SA	DG Ecfin	-	-
Employment expectations for the months ahead (industry), IT, SA	DG Ecfin	-	-
Confidence indicator in construction, IT, SA	DG Ecfin	-	-
Confidence indicator in retail, IT, SA	DG Ecfin	-	-
Consumer confidence indicator, IT, SA	DG Ecfin	-	-
Production trend observed in recent month (industry), ES, SA	DG Ecfin	-	-
Assessment of order-book levels (industry), ES, SA	DG Ecfin	-	-
Assessment of stocks of finished products (industry), ES, SA	DG Ecfin	-	-
Production expectations for the months ahead (industry), ES, SA	DG Ecfin	-	-
Employment expectations for the months ahead (industry), ES, SA	DG Ecfin	-	-
Confidence indicator in construction, ES, SA	DG Ecfin	-	-
Confidence indicator in retail, ES, SA	DG Ecfin	-	-
Consumer confidence indicator, ES, SA	DG Ecfin	-	-
Confidence Indicator in services, ES, SA	DG Ecfin	-	-
Industrial production index B-D;F, EA, SA	Eurostat	Δ_{12} log	Δ log
Industrial production index C, EA, SA	Eurostat	Δ_{12} log	Δ log
Producer price index C, EA, NSA	Eurostat	Δ_{12} log	Δ log
Turnover index in retail trade except for motor vehicles, deflated, EA, NSA	Eurostat	Δ_{12} log	Δ log
The US share price index, US, NSA	Eurostat	Δ_{12} log	Δ log
The EA share price index, EA, NSA	Eurostat	Δ_{12} log	Δ log

Table A.2: The euro area dataset.

Note: Δ - difference operator, SA - seasonally adjusted, NSA - not seasonally adjusted, EA - euro area, DE - Germany, FR - France, IT - Italy, ES - Spain, US - United States.

Appendix B Filter constraints

The first order, or amplitude, constraint imposes specific values for the amplitude functions at zero frequency. For a bandpass, one would typically set amplitudes at zero frequency to be zero ensuring that a bandpass puts zero weight at trend frequency, while for a univariate lowpass one would typically set amplitude at zero frequency to unity to ensure that a lowpass tracks the level/scale of the target.

For a multivariate filter, the optimal constrained level of the amplitude at zero frequency is less clear cut. That level can be set to an inverse of the number of explanatory variables for all the variables if all explanatory variables follow about the same trend. However, the latter might not always be the case and thus a better outcome could be obtained by differentiating the amplitude constraint at zero frequency for various explanatory variables. An example of such a differentiation of the constraint is provided in the empirical section.

This constraint is implemented by restricting:

$$b^u_{-h} + b^u_{-(h-1)} + \ldots + b^u_{L-h} = w^u, \tag{B.1}$$

where w^u is the value at which the transfer function for a variable u is set at zero frequency, and h is the targeted lag.

The second order, or time-shift, constraint restricts the time shift of the filter at zero frequency to vanish, and is related to assuming the target variable has two unit-roots at zero frequency, in which case both first and second order constraints would be implemented. In practice, however, the usage of the constraints are up to the practitioner's agenda, and one could use the time shift constraint without imposing the amplitude constraint, the combination of the constraints that cannot be straightforwardly imposed in the time domain. The second order constraint is imposed by forcing the derivative of the transfer function at zero frequency to vanish, which results in the following coefficient constraint:

$$-hb^u_{-h} + (1-h)b^u_{1-h} + (2-h)b^u_{2-h} + \ldots + b^u_1 + 2b^u_2 + \ldots + (L-h)b^u_{L-h} = 0, \tag{B.2}$$

where h is the targeted lag.

Both constraints can be implemented by selecting any two of the coefficients but is implemented by constraining b^u_0 and b^u_1, so as to avoid a conflicting situation between these constraints and the regularization.

The constrained regularized filter problem is solved by rewriting filter coefficient vector b as

$$b = Rb_f + c, \tag{B.3}$$

where b_f is the vector of freely determined filter coefficients, plugging (B.3) in (2.7), solving for b_f, and then plugging the estimate of b_f into (B.3) to get the estimate of b; see Wildi (2012) for details.