Forecasting and signal extraction with regularized multivariate direct filter approach*

Ginters Buss†
Bank of Latvia

February 22, 2013

Abstract

The paper studies the regularized direct filter approach as a tool for high-dimensional filtering and real-time signal extraction. It is shown that the regularized filter is able to process high-dimensional data sets by controlling for effective degrees of freedom and that it is computationally fast. The paper illustrates the features of the filter by tracking the medium-to-long-run component in GDP growth for euro area, including the replication of the Eurocoin-type behavior, as well as producing more timely indicators. Further robustness check is performed on a less homogeneous Latvia's dataset. The resulting real-time indicators are found to track economic activity in timely and robust manner. The regularized direct filter approach is found to be a promising tool for both concurrent estimation and forecasting using high-dimensional datasets.

Keywords: high-dimensional filtering, real-time estimation, coincident indicator, leading indicator, parameter shrinkage, business cycles, dynamic factor model
JEL code: C13, C32, E32, E37

*The author thanks Marc Wildi for valuable discussions. 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@bank.lv.

1 Introduction

Nowadays, gathering rich datasets is relatively easy. A more difficult exercise is effectively using them for a particular problem at hand. This paper adds to the forecasting and high-dimensional estimation literature (see e.g. ridge regression (e.g. Tikhonov and Arsenin, 1977; Hoerl and Kennard, 1970), lasso (Tibshirani, 1996), least angle regression (Efron et al, 2004), Bayesian shrinkage (e.g. Doan, Litterman and Sims, 1984), factor models (Stock and Watson, 2002; Forni et al., 2000, 2005)) by exploring the properties of the regularized multivariate direct filter approach (in short, RMDFA; Wildi, 2012) in signal extraction and forecasting with many variables.

Wildi (2012) derives the regularized multivariate filter as a successor to an unregularized direct filter approach (Wildi, 2011) but does not study its properties either on generated or on real-world data. This paper is the first paper known to the author that studies and implements the RMDFA on data. Furthermore, this paper studies how RMDFA can be implemented on high-dimensional real-world data sets. This paper finds that filter regularization helps in real-time signal extraction since it helps control for effective degrees of freedom, thus it allows to control for overfitting that can have degrading effects in out-of-sample performance. Another advantage of a regularized filter is that it allows for high-dimensional data to enter the filter and therefore further robustify the outcome. As it is shown in the paper, a particular regularization feature used in the paper might remind about the 'lag decay' term in Minnessota prior (e.g. Doan, Litterman and Sims, 1984) in Bayesian econometrics. Forcing more distant filter coefficients to zero both restricts degrees of freedom and effectively shortens the filter, thus making it more responsive to changing environment. Another regularization feature studied in the paper is cross-sectional shrinkage that makes filter coefficients to behave similarly for similar input series. The cross-sectional shrinkage has been found to be useful particularly if the dataset is rather homogeneous.

As an application, the filter is applied on up to 72 variables in order to track the medium-to-long-run component of the euro area GDP growth. Both yearly and quarterly growth rates of GDP are considered. The results show that the filter output is robust, and is able to both mimic and produce more timely indicators than an established Eurocoin indicator (Altissimo et al. 2010) that is based on a dynamic factor methodology (Forni et al. 2000, 2005). A comparison of RMDFA to dynamic factor methodology of Forni et al. (2005) is especially important since both methods have much in common but also feature some clear-cut differences. First, while dynamic factor methodology shrinks the dimension of the dataset to a few unobserved factors and thus has a few parameters to estimate, the RMDFA does not shrink the dimension of the dataset but rather imposes restrictions on coefficient behavior. Therefore, RMDFA can be involved in computing hundreds or even thousands of coefficients. Nonetheless, controlling for effective degrees of freedom helps avoid overparameterization problem and thus achieve good out of sample behavior. The paper illustrates this point by computing more than 800 filter coefficients on less than 150 observations long sample. Second, dynamic factor methodology of Forni et al. (2005) (as most other factor methods, including Stock and Watson, 2002) extracts factors from the explanatory dataset independently of what is the target variable. If irrelevant variables dominate the dataset, the extracted generalized principal components would have little in common with the target. Thus, careful pre-selection of explanatory

variables is a prerequisite for a successful application of factor methodology. In contrast, the RMDFA is more robust to such an error since the filter would put smaller weights on irrelevant variables, and higher weights - on more relevant ones. Therefore, RMDFA requires potentially less work in the variable pre-selection step.

As a robustness check, the filter is applied on a less homogeneous Latvia's dataset.

A possible downside of the RMDFA compared to the factor methodology is its many hyperparameters a user has to choose. The choice of hyperparameters is problem-specific. The existence of hyperparameters should not be the most critical aspect of the method. There are many popular methods widely used in applied econometrics that involve hyperparameters, e.g. the Bayesian approach. It might be helpful to endogenize the choice of some hyperparameters at least for some common problems but this aspect is left for future research.

The paper is structured as follows. Section 2 introduces new regularization features in the direct filter approach. Sections 3.1 and 3.2 illustrate the new features of the filter by creating indicators for yearly and quarterly growth rates of euro area GDP, respectively. Section 4 performs a robustness check on a less homogeneous Latvia's dataset. Section 5 concludes. Appendix lists the data and their transformations.

2 Regularized multivariate direct filter approach

The regularized multivariate direct filter approach is a regularized version of the multivariate direct filter approach (Wildi, 2011), which has been found to be useful in creating real-time indicators (Buss, 2012).

However, the unregularized multivariate direct filter contains many parameters whose number increases with filter's dimension. Thus, the filter in Wildi (2011) cannot be too long or cannot contain tens of macroeconomic variables due to the limited sample size typically observed in macroeconomics, otherwise the filter would be overparameterized and the filter output would be of poor quality in out of sample. One way of increasing the cross-sectional dimension of the filter would be to decrease the length of the filter, and that indeed has been coded in the algorithm used in Buss (2012). However, the length of the filter cannot be decreased infinitely since it is bounded to zero and a too short filter would result in deteriorating quality of its output. Therefore, similar to standard econometric practices in parameter shrinkage, it would be thoughtful to attempt to shrink filter parameters as well, in order to allow for controlling for effective degrees of freedom and using high-dimensional data sets. Such an attempt is done in Wildi (2012) that 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. The three shrinkage dimensions can be imposed in any of their combinations, or all of them can be set to zero such that the new filter replicates the one discussed in Wildi (2011).

In order to introduce the new regularization features, this paper builds on the classical filtration problem. Since details on technicalities can be found in Wildi (2011, 2012), this section just summarizes the main elements of a customized filter necessary to introduce the new regularization features later in the section.

Denote y_T as the output of a symmetric, possibly bi-infinite filter, $\sum_{j=-\infty}^{\infty} \gamma_j L^j$, ap-

plied on input series x_T:

$$\begin{aligned} y_T &= \sum_{j=-\infty}^{\infty} \gamma_j L^j x_T \\ &= \sum_{j=-\infty}^{\infty} \gamma_j x_{T-j}, \end{aligned} \tag{1}$$

where L is called the lag or backshift operator. A real-time estimate of y_T is

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

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.

For a stationary process 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:

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

where $H(\omega)$ is the unknown spectral distribution of x_T. A finite sample approximation of the MSFE, (3), is

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

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

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

see Brockwell and Davis, 1987, Ch. 10 for the reason for w_k. $S(\omega_k)$ in (4) can be interpreted as an estimate of the unknown spectral density of x_T, which can be any spectral estimate, e.g., the one of white noise (Baxter and King, 1999), random walk (Christiano and Fitzgerald, 2003, and its multivariate extension, Valle e Azevedo (2011)), an ARIMA-based spectral estimate as used in the TRAMO/SEATS seasonal adjustment procedure (Caporello, Maravall and Sanchez, 2001), or the specific ARIMA(0,2,2) process underlying the Hodrick-Prescott filter (Hodrick and Prescott, 1997; King and Rebelo, 1993; Maravall and Rio, 2001). However, 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 a 'sufficient statistic' - periodogram, $I_{Tx}(\omega_k)$ - as $S(\omega_k)$ in (4):

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

Minimizing expression (4) yields the real-time filter output optimally approximated to the ideal output in mean squared error sense. Yet, Wildi (2008) proposes a customized version of (4).

2.1 Univariate direct filter approach

Rewrite discrete version MSFE, (4), as

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

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*.

Univariate analysis is of limited usefulness thus we turn now to multiple-series analysis.

2.2 Multivariate direct filter approach

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

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

where

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

In addition to x_t, assume there are m additional explanatory variables z_{jt}, $j = 1, \ldots, m$ that might help improve the real-time estimate of y_t. Then, the second expression in the modulus on the second line of (9), $\hat{\Gamma}_X(\omega_k)\Xi_{Tx}(\omega_k)$, becomes

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

where

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

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

are the one-sided transfer functions applied to the explanatory variables, and $\Xi_{Tx}(\omega_k)$, $\Xi_{Tz_n}(\omega_k)$ are the corresponding DFTs. Then, the multivariate version of (9) can be written as

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

2.3 Rewriting filtration problem in a least squares form

In order to conveniently define the regularized filter problem, the above multivariate filtration problem is rewritten in a least squares form, see Wildi (2012) for details; this subsection explains how it is done while the next subsection introduces the regularization problem.

Define X such that its k-th row, X_k, is:

$$X_k' = (1 + I_{k>0}) \operatorname{Vec} \begin{pmatrix} \Xi_{Tx}(\omega_k) & \exp(-i\omega_k)\Xi_{Tx}(\omega_k) & \cdots & \exp(-iL\omega_k)\Xi_{Tx}(\omega_k) \\ \Xi_{Tz_1}(\omega_k) & \exp(-i\omega_k)\Xi_{Tz_1}(\omega_k) & \cdots & \exp(-iL\omega_k)\Xi_{Tz_1}(\omega_k) \\ \Xi_{Tz_2}(\omega_k) & \exp(-i\omega_k)\Xi_{Tz_2}(\omega_k) & \cdots & \exp(-iL\omega_k)\Xi_{Tz_2}(\omega_k) \\ \vdots & \vdots & \vdots & \vdots \\ \Xi_{Tz_m}(\omega_k) & \exp(-i\omega_k)\Xi_{Tz_m}(\omega_k) & \cdots & \exp(-iL\omega_k)\Xi_{Tz_m}(\omega_k) \end{pmatrix}, \tag{15}$$

where L is the filter length, and $I_{k>0} = 0$ for $k = 0$ and $I_{k>0} = 1$ for $k = 1, 2, \ldots, [T/2]$. Define vectors b and Y as

$$b = \operatorname{Vec} \begin{pmatrix} b_{x0} & b_{z_1 0} & b_{z_2 0} & \cdots & b_{z_m 0} \\ b_{x1} & b_{z_1 1} & b_{z_2 1} & \cdots & b_{z_m 1} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ b_{xL} & b_{z_1 L} & b_{z_2 L} & \cdots & b_{z_m L} \end{pmatrix}, \qquad Y = \begin{pmatrix} \Gamma(\omega_0)\Xi_{Tx}(\omega_0) \\ 2\Gamma(\omega_1)\Xi_{Tx}(\omega_1) \\ 2\Gamma(\omega_2)\Xi_{Tx}(\omega_2) \\ \vdots \\ 2\Gamma(\omega_{[T/2]})\Xi_{Tx}(\omega_{[T/2]}) \end{pmatrix}. \tag{16}$$

Neglecting a constant $2\pi/T$ and the practically negligible w_k, (14) with $W(\omega_k) = 1$ can be rewritten as

$$(Y - Xb)'(Y - Xb) \to \min_b \tag{17}$$

Since X and Y are complex-valued, the solution to (17) would be complex-valued as well. A real-valued b can be obtained by rotating X and Y s.t. the value of the metric in (17) is unaffected:

$$\begin{aligned} X_{k,rot} &= X_k \exp\left(-i \arg(\Gamma(\omega_k)\Xi_{Tx}(\omega_k)) + ih\omega_k\right) \\ Y_{rot} &= |Y|, \end{aligned} \tag{18}$$

where $X_{k,rot}$ is the k-th row of X_{rot}, and h is the lag at which filter is estimated, i.e. $h = 0$ for a concurrent filter that targets $y_{T-h} = y_T$, $h > 0$ for a smoother, and $h < 0$ for

forecasting the signal. A real-valued b thus can be obtained from solving

$$(Y_{rot} - X_{rot}b)'(Y_{rot} - X_{rot}b) \to \min_b . \tag{19}$$

For the customized multivariate filter ($W(\omega_k) \neq 1$), define

$$X^{cust}_{k,rot} = X_{k,rot}\sqrt{W(\omega_k, expw, cut)} \tag{20}$$

$$Y^{cust}_{rot} = \begin{pmatrix} |\Gamma(\omega_0)\Xi_{Tx}(\omega_0)|\sqrt{W(\omega_0, expw, cut)} \\ 2|\Gamma(\omega_1)\Xi_{Tx}(\omega_1)|\sqrt{W(\omega_1, expw, cut)} \\ \vdots \\ 2|\Gamma(\omega_{[T/2]})\Xi_{Tx}(\omega_{[T/2]})|\sqrt{W(\omega_{[T/2]}, expw, cut)} \end{pmatrix}, \tag{21}$$

where $X^{cust}_{k,rot}$ is the k-th row of X^{cust}_{rot}. Then, the least-squares form for the customized filter problem can be written as

$$(Y^{cust}_{rot} - X^{cust}_{rot}b)'(Y^{cust}_{rot} - X^{cust}_{rot}b) \to \min_b, \tag{22}$$

which collapses to (19) for $expw = 0$.

We are now ready to introduce the regularized filter problem.

2.4 Regularization

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 introduced in Wildi (2012) is of the familiar form:

$$(Y^{cust}_{rot} - X^{cust}_{rot}b)'(Y^{cust}_{rot} - X^{cust}_{rot}b) + \lambda_s b'Q_s b + \lambda_c b'Q_c b + \lambda_d b'Q_d b \to \min_b, \tag{23}$$

where the three additional expressions of bilinear form represent three different regularization directions - coefficient smoothness (subscript 's'), cross-sectional shrinkage (subsript 'c'), and shrinkage along time dimension (subscript 'd'). Let us discuss each in turn.

The idea behind the smoothness restriction is that filter coefficients should not change too erratically as functions of a lag. The Q_s matrix of size $(L+1) \times (L+1)$ is such that

$$b'Q_s b = \sum_{u=0}^{m}\sum_{l=2}^{L}\left((1-L)^2 b^u_l\right)^2, \tag{24}$$

where $(1-L)^2 b^u_l = b^u_l - 2b^u_{l-1} + b^u_{l-2}$ is the second order difference of b^u_l, $l = 0, \ldots, L$, and $u = 0, \ldots, m$. Therefore, the term in (24) is a measure for the quadratic curvature of filter coefficients - if coefficients decay linearly as functions of a lag then this term vanishes. Thus, in the limiting case when $\lambda_s \to \infty$, the filter coefficients are restricted to be linear functions of a lag.

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

$$\sum_{u=0}^{m}\left(\left(b_0^u-\frac{1}{m+1}\sum_{u'=0}^{m}b_0^{u'}\right)^2+\left(b_1^u-\frac{1}{m+1}\sum_{u'=0}^{m}b_1^{u'}\right)^2+\ldots+\left(b_L^u-\frac{1}{m+1}\sum_{u'=0}^{m}b_L^{u'}\right)^2\right) \tag{25}$$

which yields a symmetric bilinear form with

$$Q_c=\begin{pmatrix} q_{c,1} \\ q_{c,2} \\ \vdots \\ q_{c,(m+1)*(L+1)} \end{pmatrix} \tag{26}$$

where

$$\begin{aligned}
q_{c,1} &= (1-\frac{1}{m+1},0,\ldots,0|-\frac{1}{m+1},0,\ldots,0|-\frac{1}{m+1},0,\ldots,0|\ldots) \\
q_{c,2} &= (0,1-\frac{1}{m+1},0,\ldots,0|0,-\frac{1}{m+1},0,\ldots,0|0-\frac{1}{m+1},0,\ldots,0|\ldots) \\
q_{c,3} &= (0,0,1-\frac{1}{m+1},0,\ldots,0|0,0,-\frac{1}{m+1},0,\ldots,0|0,0,-\frac{1}{m+1},0,\ldots,0|\ldots) \\
&\ldots \\
q_{c,(m+1)*(L+1)} &= (0,0,\ldots,-\frac{1}{m+1}|0,0,\ldots,-\frac{1}{m+1}|0,0,\ldots,-\frac{1}{m+1}|\ldots|0,0,\ldots,1-\frac{1}{m+1})
\end{aligned} \tag{27}$$

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 (25).

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.

Finally, 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

$$b'Q_d b=\sum_{u=0}^{m}\sum_{l=0}^{L}\tilde{q}_l(b_l^u)^2, \tag{28}$$

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

$$\tilde{q}=\left(q^{0\vee h},q^{|1-0\vee h|},q^{|2-0\vee h|},\ldots,q^{|L-0\vee h|}\right), \tag{29}$$

where q is set to $q:=1+\lambda_d$, $\vee$ denotes a $\max(\cdot)$ function, and h signifies 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, (29) 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 (23) has an analytic solution. Setting $\lambda_s = \lambda_c = \lambda_d = 0$ gives the unregularized filter problem in (22). Or, setting $expw = 0$ but letting some of the regularization lambdas positive gives the regularized classical multivariate filter problem. This paper has found the lag decay shrinkage the most useful of the three regularization types for the application at hand, followed by the cross-sectional shrinkage.

The next section describes an application of the filter obtained by solving (23) subject to two potential constraints - first- and/or second-order constraints which are explained in the following subsection.

2.5 Level and time shift constraints

The first order constraint imposes specific values for the amplitude functions in frequency zero. For a bandpass, one would typically set amplitudes at frequency zero to be zero ensuring that a bandpass puts zero weight on trend frequency, while for a univariate lowpass one would typically set amplitude at frequency zero to unity to ensure that a lowpass tracks the level/scale of the target; such restriction is related to assuming the target has a unit root at frequency zero, i.e., it is a first order integrated process.

For a multivariate filter the optimal constrained level of the amplitude at frequency zero 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 frequency zero for various explanatory variables. An example of such a differentiation of the constraint is provided in the empirical section.

In practice, one can choose to use or not to use the level constraint at ones own discretion. This constraint is implemented by restricting:

$$b^u_{-h} + b^u_{-(h-1)} + \ldots + b^u_{L-h} = w^u, \tag{30}$$

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

The second order 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 in frequency zero, 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 level constraint, the combination of the constraints that can not be straightforwardly imposed in the time domain. The second order constraint is imposed by forcing the derivative of the transfer function at frequency zero to vanish, which results in the following coefficient constraint:

$$-hb^u_{-h} + (1-h)b^u_{1-h} + (2-h)b^u_{2-h} + \ldots + b^u_1 + 2b^u_2 + \ldots + (L-h)b^u_{L-h} = 0, \tag{31}$$

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

$$b = Rb_f + c, \tag{32}$$

where b_f is the vector of freely determined filter coefficients, plugging (32) in (23), solving for b_f, and then plugging the estimate of b_f into (32) to get the estimate of b; see Wildi (2012) for details.

2.6 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,

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

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

$$\hat{Y} = SY = X(X'X)^{-1}X'Y = PY. \tag{33}$$

The degrees of freedom is trace of the projection matrix:

$$d.f. = \mathrm{tr}(P), \tag{34}$$

which equals to $\mathrm{rank}(X)$.

For a regularized problem as in expression (23),

$$(Y_{rot}^{cust} - X_{rot}^{cust}b)'(Y_{rot}^{cust} - X_{rot}^{cust}b) + \lambda_s b'Q_s b + \lambda_c b'Q_c b + \lambda_d b'Q_d b \rightarrow \min_b,$$

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

$$\tilde{S} = \mathrm{Re}(X_{rot}^{cust}) \left((X_{rot}^{cust})'X_{rot}^{cust} + \lambda_s Q_s + \lambda_c Q_c + \lambda_d Q_d\right)^{-1} \mathrm{Re}(X_{rot}^{cust})', \tag{35}$$

such that $\hat{Y}_{rot}^{cust} = \tilde{S}Y_{rot}^{cust}$, the effective degrees of freedom (or, effective number of parameters) is the trace of $\tilde{S}$:

$$e.d.f. = \mathrm{tr}(\tilde{S}), \tag{36}$$

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

Effective degrees of freedom are useful to use for controlling for overfitting and thus for controlling for out-of-sample performance.

3 Tracking economic activity in the euro area

3.1 Tracking trendcycle in yearly growth of GDP

This section discusses the new regularization features of the multivariate filter by creating two real-time indicator designs for the euro area GDP. The two indicator designs differ by the input data transformation and according modifications in their filter designs. The first design discussed in this subsection considers yearly growth rates of real GDP, while the second design discussed in the next subsection considers quarterly growth rates of real GDP. The potential user of the indicators then can form a subjective preference between the two. More discussion follows in the respective subsections discussing each design separately, starting with the yearly growth design.

3.1.1 Target

The filter is set to target an ideal (i.e., filter's amplitude is unity in the passband and zero in the stopband) lowpass of yearly growth of real GDP with cut-off wave length 12 months. The quarterly GDP data are taken from 1995Q1 till 2011Q4, as published by Eurostat. The data are linearly interpolated to monthly frequency, logged, yearly differenced and demeaned before their spectral content enters the filter.

3.1.2 Explanatory variables

Monthly business and consumer confidence indicators published by DG Ecfin and other monthly variables 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 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[1]. The above-mentioned considerations make Ecfin data comfortable for real-time filtration. 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 2012M4, standardized to zero mean and unit variance. Integrated data are made non-integrated by suitable transformations. Appendix lists the data and their transformations.

3.1.3 Regularization features

We now study the regularization features of the filter. For visual tractability and due to numerical issues (an unregularized filter crashes for a high-dimensional input data when the number of estimated filter parameters reaches the number of sample observations)

[1]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

only nine survey variables are used to analyze the filter effect. More data are added later in the section. The nine variables are business and consumers confidence data: 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, consumers confidence indicator, and confidence indicator in services.

In order to motivate the chosen transformation of data, it is illustrative to plot the transformed target variable and explanatory variables. Figure 1(a) shows standardized yearly growth of EA GDP versus standardized business and consumers data. Explanatory data are well aligned with the yearly growth of GDP. Extracting the cross-sectional mean and the first principal component of the standardized explanatory data and plotting against standardized yearly growth of GDP shows that both the mean and the first principal component explain yearly changes in GDP well, and there is not much difference in the performance of the mean versus the principal component, see Figure 1(b).

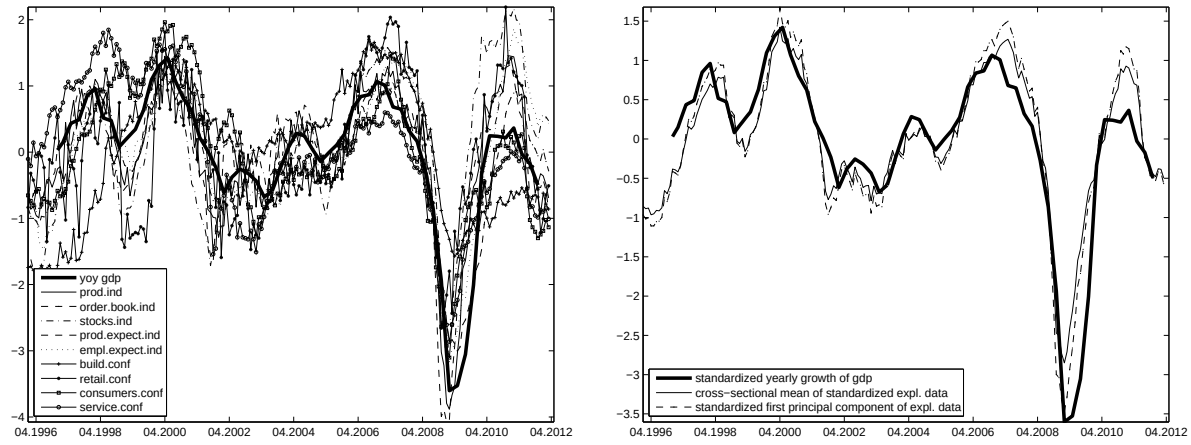

(a) (b)

Figure 1: (a) Yearly growth of GDP versus business and consumers data, all normalized to zero mean and unit variance. (b) Yearly growth of GDP versus the cross-sectional mean and the first principal component of business and consumers data, all normalized to zero mean and unit variance.

Clearly, there is not much to improve upon the simple cross-sectional mean or the first principal component of the explanatory variables as it comes to tracking cyclical developments in normalized yearly growth of euro area GDP; slightly more difficult is to track non-normalized target, see the results below. The cross-sectional mean or principal components could be used as filter inputs but this paper shows that it is not necessary to do so and that one can use the original, possibly high-dimensional data as the input and potentially benefit from the richness of data.

In order to understand the extent of overparameterization in an unregularized multivariate filter, consider an unconstrained filter applied on the considered nine variables targeting an ideal lowpass of yearly growth of EA GDP with the cut-off wave length 12 months. The filter length is set to be fixed 12 observations, for simplicity. While the esti-

mation routine can estimate a 9-variable filter on the full sample (178 observations long), it crashes for smaller subsamples because of the degrees of freedom exceed the number of observations for all subsamples shorter than 9*12=108 observations. A further reduction of filter length might be a temporary solution but not for long and not without consequences on output quality. Therefore, an unconstrained 9-variable filter output is infeasible for the considered data samples. Thus, some sort of parameter shrinkage is necessary. In order to illustrate the effect of the parameter shrinkage induced by the regularized filter, consider the estimated filter coefficients for an unconstrained and unregularized 9-variable filter on the full sample. The number of estimated parameters is 9 variables times 12 observations long filter which gives 108 parameters to estimate on a 178-observations long sample, which gives only 70 residual degrees of freedom. Figure 2(a) shows that the estimated filter coefficients look erratic, unsmooth and do not show a similar behaviour between the variables nor an evident decay towards zero with an increasing lag. Figure 2(b) shows the (rather chaotic) filter amplitudes corresponding to the coefficients in Figure 2(a); it will be useful to analyze how the amplitudes change with various constraints and regularization restrictions.

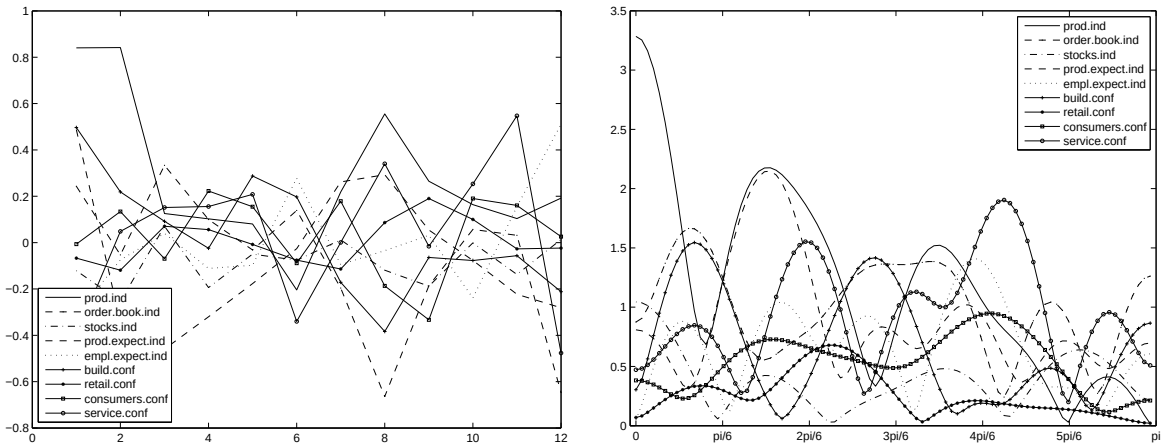

(a) (b)

Figure 2: (a) 9-variable filter coefficients without regularization and filter constraints. The estimated filter coefficients look erratic, unsmooth and do not show a similar behaviour between the variables nor an evident decay towards zero with an increasing lag coefficients. (b) Filter amplitudes corresponding to the coefficients in Figure 2(a).

We will now witness the effect of filter constraints and the regularization features first applied each one at a time and then in a potentially useful combination.

The first order restriction imposes filter amplitude to be a specific value at frequency zero. For a univariate lowpass a natural value of the amplitude at frequency zero is unity in order to ensure that the scale of the output is comparable to the scale of the target signal. For a multivariate filter, things are not that straightforward since all the input series generally do not possess the same trend, therefore restricting all amplitudes to be of the same value at frequency zero might be suboptimal. If all the input series do follow a common trend then it would be natural for a multivariate lowpass to set amplitudes at frequency zero to be inverse of the number of input series, so that summing

over the amplitudes would result in unity at frequency zero. Since the input series used in this exercise have a somewhat similar behaviour between each other, the latter approach is used in this exercise; however, there might be potential gains by using a more sophisticated amplitude constraint that would differentiate amplitude values at frequency zero for different input series; one such approach is discussed later in the section when applying the filter on a higher-dimensional set of explanatory variables.

The first order constraint restricts one d.f. per input series, thus nine d.f. are restricted for an unregularized nine-variable filter.

Figures 3(a) and 3(b) show that the effect of amplitude constraint is slightly more dispersed coefficients (the scale of the graph has changed) as well as slightly more exploded amplitudes. Thus, the first order constraint per se does not seem to be of much help for an otherwise ill-posed high-dimensional filter. Note that the amplitude constraint is binding for almost all series since the unconstrained amplitudes at frequency zero are dispersed far from the constrained value (1/9).

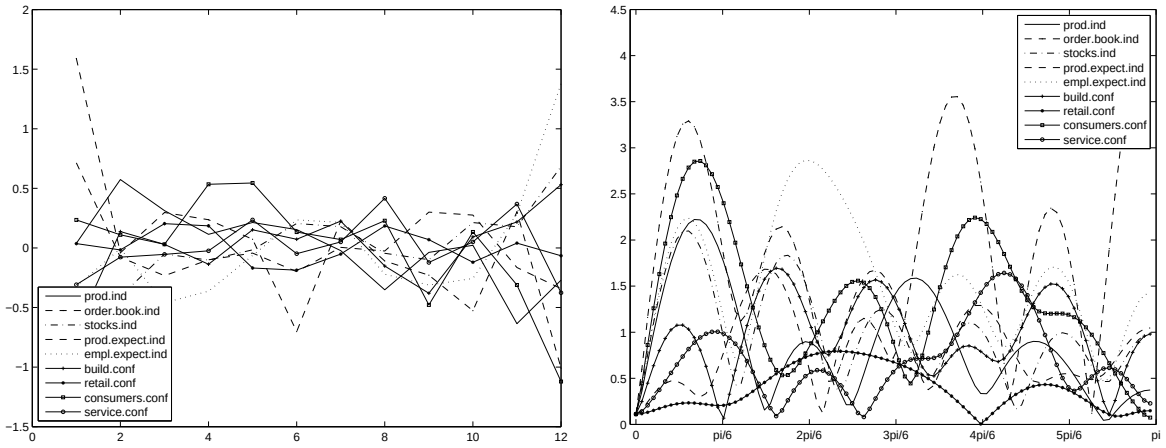

(a) (b)

Figure 3: (a) Coefficients for a 1st-order constrained lowpass filter. (b) Filter amplitudes corresponding to the coefficients in Figure 3(a).

The second order restriction imposes a vanishing phase shift at frequency zero for a targeted lead or lag, and also restricts a d.f. per input variable in an unregularized problem. This constraint is related to assuming the target variable follows the second order integrated process, in which case there are two unit roots at frequency zero and therefore both first and second order constraints would be implemented. However, the time-shift constraint can be used without the first order constraint in order to ensure the output is coincident with the target signal but not necessarily assuming that the target signal follows a second order integrated process. Therefore, such a combination of constraints goes beyond the one typically seen in the time-domain applications.

The corresponding filter coefficient and amplitude Figures 4(a) and 4(b) show that the coefficients are back to their original scale and also amplitudes look less exploded compared to the ones of the 1st order constrained filter. (Evidently higher amplitudes at the high-frequency content indicates that zero time shift at frequency zero is obtained by putting higher weight on high-frequency content which is typically the case when the

explanatory variables are lagging with respect to the target variable, which is in line with the observation from Figures 1(a) and 1(b).) Still, the second order constrain is not a panacea since the amplitudes are still erratic and since the number of degrees of freedom exceeds the number of observations for samples smaller than 9*(12-1)=99 months which is 8 years of data.

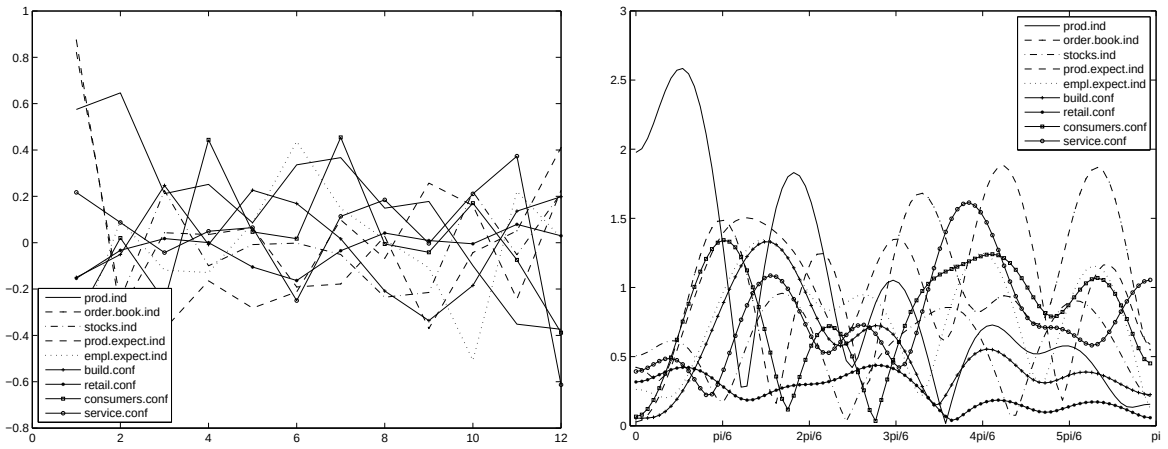

(a) (b)

Figure 4: (a) Coefficients for a 2nd-order constrained concurrent filter. (b) Filter amplitudes corresponding to the coefficients in Figure 4(b).

Turning to the new regularization features, Figures 5(a) to 5(f) show the effect of coefficient smoothness restriction of various extent corresponding to λ_s being 0.01, 0.1 and 1, which correspond to the effective degrees of freedom 66, 43 and 30, respectively.

Figures 5(a) to 5(f) show that the filter coefficients are no longer erratic; they are nice and smooth and they are getting more linear as the smoothness parameter λ_s increases. If the smoothness parameter is increased still further, the filter coefficients converge to horizontal straight lines. However, such an over-regularization is not necessary nor welcome since the considered small values of the smoothness tuning coefficient already reduces a lot of degrees of freedom and the corresponding amplitudes look much closer to those that would be expected, i.e., most of their weights concentrate in the passband $[0, \pi/6]$ and converge to zero in the stopband. Nonetheless, the filter coefficients show neither convergence to zero with higher lags, nor similarity across series.

Figures 6(a) to 6(f) show the (partial) effect of cross-sectional restriction of various extent corresponding to λ_c being 0.01, 0.1 and 1 (the rest of shrinkage parameters being zero), which correspond to the effective degrees of freedom 85, 48 and 24, respectively, which is close to what we have observed with parameter smoothness restriction.

The effects of cross-sectional restriction differ from those of parameter smoothness restriction - mild cross-sectional restriction seemingly improves the behaviour of filter coefficients and amplitudes (see Figures 6(a) and 6(b)) but further cross-sectional restriction can be harmful if applied alone (see amplitude behaviour in Figure 6(f)). Such a cross-sectional restriction analysis might help understand which series or clusters of series are different from the others. In our exercise, no series clearly stands out from the rest.

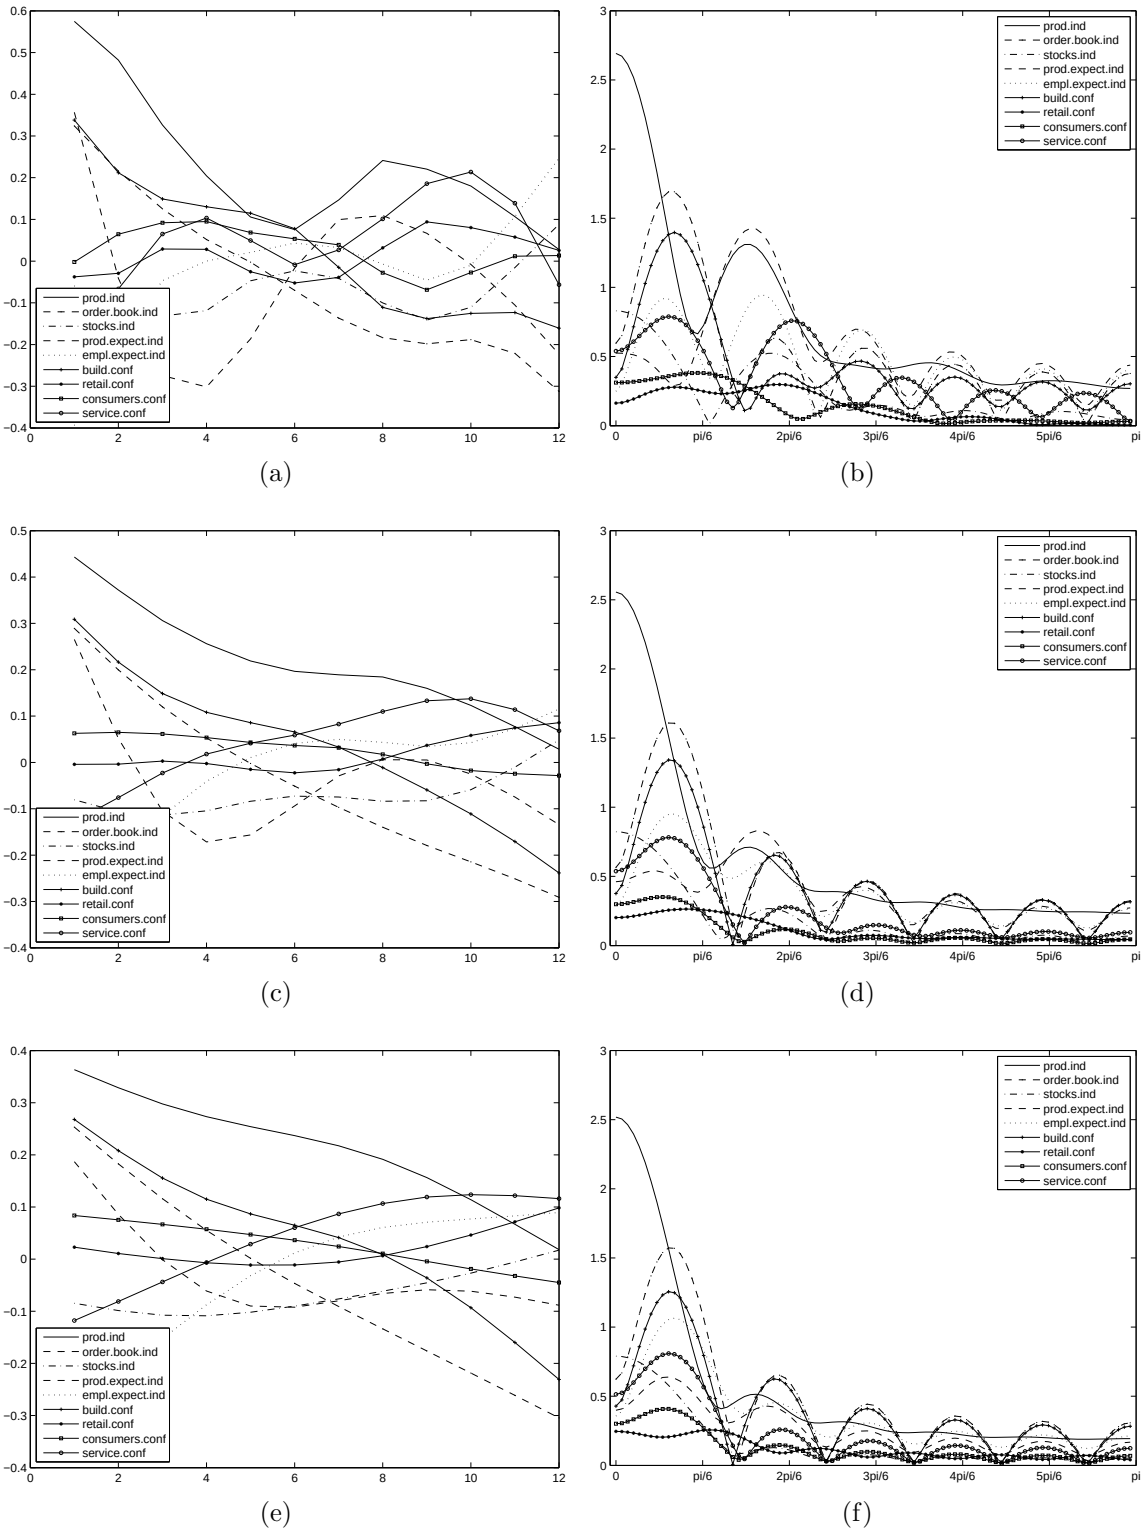

Figure 5: (a) Coefficients for an unconstrained filter if $\lambda_s = 0.01$. (b) Filter amplitudes corresponding to the coefficients in Figure 5(a). (c) Coefficients for an unconstrained filter if $\lambda_s = 0.1$. (d) Filter amplitudes corresponding to the coefficients in Figure 5(c). (e) Coefficients for an unconstrained filter if $\lambda_s = 1$. (f) Filter amplitudes corresponding to the coefficients in Figure 5(e).

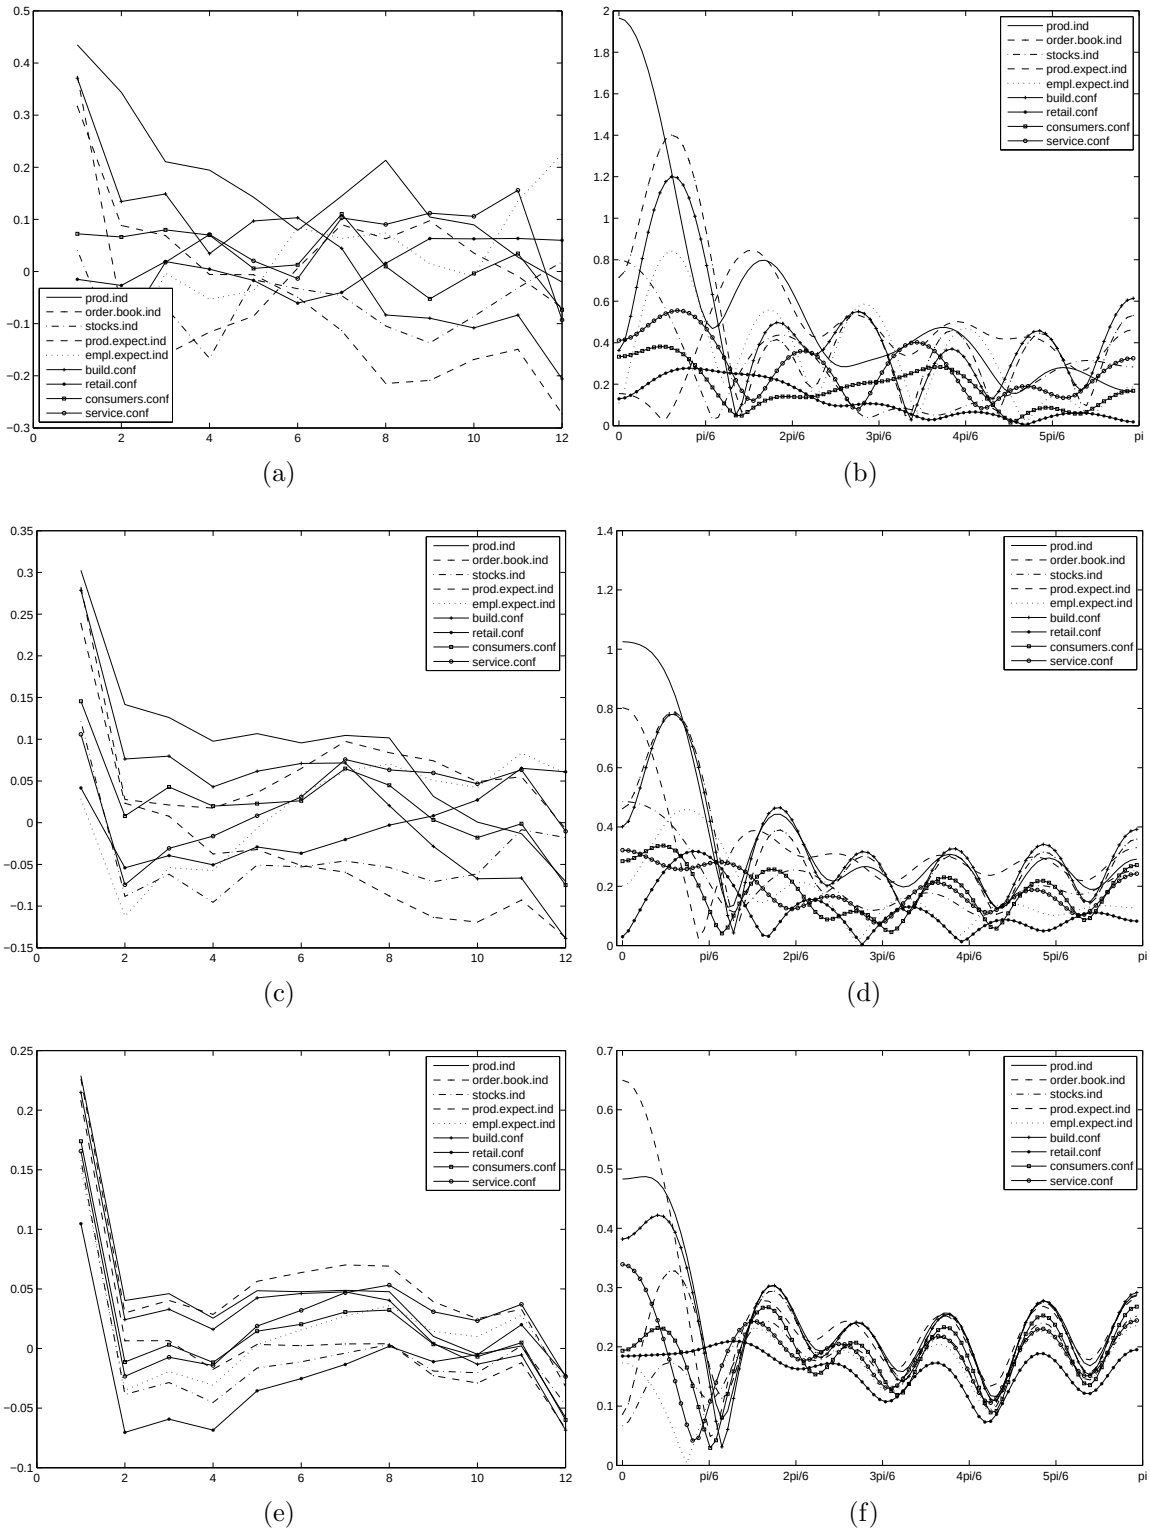

Figure 6: (a) Coefficients for an unconstrained filter if $\lambda_c = 0.01$. (b) Filter amplitudes corresponding to the coefficients in Figure 6(a). (c) Coefficients for an unconstrained filter if $\lambda_c = 0.1$. (d) Filter amplitudes corresponding to the coefficients in Figure 6(c). (e) Coefficients for an unconstrained filter if $\lambda_c = 1$. (f) Filter amplitudes corresponding to the coefficients in Figure 6(e).

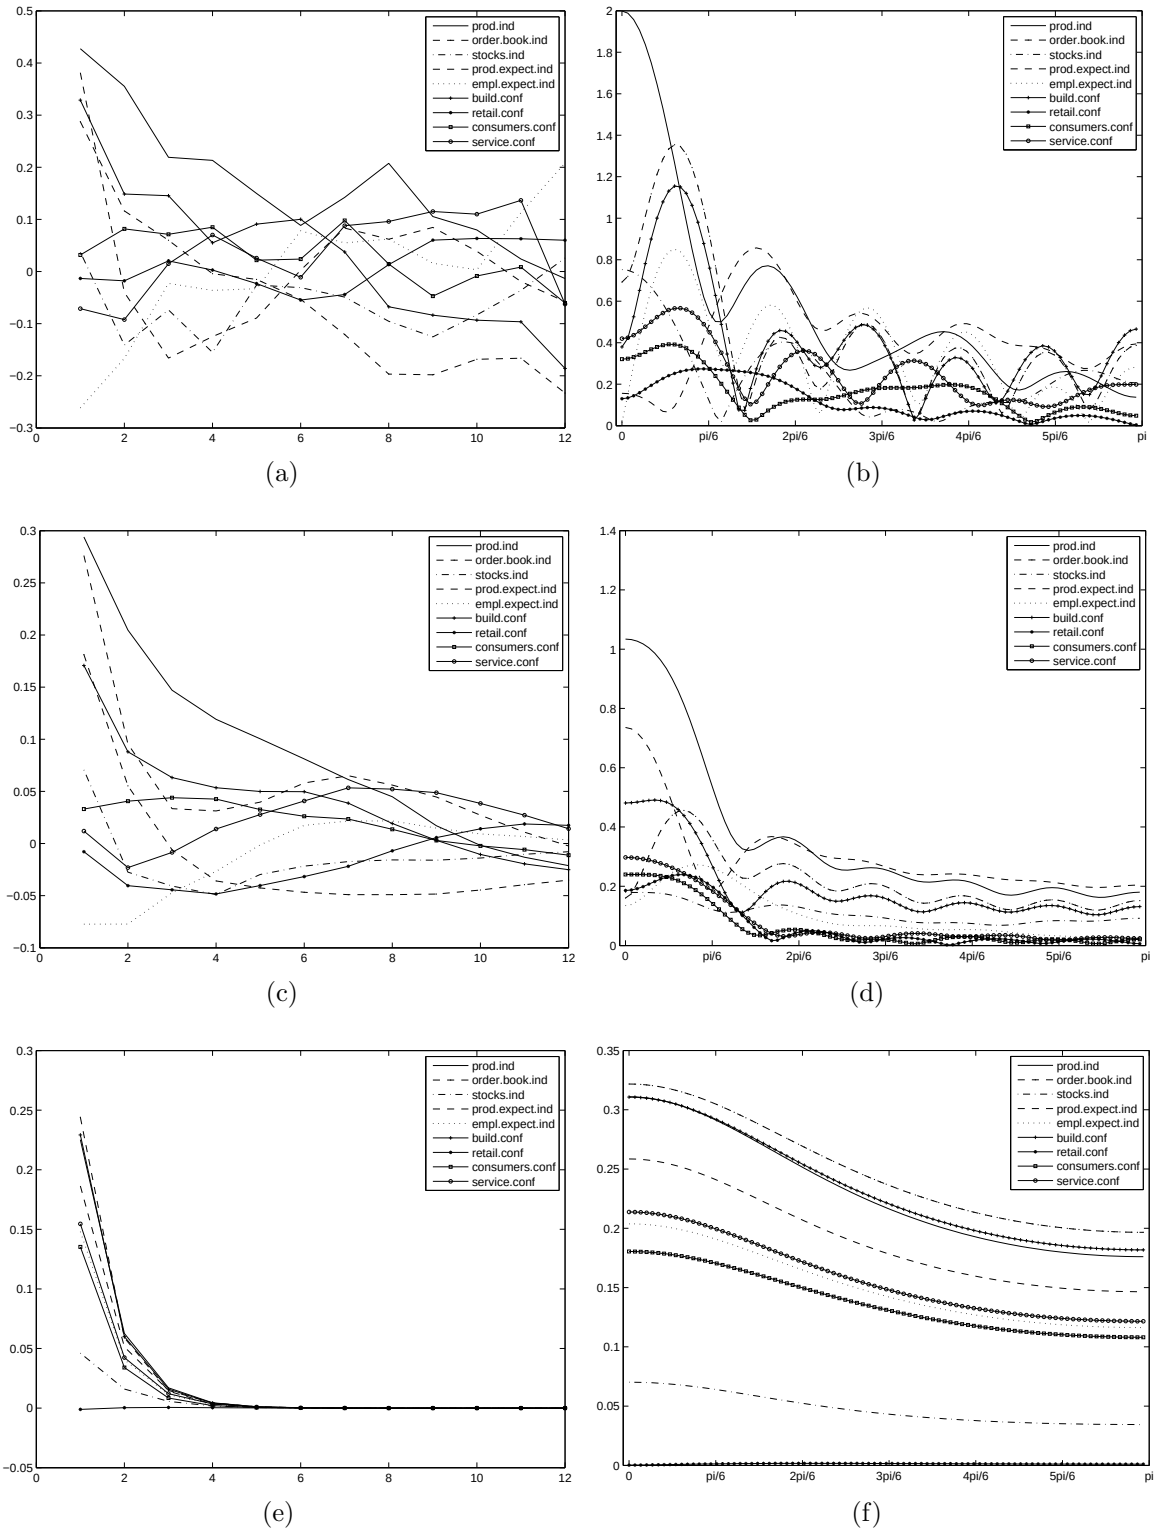

Figure 7: (a) Coefficients for an unconstrained filter if $\lambda_d = 0.01$. (b) Filter amplitudes corresponding to the coefficients in Figure 7(a). (c) Coefficients for an unconstrained filter if $\lambda_d = 0.1$. (d) Filter amplitudes corresponding to the coefficients in Figure 7(c). (e) Coefficients for an unconstrained filter if $\lambda_d = 1$. (f) Filter amplitudes corresponding to the coefficients in Figure 7(e).

As for the third regularization feature, Figures 7(a) to 7(f) show the effect of longitudinal, i.e. a lag decay restriction of various extent corresponding to λ_d being 0.01, 0.1 and 1, which correspond to the effective degrees of freedom 82, 30 and 5, respectively, which is a stronger shrinkage than what we have observed with parameter smoothness or cross-sectional restriction.

Figures 7(a) to 7(f) show that a lag decay restriction forces filter coefficients to shrink towards zero as functions of lag and that a sufficiently high shrinkage parameter yield filter coefficients to be non-zero for a small number of lags. Figure 7(f) shows that a sufficiently high longitudinal shrinkage forces filter amplitudes to shrink towards zero (see the scale of Figure 7(f)) and flatten, resembling those of an allpass filter, which is an expected behaviour since a short filter cannot discriminate between frequencies effectively.

Coefficients in Figures 7(a) and 7(c) are rather smooth which resembles the effect of parameter smoothness restriction. Also, Figures 7(a) and 7(c) show that longitudinal restriction forces filter coefficients to behave somewhat similarly across series, which reminds of the cross-sectional shrinkage. These effects might suggest that the lag decay shrinkage is the most useful of all three shrinkages. Still, the longitudinal shrinkage might conflict with e.g. parameter smoothness restriction for a sufficiently high lag decay restriction, see Figure 7(e). But, instead of using both longitudinal and parameter smoothness regularization features, one might just loosen the lag decay restriction.

The findings in this paper indeed suggest that the longitudinal shrinkage might be the most useful of the three regularization features. Moreover, this paper will use only the longitudinal and the cross-sectional shrinkages from the considered regularization 'troika' since the parameter smoothness restriction can be obtained implicitly by the former two.

Recall that setting the longitudinal shrinkage to $\lambda_d = 1$ yields only five e.d.f. which might suggest that a slight change in the sample size or in the number of explanatory series could yield close to zero e.d.f. Indeed, the estimation routine can break up if severe regularization is imposed. Therefore, a caution should be taken in empirical work so that a sufficient number of effective degrees of freedom are given to the estimation routine. Otherwise, the estimation routine will not work not because of overparameterization but because of 'underparameterization'.

Filter constraints have been found to be useful in real-time signal extraction (see e.g. Buss, 2012). Therefore, consider the effect of longitudinal shrinkage combined with 1st-order constraint or 2nd-order constraint or both 1st- and 2nd-order constraints.

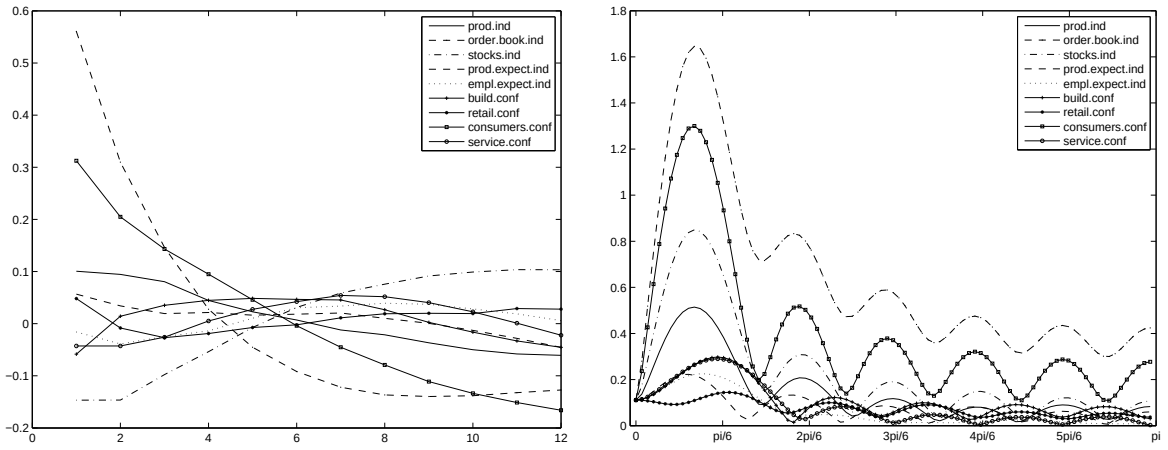

(a) (b)

Figure 8: (a) Coefficients if longitudinal regularization with $\lambda_d = 0.1$ and the 1st order constraint are implemented. (b) Filter amplitudes corresponding to the coefficients in Figure 8(a).

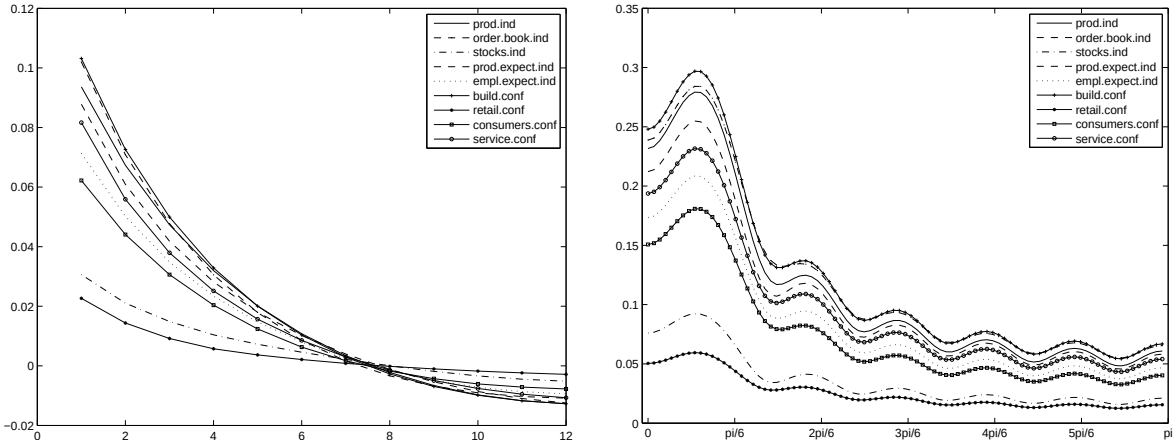

(a) (b)

Figure 9: (a) Coefficients if longitudinal regularization with $\lambda_d = 0.1$ and the 2nd order constraint are implemented. (b) Filter amplitudes corresponding to the coefficients in Figure 9(a).

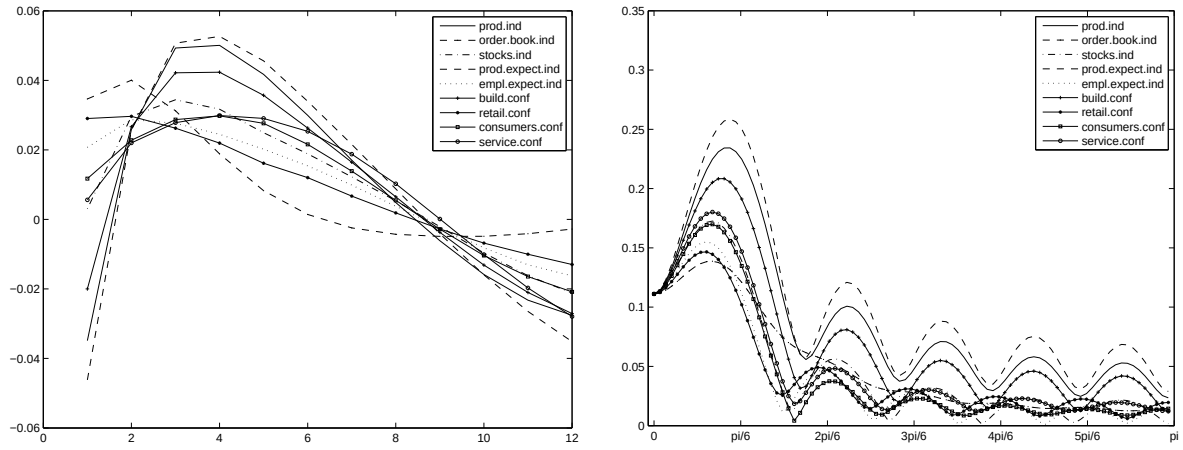

(a) (b)

Figure 10: (a) Coefficients if longitudinal regularization with $\lambda_d = 0.1$ and both the 1st- and the 2nd-order constraints are implemented. (b) Filter amplitudes corresponding to the coefficients in Figure 10(a).

Implementing the 1st-order constraint together with the longitudinal shrinkage yields similarly-behaved coefficients and amplitudes whose values at frequency zero are an inverse of the number of input variables, i.e 1/9. Amplitude values tend to diverge sharply and mostly increase for passband frequencies after which they tend to converge and decrease. The instability of the amplitudes at low frequencies might be explained by the restrictive nature of the 1st-order constraint - it forces all amplitudes to be of the same small value although the unrestricted amplitudes are somewhat dispersed around frequency zero. Also, some of the coefficients are negative at low lags which can be considered as an unwelcome effect for the dataset where each series correlates positively with the target.

The second order constraint slightly increases the dispersion of the coefficients but otherwise does not add drastic changes to the regularized filter.

Implementing both constraints simultaneously is the most restrictive case. Figures 10(a) and 10(b) show that filter coefficients behave more similarly among series than in the case of no constraints or just 1st-order constraint (notice the scale of graphs), and so the corresponding amplitudes are less dispersed than in the case of no constraints or just 1st order constraint. Still, negative coefficient values implied by the 1st-order constraint might be considered as somewhat implausible/unwanted, as well as the cause of their implausibility - the restrictive and somewhat arbitrary amplitude constraint. Therefore, if the 1st-order constraint is to be used, one should think of plausible values for amplitudes at frequency zero. Otherwise, the practitioner might be willing to use the cross-sectional shrinkage as a tool to help controlling for degrees of freedom (at least for rather homogeneous datasets), instead of using the amplitude constraint.

3.1.4 Indicator design

The chosen real-time filter design for the yearly growth rate of the euro area GDP is thus a regularized, 2nd-order constrained lowpass filter with possibly

positive longitudinal and cross-sectional shrinkages ($\lambda_d \geq 0$, $\lambda_c \geq 0$) **and no parameter smoothness restriction** ($\lambda_s = 0$).

Applying the filter on all 72 variables requires more stringent shrinkage. This is done by increasing the longitudinal shrinkage parameter to $\lambda_d = 0.2$ and the cross-sectional shrinkage parameter to $\lambda_c = 5$. The rationale for the chosen shrinkage parameters is the following. The previous subsection shows that the longitudinal shrinkage is more aggressive than the cross-sectional one. Thus, the longitudinal shrinkage parameter cannot be set too high since the filter will be effectively too short (filter coefficients will be zero for larger lags). Therefore, in order not to reduce the filter length to inappropriate value (since a too short filter cannot discriminate between frequencies effectively), the rest of d.f. reduction can be achieved by the cross-sectional shrinkage. Increasing the cross-sectional shrinkage parameter to infinity yields filters for all variables to converge, and d.f. to reduce. Thus, increasing the extent of the cross-sectional shrinkage does not yield fatal outcome and thus is less harmful than increasing the extent of the longitudinal shrinkage. This consideration can be considered as satisfactory at least for homogeneous enough datasets, which is the case for the euro area dataset because it is dominated by a large number of survey data. Indeed, for the EA dataset, increasing the cross-sectional shrinkage parameter to, say, $\lambda_c = 20$, would yield less d.f. but hardly any difference in filter output. Yet, there is a good reason to allow some d.f. for the filter if the input dataset is heterogeneous.

Filter coefficients and amplitudes are shown in Figures 11(a) and 11(b), respectively. Figure labels are removed due to over-cluttering.

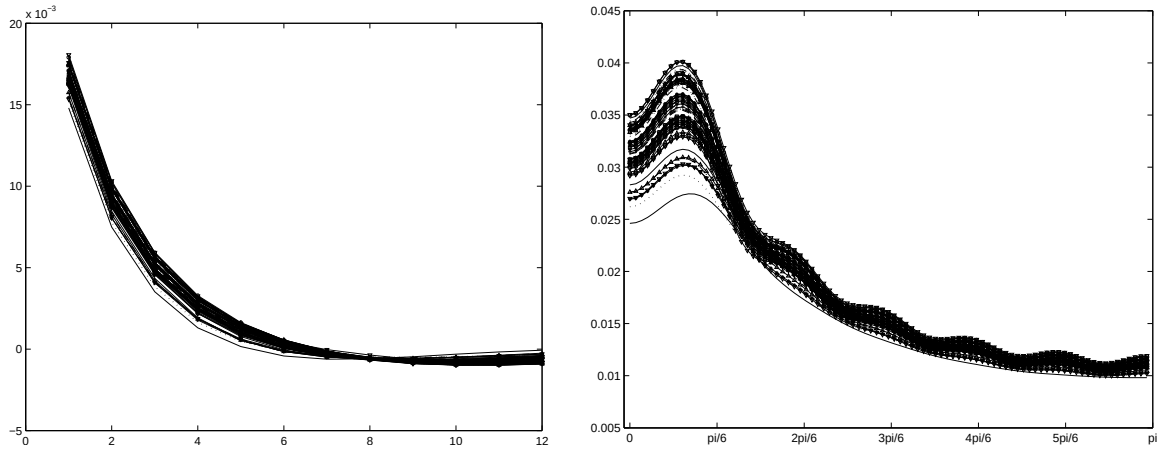

(a) (b)

Figure 11: (a) Coefficients for a concurrent 72-variables filter with 2nd-order constraint, $h = 0$, $\lambda_d = 0.2$, $\lambda_c = 5$. (b) Filter amplitudes corresponding to the coefficients in Figure 11(a).

The resulting filter coefficients and amplitudes look plausible. The coefficients for small lags are positive and decay smoothly to zero with a higher lag order. Filters with coefficients that do not shrink to zero at higher lag orders can be argued to be suboptimal/incomplete. The amplitudes also look plausible - there is some d.f. such that they are not the same for all variables but still they are close to each other and have the

most weight in the passband, and decay towards zero in stopband.

The filter's simulated real-time output for the last ten years is shown in Figure 12 along with another established indicator Eurocoin (Altissimo et al. 2010), the latter being transformed to yearly growth rates. The particular parameter setting results in about three e.d.f. on average over the whole sample.

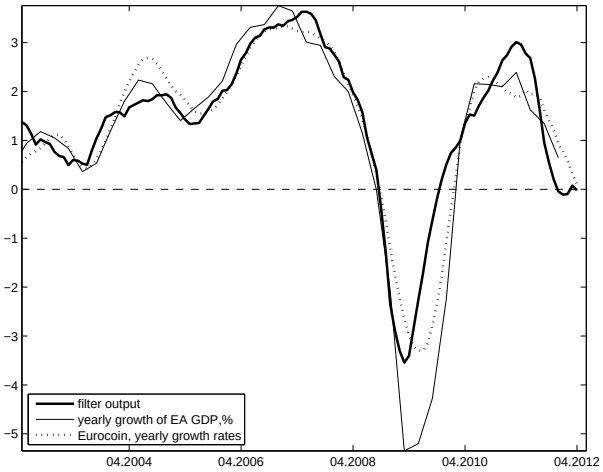

Figure 12: Output of a regularized 72-variable filter with $h = 0$, $\lambda_d = 0.2$, $\lambda_c = 5$ tracking a trendcycle in yoy EA GDP versus Eurocoin transformed to represent yearly growth rates.

Figure 12 shows that the filter output precedes Eurocoin on several occasions, and that Eurocoin is actually lagging w.r.t. GDP growth in several episodes. 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, dynamic correlation between an indicator and the GDP is used. The peak correlation between Eurocoin and GDP is found to be at a zero lag w.r.t. GDP, while the second highest correlation being at a one month lag w.r.t. GDP. For the output of RMDFA as in Figure 12, the peak correlation is at one month lead w.r.t. GDP, and the second highest correlation located at a two months lead w.r.t. GDP (Table 1 in Appendix).

Forecasting. This paragraph shows that the regularized filter can be used not only for a concurrent signal extraction but also for forecasting. Figures 13(a) and 13(b) show coefficients and amplitudes for the filter targeting a lead of three months with respect to the target signal. This is done by setting $h = -3$ in formulas (18), (29), (30), and (31). The rest of filter parameters are left unchanged, i.e., $\lambda_d = 0.2$ and $\lambda_c = 5$. This is an example of a direct forecasting, as opposed to iterated forecasting. Figures 13(a) and 13(b) show that filter coefficients and amplitudes are slightly more dispersed than in the coincident case. Repeating the exercise with an increased target lead of six months ($h = -6$) (with other filter parameters unchanged), yields filter coefficients and amplitudes as plotted in Figures 14(a) and 14(b), which show even more dispersed coefficients and amplitudes.

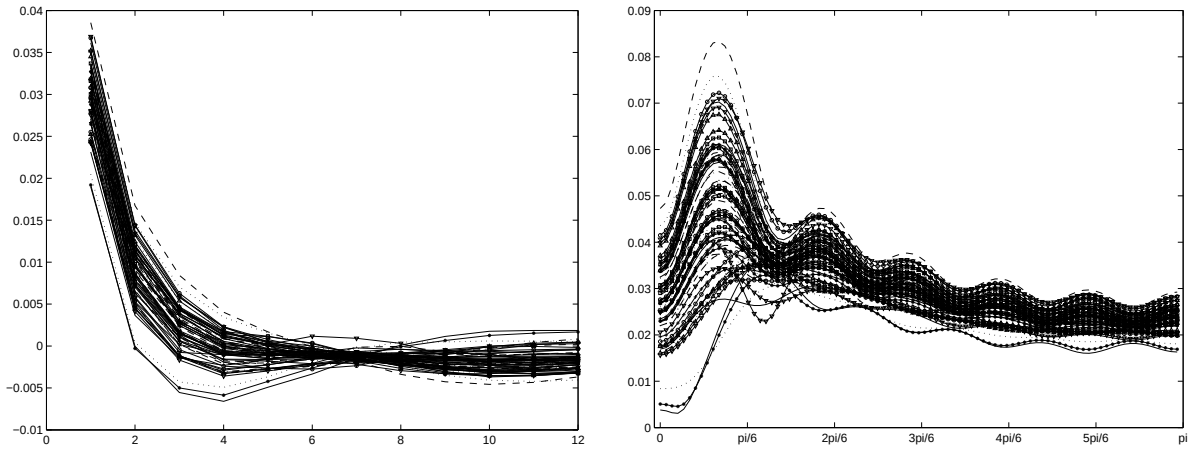

(a) (b)

Figure 13: (a) Coefficients for a 72-variables filter with 2nd-order constraint, $h = -3$, $\lambda_d = 0.2$, $\lambda_c = 5$. (b) Filter amplitudes corresponding to the coefficients in Figure 13(a).

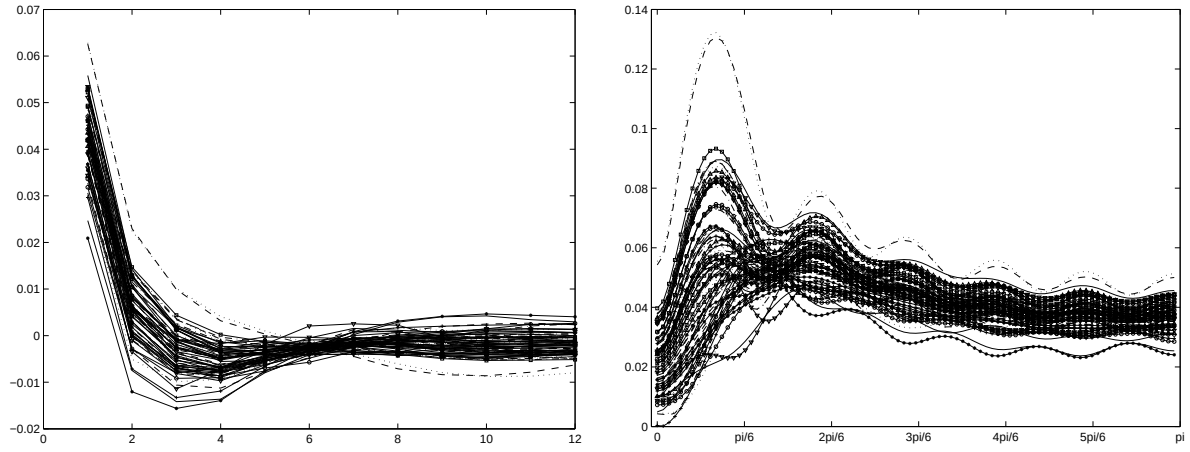

(a) (b)

Figure 14: (a) Coefficients for a 72-variables filter with 2nd-order constraint, $h = -6$, $\lambda_d = 0.2$, $\lambda_c = 5$. (b) Filter amplitudes corresponding to the coefficients in Figure 14(a).

The corresponding effective degrees of freedom are 16 (for three-months lead) and 23 (for six-months lead), as opposed to five e.d.f. for the concurrent filter. The increase of e.d.f. with the targeted lead can be explained intuitively by the fact that the filter has more freedom to chose which series will have what weight at what lead/lag in order to achieve the desired outcome. The longer way to go, the more possible ways can be chosen in order to get to the predestined place. A practitioner might set a more stringent shrinkage with higher targeted lead in order to achieve the desired degrees of freedom but it might be argued that it is intuitively unappealing to do so, since filter should be free enough to differentiate between series when it comes to targeting high leads.

The resulting real-time outputs of filters targeting three- and six-months leads are shown in Figures 15(a) and 15(b), respectively, together with Eurocoin.

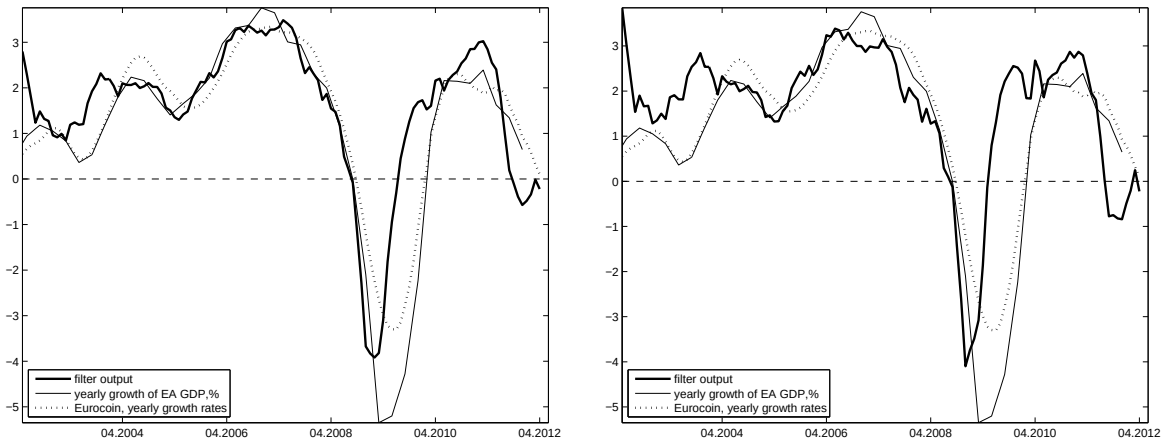

(a) (b)

Figure 15: (a) Filter output corresponding to filter coefficients in Figure 13(a) (targeting a 3-months lead) versus Eurocoin (yearly growth rates). (b) Filter output corresponding to filter coefficients in Figure 14(a) (targeting a 6-months lead) versus Eurocoin (yearly growth rates).

Figures 15(a) and 15(b) show that the resulting lead of the filter output is moderate but existent, with biggest noticeable gains in signalling recovery during the 2009-recession and the downward movement in the 2012-downturn. The level fit worsens with a higher targeted lead but this is an expected result in any forecasting exercise. For the filter output in Figure 15(a), the peak correlation is at a three months lead w.r.t. GDP, while the second highest correlation being located at a four months lead w.r.t. GDP. For the filter output in Figure 15(b), the peak correlation is located at a five months lead w.r.t. GDP, and the second highest correlation being at a six months lead w.r.t. GDP (Table 1 in Appendix).

Having created the filter design for tracking a trendcycle in yearly growth of EA GDP, we now turn to designing filter for tracking a trendcycle in quarterly growth of EA GDP.

3.2 Tracking trendcycle in quarterly growth of GDP

3.2.1 Target and data

The filter is set to target an ideal lowpass of quarterly growth of real GDP with cut-off wave length 12 months. The GDP data are linearly interpolated to monthly frequency, logged and quarterly differenced. A full list of data transformations is presented in Appendix.

3.2.2 Indicator design

There are two main differences of this design w.r.t. the yearly growth design. First, monthly differenced data are more volatile than the yearly differenced ones. Thus, a smooth signal extraction requires more noise suppression/tighter regularization. Second, the main explanatory variables are business and consumer survey data, since they are published with almost no delay and have been found to correlate well with the GDP. In

the previous subsection, it was shown that undifferenced survey data are about coincident with yearly growth of GDP. Thus, undifferenced survey data are lagging w.r.t. quarterly growth of GDP. Therefore, forecasting ($h < 0$) should be involved in order to get a coincident quarterly growth signal. (Otherwise, a practitioner could difference survey data, but regularly differenced survey data overshoot after the great recession and, strictly speaking, are over-differenced since undifferenced survey data are not integrated.).

Given the above considerations, we will show the results of two different specifications - with and without the amplitude constraint. More noise suppression can be accomplished with a tighter shrinkage, specifically by raising lag decay and cross-sectional shrinkage parameters. However, it was argued in the previous subsection that tight cross-sectional shrinkage might be suboptimal if forecasting is involved. Therefore, amplitude constraint might be used as an additional constraint that reduces degrees of freedom, to which we now turn.

Filter with an amplitude constraint. Amplitude constraint can help contain the filter output on the right level 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, given that the data set can be heterogeneous, the value of amplitude constraint at frequency zero equal for all series might be suboptimal. Therefore, we here differentiate that value to be proportional to the in-sample correlation of explanatory series with the GDP (though the result is close to what would be obtained with equal weights). The lag decay parameter has been increased to $\lambda_d = 0.4$ and the cross-sectional shrinkage parameter has been decreased to $\lambda_c = 1$. This setting gives about three e.d.f., thus more cross-sectional shrinkage is unnecessary. Filter coefficients and amplitudes are plotted in Figures 16(a) and 16(b).

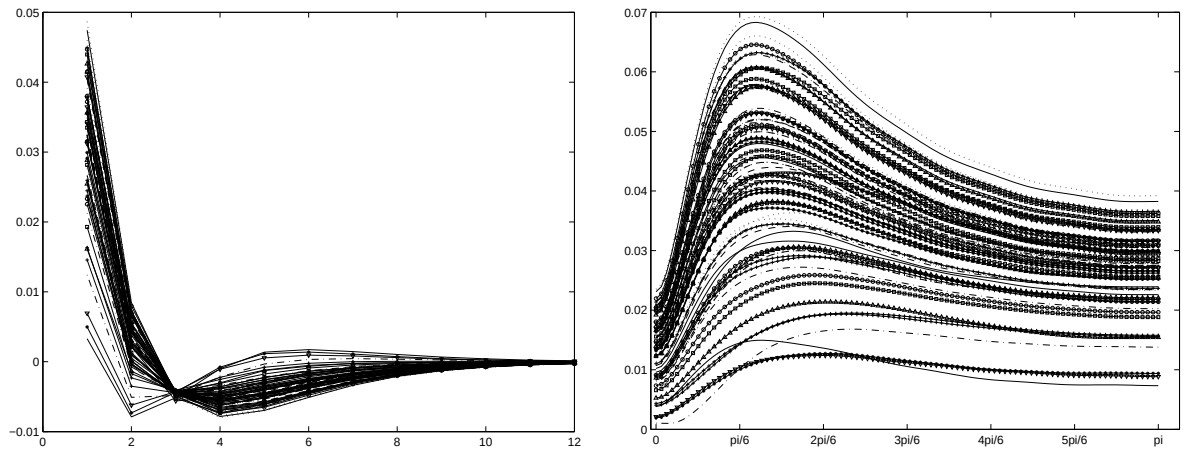

(a) (b)

Figure 16: (a) Coefficients for a 72-variables filter with both 1st- and 2nd-order constraints, $h = -6$, $\lambda_d = 0.4$, $\lambda_c = 1$. (b) Filter amplitudes corresponding to the coefficients in Figure 16(a).

The resulting real-time filter output is plotted in Figure 17 along with Eurocoin.

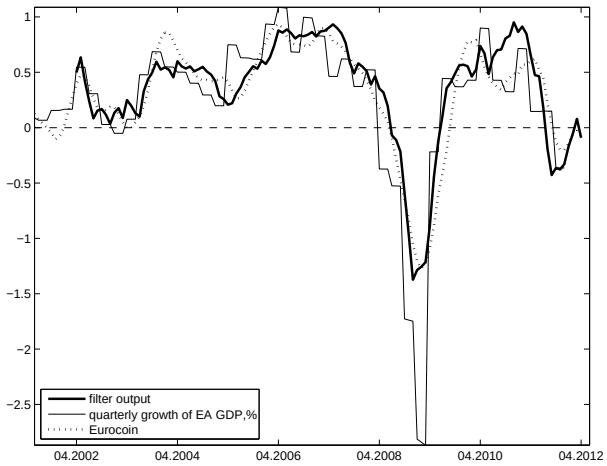

Figure 17: Filter output corresponding to filter coefficients in Figure 16(a) versus Eurocoin.

Figure 17 shows that the filter output tracks the level of the target well and precedes Eurocoin on several occasions. The peak correlation of Eurocoin with GDP is located at a two months lag w.r.t. GDP, and the second highest correlation being located at a one month lag w.r.t. GDP. For the RMDFA output in Figure 17, 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 (Table 1 in Appendix).

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.

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 1st-order constraint means more degrees of freedom, therefore shrinkage should be tightened by increasing the cross-sectional shrinkage parameter back to $\lambda_c = 5$. This setting gives about 8 e.d.f. Filter coefficients and amplitudes are plotted in Figures 18(a) and 18(b).

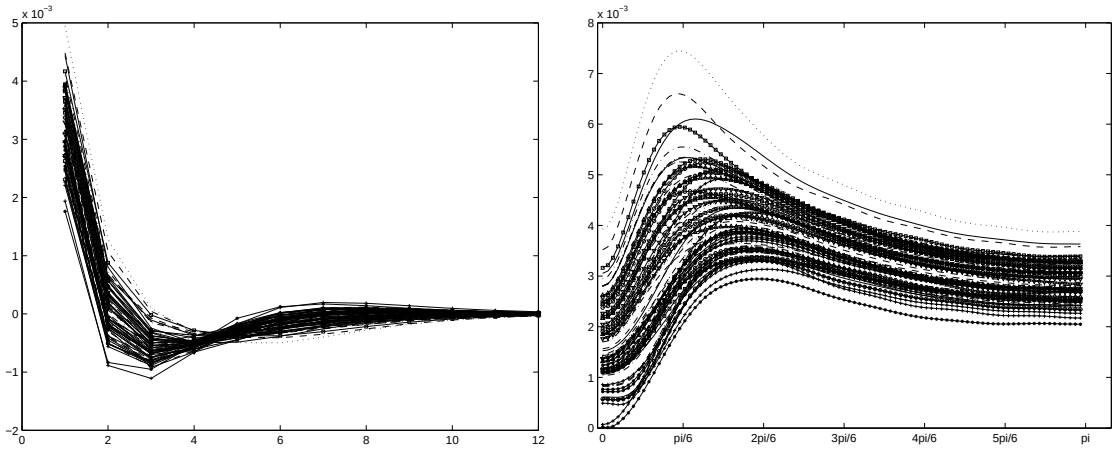

(a) (b)

Figure 18: (a) Coefficients for a 72-variables filter with 2nd-order constraint, $h = -3$, $\lambda_d = 0.4$, $\lambda_c = 5$. (b) Filter amplitudes corresponding to the coefficients in Figure 18(a).

Additional noise suppression can be achieved by suppressing amplitudes in the stopband with positive *expw* parameter, see expression (8). Particularly, noise suppression parameter is set to $expw = 0.5$, which is a standard value across applications, see Buss (2012) for a similar application. Since *expw* is not among regularization parameters, it counteracts to some extent to the regularization such that the e.d.f. increase to about 12. Filter coefficients and amplitudes are plotted in Figures 19(a) and 19(b).

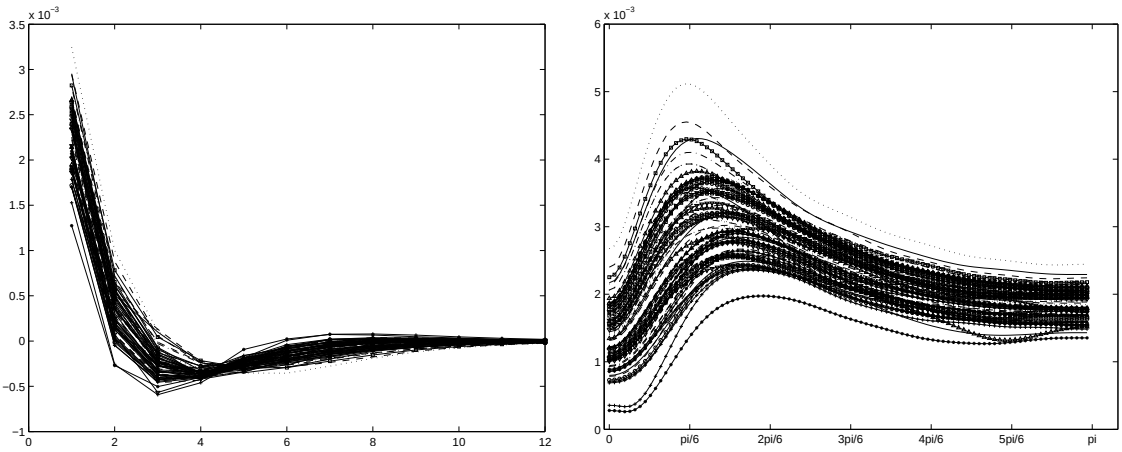

(a) (b)

Figure 19: (a) Coefficients for a 72-variables filter with 2nd-order constraint, $h = -3$, $\lambda_d = 0.4$, $\lambda_c = 5$, $expw = 0.5$. (b) Filter amplitudes corresponding to the coefficients in Figure 19(a).

The differences between the two cases are small but evident - an additional noise suppression in the stopband slightly reduces the amplitude dispersion and lowers their weights on higher frequencies. The result is a slightly slower but smoother filter output, see Figures 20(a) and 20(b) for without noise suppression ($expw = 0$) and with moderate

noise suppression ($expw = 0.5$), respectively. In both cases, the peak correlation 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 (Table 1 in Appendix).

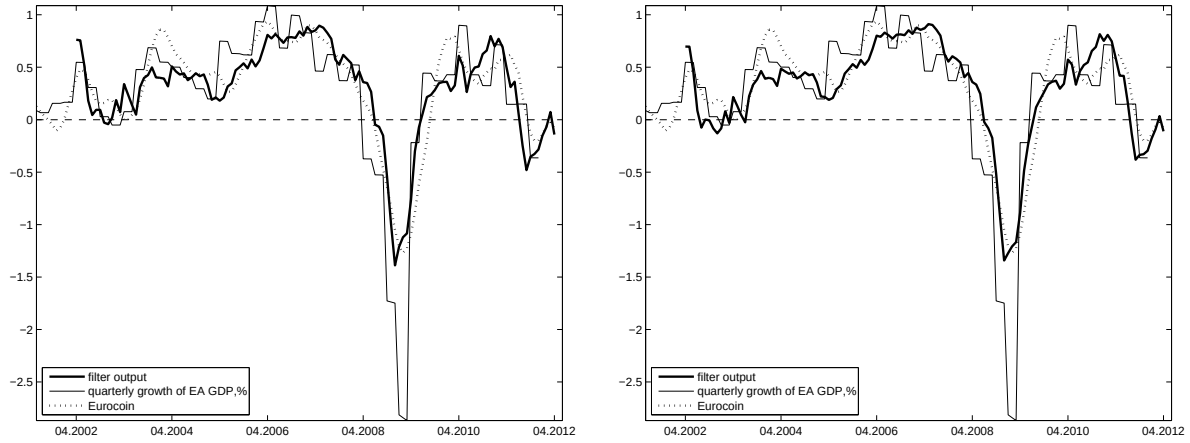

(a) (b)

Figure 20: (a) Filter output corresponding to filter coefficients in Figure 18(a) versus Eurocoin. (b) Filter output corresponding to filter coefficients in Figure 19(a) versus Eurocoin.

Increasing targeted lead to $h = -6$ and noise suppression to $expw = 1$, yields filter coefficients and amplitudes shown in Figures 21(a) and 21(b), respectively, and the real-time filter output in Figure 22.

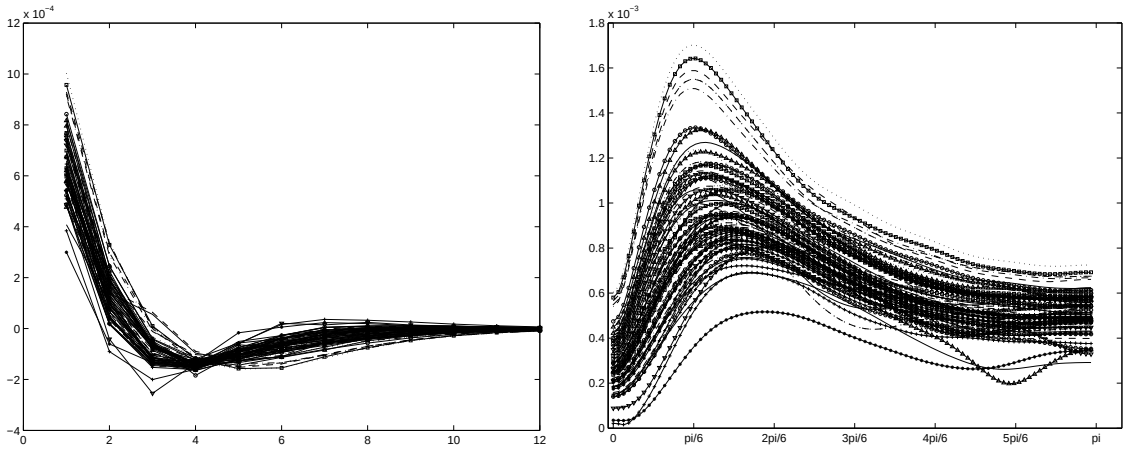

(a) (b)

Figure 21: (a) Coefficients for a 72-variables filter with 2nd-order constraint, $h = -6$, $\lambda_d = 0.4$, $\lambda_c = 5$, $expw = 1$. (b) Filter amplitudes corresponding to the coefficients in Figure 21(a).

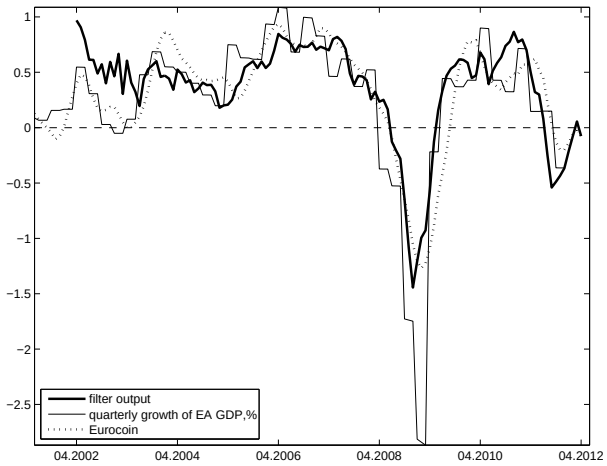

Figure 22: Filter output corresponding to filter coefficients in Figure 21(a) versus Eurocoin.

The in-sample size at the beginning of estimation is evidently too short but soon after the filter output stabilizes and outperforms Eurocoin w.r.t. timeliness on several occasions but clearly after 2009. The peak correlation with GDP is located at a zero months lag w.r.t. GDP, with the second highest correlation being at a one month lead w.r.t. GDP (Table 1 in Appendix).

The following section checks the filter performance on a less homogeneous Latvia's dataset.

4 Robustness check on a less homogeneous Latvia's dataset

Latvia's dataset contains 40 explanatory variables - 30 business and consumer variables for Latvia, EA, Estonia and Lithuania, 3 industrial production indices for EA, registered unemployment, job vacancies, monetary aggregates M1 and M3, currency in circulation, volume index of exports of goods, and budget income variable. Many relevant variables are not included due to their short sample size or strong seasonality; those variables could be seasonally adjusted and extrapolated with an expectation-maximization algorithm but this is not done for simplicity. In fact, some of the included variables (volume of exports, budget income, monetary aggregates, employment variables) are seasonal but seasonality is not deal with, neither checking for outliers is performed. Several variables were found to have low correlation with the target variable but none was excluded. Integrated series were made non-integrated by suitable transformations. Appendix lists the data and transformations.

There is a difference between applying filter on EA or Latvia's datasets. Particularly, Latvia's survey data have been found to yield better explanatory power for quarterly growth of Latvia's GDP when they are regularly differenced. Thus, all data, including survey data, are regularly differenced when the target is quarterly growth of GDP. However, differenced survey data overshoot after the great recession. Therefore, in order to lessen the impact of the differenced survey data on the outcome, Latvia's GDP series is

included in the set of explanatory variables, as well. In order to produce close to real-time performance, flash GDP values (released about 45 days after the reference period) are not used, as well as the first releases (published about 65 days after the reference period) are dropped off. Thus, the GDP series lags survey data by seven months. The existence of GDP in the set of explanatory variables makes Latvia's dataset particularly less homogeneous than the one for the euro area.

No filter constraints are imposed. Regarding the regularization parameters, the lag decay parameter is set to the value which is used in the EA application, $\lambda_d = 0.2$; the cross-sectional shrinkage parameter, however, is set to be much lower, $\lambda_c = 0.2$, which can be explained by the more heterogeneous dataset, particularly, the existence of GDP series in the dataset. If the cross-sectional shrinkage parameter is increased, the filter coefficients on GDP series are shrunk towards the rest of filter parameters; since the latter are dominated by survey data that overshoot after the great recession, this means that the increase of the cross-sectional shrinkage leads to a more timely extracted signal but also that it overshoots more after the great recession. There is also another reason for keeping the cross-sectional shrinkage parameter low, particularly - heterogeneous data might contain irrelevant variables (the variables were not subject to scrutinized pre-screening except for changing the signs of negatively correlated variables) and thus forcing all filters to have the same coefficients might be considered to be suboptimal.

Given the absence of filter constraints and the small values of regularization parameters, the effective degrees of freedom are quite large (60) compared to the EA application in the previous section, but which is still much less than the number of estimated filter coefficients ($41 \times 12 = 492$). Filter coefficients and amplitudes are shown in Figures 23(a) and 23(b).

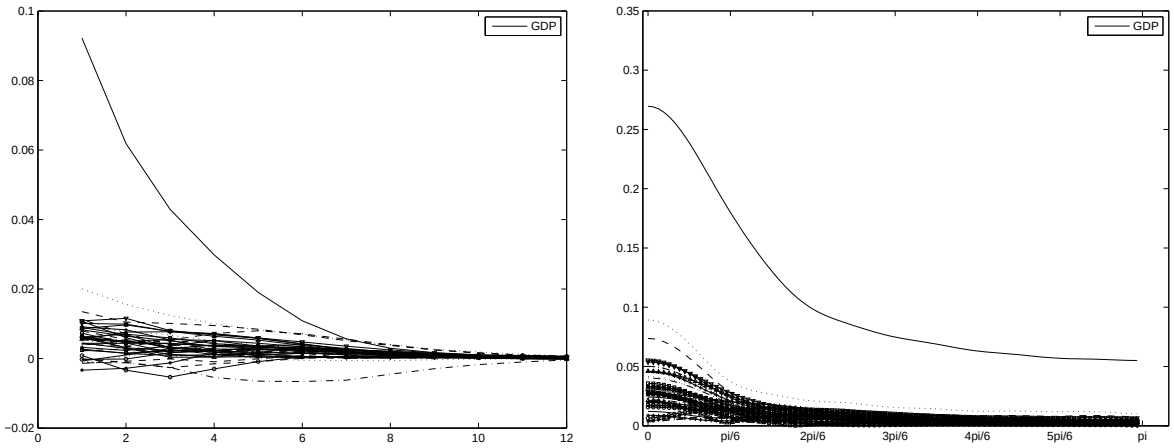

(a) (b)

Figure 23: (a) Coefficients for a 41-variables filter, $\lambda_d = 0.2$, $\lambda_c = 0.2$. (b) Filter amplitudes corresponding to the coefficients in Figure 23(a).

One series clearly stands out and it is the GDP series, which has a higher weight than the rest. There are also a couple of series with practically zero coefficients and amplitudes - these can be considered to be irrelevant series, but they are not excluded from the dataset both for simplicity and for the fact that this is a high-dimensional

filtering exercise which by its name suggests that there might be irrelevant variables which should not be necessarily excluded in order to get a decent outcome.

The resulting real-time filter output for the last 10 years is shown in Figure 24, along with quarterly growth of Latvia's GDP and pseudo real-time values of Latcoin indicator (see Benkovskis (2010), though its design has been slightly changed since then) which is a real-time indicator for Latvia's GDP with a Eurocoin-type methodology (see Altissimo et al (2010)).

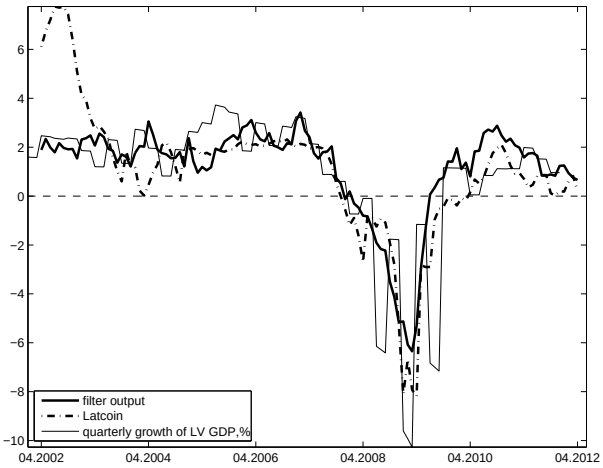

Figure 24: Filter output corresponding to filter coefficients in Figure 23(a) versus Latcoin.

Figure 24 shows that the filter output is about coincident with Latcoin but is smoother during the great recession period and slightly faster during the recovery phase. It also appears to be more robust for smaller samples although contains more parameters. The peak correlations indicate that both indicators are about coincident with GDP (Table 1 in Appendix).

5 Conclusions

Nowadays, information is abundant. Statistical tools are being developed that are suitable to process large information for a particular problem at hand. This paper considers regularized multivariate direct filter approach (Wildi, 2012) as a tool for signal extraction and forecasting using high-dimensional datasets. The paper studies filter's properties by tracking the medium-to-long-run component of EA GDP growth using up to 72-variable filter. It is shown that the filter can be successfully applied on high-dimensional datasets. The particular application replicates the behavior of an established Eurocoin indicator, as well as produces more timely indicators.

As a robustness check, another application considers a more heterogeneous Latvia's dataset. It is shown that a more heterogeneous dataset requires less stringent cross-sectional shrinkage but that moderate longitudinal and cross-sectional shrinkage on a 41-variable dataset can yield a satisfactory outcome.

Overall, the RMDFA is found to be a promising tool for both signal extraction and forecasting using high-dimensional datasets; particularly, it might become a decent competitor to such established methods as dynamic factor methodology.

A possible downside of the RMDFA compared to the factor methodology is its many hyperparameters a user has to choose. The choice of hyperparameters is problem-specific. The existence of hyperparameters should not be the most critical aspect of the method. There are many popular methods widely used in applied econometrics that involve hyperparameters, e.g. the Bayesian approach. It might be helpful to endogenize the choice of some hyperparameters at least for some common problems but this aspect is left for future research.

Appendix

Indicators	Dynamic correlation at lag:									
	-3	-2	-1	0	1	2	3	4	5	6
Eurocoin, yoy	0.921	0.963	0.985*	0.986**	0.967	0.930	0.874	0.804	0.722	0.631
RMDFA in Fig. 12	0.775	0.840	0.893	0.930	0.949**	0.946*	0.921	0.876	0.812	0.736
RMDFA in Fig. 15(a)	0.551	0.642	0.730	0.807	0.872	0.918	0.941**	0.938*	0.912	0.867
RMDFA in Fig. 15(b)	0.310	0.408	0.511	0.610	0.704	0.787	0.852	0.893	0.907**	0.896*
Eurocoin	0.852	0.882**	0.879*	0.849	0.792	0.706				
RMDFA in Fig. 17	0.774	0.846	0.883**	0.870*	0.819	0.736				
RMDFA in Fig. 20(a)	0.751	0.824	0.869**	0.862*	0.817	0.740				
RMDFA in Fig. 20(b)	0.790	0.849	0.876**	0.856*	0.799	0.709				
RMDFA in Fig. 22	0.680	0.780	0.849	0.874**	0.858*	0.810				
LATCOIN in Fig. 24	0.685	0.716	0.747**	0.744*	0.713	0.668				
RMDFA in Fig. 24	0.664	0.723	0.769**	0.762*	0.755	0.742				

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

Table 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 U.S. share price index, U.S., NSA	Eurostat	Δ_{12} log	Δ log
The EA share price index, EA, NSA	Eurostat	Δ_{12} log	Δ log

Table 2: The EA dataset.

Variable	Source	transformation
Real gross domestic product, chain-linked, LV, SA	Eurostat	Δ log,lin.interp.
Production trend observed in recent month (industry), LV, SA	DG Ecfin	Δ
Assessment of order-book levels (industry), LV, SA	DG Ecfin	Δ
Assessment of export order-book levels (industry), LV, SA	DG Ecfin	Δ
Assessment of stocks of finished products (industry), LV, SA	DG Ecfin	Δ
Production expectations for the months ahead (industry), LV, SA	DG Ecfin	Δ
Selling price expectations for the months ahead (industry), LV, SA	DG Ecfin	Δ
Employment expectations for the months ahead (industry), LV, SA	DG Ecfin	Δ
Consumer confidence indicator, LV, SA	DG Ecfin	Δ
Confidence indicator in retail, LV, SA	DG Ecfin	Δ
Business activity (sales) development over the past 3 months (retail), LV, SA	DG Ecfin	Δ
Volume of stock currently hold (retail), LV, SA	DG Ecfin	Δ
Orders expectations over the next 3 months (retail), LV, SA	DG Ecfin	Δ
Business activity expectations over the next 3 months (retail), LV, SA	DG Ecfin	Δ
Employment expectations over the next 3 months (retail), LV, SA	DG Ecfin	Δ
Confidence indicator in construction, LV, SA	DG Ecfin	Δ
Building activity development over the past 3 months (construction), LV, SA	DG Ecfin	Δ
Confidence indicator in industry, EU, SA	DG Ecfin	Δ
Consumer confidence indicator, EU, SA	DG Ecfin	Δ
Confidence indicator in retail, EU, SA	DG Ecfin	Δ
Confidence indicator in construction, EU, SA	DG Ecfin	Δ
Economic sentiment indicator, EU, SA	DG Ecfin	Δ
Confidence indicator in industry, EE, SA	DG Ecfin	Δ
Consumer confidence indicator, EE, SA	DG Ecfin	Δ
Confidence indicator in retail, EE, SA	DG Ecfin	Δ
Confidence indicator in construction, EE, SA	DG Ecfin	Δ
Economic sentiment indicator, EE, SA	DG Ecfin	Δ
Confidence indicator in industry, LT, SA	DG Ecfin	Δ
Confidence indicator in retail, LT, SA	DG Ecfin	Δ
Confidence indicator in construction, LT, SA	DG Ecfin	Δ
Economic sentiment indicator, LT, SA	DG Ecfin	Δ
Industrial production index B-D;F, euro area, SA	Eurostat	Δ log
Industrial production index C, euro area, SA	Eurostat	Δ log
Industrial production index D, euro area, SA	Eurostat	Δ log
Registered unemployment, LV, NSA	CSB	Δ
Job vacancies, LV, NSA	CSB	Δ
Monetary aggregate M1, LV, NSA	Bank of Latvia	Δ log
Monetary aggregate M3, LV, NSA	Bank of Latvia	Δ log
Currency in circulation (average), LV, NSA	Bank of Latvia	Δ log
Volume index of exports of goods, LV, NSA	CSB	Δ log
Budget income, LV, NSA	State revenue service	Δ log

Table 3: The LV dataset.

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.

[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.

[3] Benkovskis K. 2010. LATCOIN: determining medium to long-run tendencies of economic growth in Latvia in real time, *Baltic Journal of Economics*, Baltic International Centre for Economic Policy Studies, vol. 10(2), 27-48.

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

[5] Buss G. 2012. A new real-time indicator for the Euro Area GDP, Working Papers 2012/02, Latvijas Banka.

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

[7] Christiano LJ, Fitzgerald TJ. 2003. The band pass filter, *International Economic Review*, vol. 44(2), 435-465.

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

[9] Efron B, Hastie T, Johnstone I, Tibshirani R. 2004. Least angle regression, *The Annals of Statistics*, vol. 32(2), 407-499.

[10] 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.

[11] 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.

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

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

[14] Hoerl AE, Kennard RW. 1970. Ridge regression: applications to nonorthogonal problems, *Technometrics*, vol. 12(1), 69-82.

[15] King RG, Rebelo ST. 1993. Low frequency filtering and real business cycles, *Journal of Economic Dynamics and Control*, vol. 17(1-2), 207-231.

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

[17] 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.

[18] Stock JH, Watson MW. 2002. Macroeconomic forecasting using diffusion indexes, *Journal of Business & Economic Statistics*, vol. 20(2), 147-162.

[19] Tibshirani R. 1996. Regression shrinkage and selection via the lasso, *Journal of the Royal Statistical Society, Series B (Methodological)*, vol. 58(1), 267-288.

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

[21] Valle e Azevedo J. 2011. A multivariate band-pass filter for economic time series, *Journal of the Royal Statistical Society Series C*, vol. 60(1), 1-30.

[22] 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.

[23] 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.

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