RIGA TECHNICAL UNIVERSITY

G. Buss

Robust time series forecasting methods

THESIS

2013

RIGA TECHNICAL UNIVERSITY
Faculty of computer science and information technology
Institute of computer control, automation and computer engineering

Ginters BUSS
Student of doctoral program "Automation and computer engineering"

ROBUST TIME SERIES FORECASTING METHODS

Thesis

Supervisor
Dr. habil. math., Professor
J. CARKOVS

Supervisor
Dr. math., Assoc. Professor
V. AJEVSKIS

Riga 2013

This work has been supported by the European Social Fund within the project "Support for the implementation of doctoral studies at Riga Technical University".

Šis darbs izstrādāts ar Eiropas Sociālā fonda atbalstu projektā "Atbalsts RTU doktora studiju īstenošanai".

ABSTRACT

Forecasting is prevalent. Tourism industry forecasts the number of tourists (Athanasopoulos, Hyndman, Song and Wu, 2011). Energy industry forecasts the demand and price of energy (Raviv, Bouwman and van Dijk, 2013). Finance industry forecasts the prices for crude oil, grain, currency and securities (Asai, Caporin and McAleer, 2012). Policy makers make their decisions based on economic forecasts (Amisano and Geweke, 2013). One of the most widespread forecasting tools is the Box-Jenkins (Box and Jenkins, 1970) autoregressive integrated moving average (ARIMA) model. However, Box-Jenkins methodology has its drawbacks. First, a modern forecaster is faced with lots of potentially useful data which ARIMA models cannot handle. Second, in many areas of forecasting, e.g. in economics, data are noisy. Thus, forecasting methods should be used that are robust against the noise. Third, data dynamics may be subject to sudden change. The forecasting methods should be able to forecast robustly during such changes in dynamics.

The increasing demand for forecasting methods that would be able to handle potentially large sets of data subject to noise and changes in dynamics makes the topic of the thesis pertinent.

The main objective of the thesis thus is to develop robust forecasting methods that are able to work with noisy and high-dimensional data, with application to macroeconomics.

The key novelties of the thesis are the following. First, an asymmetric filter has been developed for frequency band extraction at the end-points of univariate series. Second, a method has been developed for signal extraction and forecasting using high-dimensional and noisy data sets. Third, robustness issues of Bayesian and factor forecasting models have been investigated when the dynamics of the target change rapidly.

The approbation of the thesis has been achieved by presenting the results at 11 international scientific conferences and seminars, by publishing 11 articles in international scientific journals or conference proceedings, by implementing the methods at the Central statistical bureau of Latvia and the Bank of Latvia.

The thesis consists of an introduction, three chapters and conclusions. It contains 136 pages, 69 figures, 9 tables and 92 references.

ANOTĀCIJA

Prognozēšana ir izplatīta. Tūrisma nozare prognozē potenciālo tūristu skaitu (Athanasopoulos, Hyndman, Song and Wu, 2011). Enerģētikas nozare prognozē enerģijas pieprasījumu un cenu (Raviv, Bouwman and van Dijk, 2013). Finanšu nozare prognozē jēlnaftas, graudu, valūtas un vērtspapīru cenas (Asai, Caporin and McAleer, 2012). Politikas lēmēji pieņem lēmumus, balstoties uz ekonomikas prognozēm (Amisano and Geweke, 2013). Viens no visizplatītākajiem prognozēšanas rīkiem ir Boksa-Dženkinsa (Box and Jenkins, 1970) autoregresīvais integrētais slīdošā vidējā (ARIMA) modelis. Taču Boksa-Dženkinsa metodoloģijai ir savi trūkumi. Pirmkārt, mūsdienu prognozētājam ir potenciāli daudz pieejamās informācijas, taču ARIMA metodoloģija nav spējīga apstrādāt liela apjoma informāciju. Otrkārt, daudzās prognozēšanas jomās, piemēram ekonomikā, dati ir trokšņaini. Tādējādi, prognozētājam būtu jālieto prognozēšanas metodes, kas ir robustas pret trokšņiem. Treškārt, datiem mēdz būt pēkšņas izmaiņas dinamikā. Būtu jālieto prognozēšanas metodes, kas ir robustas pret straujām izmaiņām dinamikā.

Strauji pieaugošais pieprasījums pēc prognozēšanas metodēm, kas ir piemērotas strādāt ar liela apjoma datiem, kas ir trokšņaini un pakļauti straujai dinamikas maiņai dara promocijas darba tēmu aktuālu.

Promocijas darba galvenais mērķis ir izstrādāt robustas laikrindu prognozēšanas metodes, kas ir piemērotas strādāt ar trokšņainiem un liela apjoma datiem, ar pielietojumu makroekonomikā.

Galvenie darba jaunieguvumi ir šādi. Pirmkārt, izstrādāts joslas filtrs viendimensionālu datu gala punkta novērtēšanas problēmām, kas atšķiras no alternatīvas ar citu ideālā filtra koeficientu korekciju. Otrkārt, izstrādāta metode signāla vienlaicīgai novērtēšanai un prognozēšanai daudzdimensionāliem un trokšņainiem datiem, kas ir pirmā šāda veida metode. Treškārt, novērtētas Beijesa un faktoru metožu robustuma problēmas, kad atkarīgā rādītāja dinamika strauji mainās, kas ir nepieciešams savlaicīgai dinamikas prognozēšanai.

Darba aprobācija tika veikta, prezentējot darba rezultātus 11 starptautiskajās zinātniskajās konferencēs un semināros, publicējot 11 zinātniskos rakstus starptautiskajos zinātniskajos izdevumos, pielietojot metodes LR Centrālajā statistikas pārvaldē un Latvijas Bankā.

Darbs sastāv no ievada, trīs nodaļām un nobeiguma. Darbs satur 136 lappuses, 69 attēlus, 9 tabulas un 92 nosaukumus literatūras sarakstā.

CONTENTS

INTRODUCTION

Pertinence of the thesis

Forecasting is prevalent. Tourism industry forecasts the number of tourists (Athanasopoulos, Hyndman, Song and Wu, 2011). Energy industry forecasts the demand and price of energy (Raviv, Bouwman and van Dijk, 2013). Finance industry forecasts the prices for crude oil, grain, currency and securities (Asai, Caporin and McAleer, 2012). Policy makers make their decisions based on economic forecasts (Amisano and Geweke, 2013); their decisions affect many people's lives.

One of the most widespread forecasting tools is the Box-Jenkins (Box and Jenkins, 1970) autoregressive integrated moving average (ARIMA) model. However, Box-Jenkins methodology has its drawbacks. First, a modern forecaster is faced with lots of potentially useful data which ARIMA models cannot handle. Second, in many areas of forecasting, e.g. in economics, data are noisy. Thus, forecasting methods should be used that are robust against the noise. Third, data dynamics may be subject to sudden change. The forecasting methods should be able to forecast robustly during such changes in dynamics.

The increasing demand for forecasting methods that would be able to handle potentially large sets of data subject to noise and changes in dynamics makes the topic of the thesis pertinent.

Objective and tasks of the thesis

The main objective of the thesis thus is to develop robust forecasting methods that are able to work with noisy and high-dimensional data, with applications in macroeconomics.

In order to fulfill the objective of the thesis, the following tasks are proposed:

- develop a univariate asymmetric bandpass filter for end-point estimation problems,
- compare the performance of the developed asymmetric filter to the currently most popular alternative in macroeconomics,
- develop a method suitable for forecasting and signal extraction using high-dimensional and noisy data,

- assess the properties of the above method and compare with the currently best alternative in macroeconomics,
- investigate the robustness issues for Bayesian and factor forecasting models.

Object and subject of the thesis

The object of the thesis is forecasting process of noisy and high-dimensional time series.

The subject of the thesis is the set and the system of filter and model algorithms for short-term forecasting that are suitable to work with noisy and high-dimensional macroeconomic data.

Research methods

The following methods are used in the preparation of the thesis: mathematical statistics and probability theory, optimization theory, frequency domain analysis and filtration theory, computer visualization method, and algorithm theory.

Novelty of the thesis

The main novelties of the thesis are:

1. An asymmetric filter has been developed for frequency band extraction at the endpoints of univariate series.
2. A method has been developed for signal extraction and forecasting using high-dimensional and noisy data sets.
3. Robustness issues of Bayesian and factor forecasting models have been investigated when the dynamics of the target change rapidly.

Practical applicability

1. Precise and timely estimate of business cycle conditions helps adopt the right decisions in monetary and fiscal policy that affect many people's lives;
2. The tighter link between the dependent and explanatory variables in the regularized filter methodology i) makes its estimates more robust against the presence of irrelevant explanatory variables thus making the variable pre-selection step easier, ii) makes forecasting easier, and iii) makes the decomposition of individual effects easier;

3. The results on the robustness of the Bayesian and factor methods helps to choose robust forecasting methods in real-time environment.

Approbation of the thesis

The approbation of the thesis has been achieved by presenting the results at 11 international scientific conferences and seminars (including 1 poster), by publishing 11 articles in international scientific journals or conference proceedings, by implementing the methods at the Central statistical bureau of Latvia for producing the official statistics of seasonally adjusted data and the flash release of Latvia's GDP since year 2009. The delivered models and filters have been used for forecasting purposes at Bank of Latvia since 2011.

Publications:

1. Buss, G. (2010), "A Note on Now-/Forecasting with Dynamic Versus Static Factor Models along a Business Cycle", 10th International Vilnius Conference on Probability Theory and Mathematical Statistics: Abstracts of Communications, Vilnius, Lithuania, 28 June - 2 July, 2010, p 119.

2. Buss, G. (2010), "Asymmetric Baxter-King Filter", Scientific Journal of RTU, 5th series, Computer Science, 42. vol, pp 95-99. (Indexed in: EBSCO, RePEc, SciVerse, Scirus, Econlit, Google Scholar, Microsoft Academic Search)

3. Buss, G. (2010), "Comparing forecasts of Latvia's GDP using simple seasonal ARIMA models and direct versus indirect approach: an overview", The results of statistical scientific research 2010, Research papers, Ed. O. Krastins, I. Vanags, Riga: Central Statistical Bureau of Latvia, pp 50-56. (Indexed in: RePEc, SciVerse, Scirus, Econlit, Google Scholar, Microsoft Academic Search)

4. Buss, G. (2010), "Economic Forecasts with Bayesian Autoregressive Distributed Lag Model: Choosing Optimal Prior in Economic Downturn", Aplimat: Journal of Applied Mathematics, vol 3, pp. 191-200. (Indexed in: RePEc, SciVerse, Scirus, Econlit, Google Scholar, Microsoft Academic Search)

5. Buss, G. (2010), "Economic Forecasts with Bayesian Autoregressive Distributed Lag Model: Choosing Optimal Prior in Economic Downturn", 6th Colloquium on Modern Tools for Business Cycle Analysis: "The Lessons from Global Economic Crisis", Book of Abstracts, Luxembourg, 26-29 September, 2010, pp 53.

6. Buss, G. (2010), "Forecasts with Single-Equation Markov-Switching Model: an Application to the Gross Domestic Product of Latvia", Journal of Applied Economic Sciences, Vol 5, Issue 2, pp 49-59. (Indexed in: Scopus, RePEc, SciVerse, Scirus, Econlit, Google Scholar, Microsoft Academic Search)

7. Buss, G. (2010), "Forecasts with Single-Equation Markov-Switching Model: an Application to the Gross Domestic Product of Latvia", Acta Societatis Mathematicae Latviensis: Abstract of the 8th Latvian Mathematical Conference, Valmiera, Latvia, 9-10 April, 2010, p 17.

8. Buss, G. (2011), "An Application of Direct Filter Approach: New Economic Indicators for Latvia", Scientific Journal of RTU, 5th series, Computer Science, 48 vol, pp 75-81. (Indexed in: EBSCO, Google Scholar)

9. Buss, G. (2011), "A Band Pass Filter for Real-Time Signal Extraction", Abstracts of 16th International Conference on Mathematical Modelling and Analysis, Sigulda, Latvia, 25-28 May, 2011, p 22.

10. Buss, G. (2011), "Preliminary Results on Asymmetric Baxter-King Filter", Aplimat 2011: 10th International conference on applied mathematics: Proceedings, Bratislava, Slovakia, 1-4 February, 2011, pp 1499-1508. (Indexed in: Google Scholar)

11. Buss, G. (2012), "Introduction to regularized DFA", Scientific Journal of RTU, series 5, vol. 48, pp 48-56. (Indexed in: EBSCO, Google Scholar)

Conferences:

1. Buss, G. "Forecasts with single-equation Markov-switching model: an application to the gross domestic product of Latvia" 10th International Vilnius Conference on Probability Theory and Mathematical Statistics, Lithuania, Vilnius, 28. June - 2. July, 2010

2. Buss, G. "Economic forecasts with Bayesian autoregressive distributed lag model: choosing optimal prior in economic downturn", Aplimat, 9th International Conference, Slovakia, Bratislava, 2.-5. February, 2010

3. Buss, G. "Forecasts with single-equation Markov-switching model: an application to the gross domestic product of Latvia", 8th Latvian Mathematical Conference, Latvia, Valmiera, 9.-10. April, 2010

4. Buss, G. "Economic forecasts with Bayesian autoregressive distributed lag model: choosing optimal prior in economic downturn", 8th Latvian Mathematical Conference, Latvia, Valmiera, 9.-10. April, 2010

5. Buss, G. "Asymmetric Baxter-King Filter", 51. RTU International Scientific conference, Section: Computer Science, Subsection: Technologies of computer control, 11-15 October, 2010, Riga, Latvia

6. Buss, G. "Asymmetric Baxter-King filter: business cycle estimation in real time", Finance and economics conference 2011, Lupcon Center for Business Research, 5-6 July, 2011, Frankfurt, Germany

7. Buss, G. "Asymmetric Baxter-King filter", 26th Annual congress of the European Economic Association and the 65th European meeting of the Econometric Society, 25-29 August, 2011, Oslo, Norway

8. Buss, G. "An application of direct filter approach: new economic indicators for Latvia", 52. RTU International Scientific conference, Section: Computer Science, Subsection: Technologies of computer control, 13 October, 2011, Riga, Latvia

9. Buss, G. "A new real-time indicator for the euro area GDP", Working Group on Forecasting, European Central Bank, 28-29 June, 2012, Tallinn, Estonia

10. Buss, G. "Introduction to regularized direct filter approach", 53. RTU International Scientific conference, Section: Computer Science, Subsection: Technologies of computer control, 13 October, 2012, Riga, Latvia

11. Buss, G. "Forecasting and signal extraction with regularized multivariate direct filter approach", 28th Annual congress of the European Economic Association and the 67th European meeting of the Econometric Society, 26-30 August, 2013, Gothenburg, Sweden

Structure and volume of the thesis

The thesis consists of an introduction, three chapters, conclusions and a bibliography. It contains 136 pages, 69 figures, 9 tables and 92 references. The structure of the thesis is the following:

The ***Introduction*** describes the pertinence of the thesis, the objective and tasks of the thesis, the object and the subject of the thesis, research methods used, its practical applicability, and the approbation of the thesis.

The first chapter ***"Asymmetric Baxter-King filter for end-point estimation"*** proposes an extension of the symmetric Baxter-King band pass filter to an asymmetric Baxter-King filter. Since the symmetric Baxter-King filter is unable to extract the desired signal at the very ends of the series, the extension to an asymmetric filter is useful whenever the real time estimation is needed. The chapter uses monte carlo simulation to compare the proposed filter's properties in extracting business cycle frequencies to the ones of the original Baxter-King filter and Christiano-Fitzgerald filter. Simulation results show that the asymmetric Baxter-King filter is superior to the asymmetric default specification of Christiano-Fitzgerald filter in real time signal extraction exercises.

The second chapter ***"Multivariate filter for high-dimensional and noisy datasets"*** develops an algorithm for high-dimensional filtering and real-time signal extraction. It is shown that the filtration algorithm is able to process high-dimensional and noisy data sets. The chapter illustrates the features of the filter by tracking the medium-to-long-run component in GDP growth for euro area, including the replication of an established indicator's behavior, as well as producing more timely indicators. Further robustness check is performed on a less homogeneous Latvia's dataset. The method is found to be a promising tool for both concurrent estimation and forecasting using high-dimensional datasets, and a decent alternative to dynamic factor methodology.

The third chapter ***"Robustness of traditional methods and forecasting system overview"*** studies the robustness issues of the Bayesian and factor methodologies. It finds that the Bayesian Minnessota prior and the exact dynamic factors are not robust against a rapid change in the dynamics of the target variable. The chapter also summarizes the methods considered in the thesis and describes a forecasting system involving methods developed in the thesis.

Main conclusions

Bibliography

1 ASYMMETRIC BAXTER-KING FILTER FOR END-POINT ESTIMATION

This chapter proposes an extension of the symmetric Baxter-King band pass filter to an asymmetric Baxter-King filter. The optimal correction scheme of the ideal filter weights is the same as in the symmetric version, i.e, cut the ideal filter at the appropriate length and add a constant to all filter weights to ensure zero weight on zero frequency. Since the symmetric Baxter-King filter is unable to extract the desired signal at the very ends of the series, the extension to an asymmetric filter is useful whenever the real time estimation is needed. The chapter uses monte carlo simulation to compare the proposed filter's properties in extracting business cycle frequencies to the ones of the original Baxter-King filter and Christiano-Fitzgerald filter. Simulation results show that the asymmetric Baxter-King filter is superior to the asymmetric default specification of Christiano-Fitzgerald filter in real time signal extraction exercises.

1.1 Introduction

This section proposes an extension of the symmetric Baxter-King band pass filter (Baxter and King, 1999) to an asymmetric Baxter-King filter. Such modification, to the best of my knowledge, has not been discussed in the literature. Symmetric filters are not applicable at the very ends of an input signal without the extension of the ends with forecasts. Thus, asymmetric band pass filters are necessary to extract the desired band of frequencies at the ends of an input signal, if forecasting is not used for extending the ends of the input signal.

The closest band pass filter to the Baxter-King filter is Christiano-Fitzgerald band pass filter (Christiano and Fitzgerald, 2003) which, in general, is asymmetric, and whose default specification is optimized for an input signal following a random walk (RW) process, but it allows the input signal to follow other data generating processes (DGP). However, Christiano and Fitzgerald (2003) argue that their default specification of the filter is a good approximation to many DGPs observed in macroeconomic time series and, thus, macroeconomists may opt for it. Although Christiano and Fitzgerald (2003) compares their filter to the symmetric Baxter-King filter, they do not elaborate on an asymmetric version of the Baxter-King filter.

This section formally develops an asymmetric version of the Baxter-King filter and assesses its properties in extracting business cycle frequencies, in comparison to the symmetric Baxter-King filter, and symmetric and asymmetric default specification of Christiano-Fitzgerald filter, by using monte carlo simulation. The results show that, given the considered DGP and metric (estimated correlation of the true and extracted cycles at time t), the asymmetric Baxter-King filter is superior to the asymmetric Christiano-Fitzgerald filter at the very ends of a sample, thus indicating that the asymmetric Baxter-King filter should be preferred over the asymmetric Christiano-Fitzgerald filter in real time signal extraction exercises. Several other 'interesting' results are obtained, like, fixed-length symmetric filters outperforming their asymmetric counterparts, and shorter-length symmetric filters outperforming longer-length symmetric filters (which is contrary to the assertion by Christiano and Fitzgerald, 2003, and *a priori* belief of mine).

The chapter is organized as follows. Section 1.2 develops the filter, section 1.3 assesses the performance of the filter by means of monte carlo simulation, and section 1.4 concludes.

1.2 Deriving the filter

Consider the following orthogonal decomposition of the zero-mean covariance stationary stochastic process, x_t:

$$x_t = y_t + \tilde{x}_t. \tag{1.1}$$

The process, y_t, has power only in frequencies (measured in radians) belonging to the interval $\{[a_1, a_2] \cup [-a_2, -a_1]\} \subset (-\pi, \pi)$, where $0 < a_1 < a_2 < \pi$. The process, $\tilde{x}_t$, has power only in the complement of this interval in $(-\pi, \pi)$. By the spectral representation theorem,

$$y_t = b(L)x_t, \tag{1.2}$$

where the ideal band pass filter, $b(L)$, is

$$b(L) = \sum_{h=-\infty}^{\infty} b_h L^h, \quad L^h x_t = x_{t-h}, \tag{1.3}$$

with

$$\begin{aligned} b_h &= \frac{\sin(ha_2) - \sin(ha_1)}{\pi h}, \quad h = \pm 1, \pm 2, \dots \\ b_0 &= \frac{a_2 - a_1}{\pi}, \quad a_1 = \frac{2\pi}{p_u}, \quad a_2 = \frac{2\pi}{p_l}, \end{aligned} \tag{1.4}$$

and $p_u, p_l \in (2, \infty)$ define the upper and lower bounds of the wave length of interest. With b_h's specified as in (1.4), the frequency response function of the ideal filter at frequency

ω is

$$\begin{aligned} \beta(\omega) &= 1 \quad \text{for } \omega \in [a_1, a_2] \cup [-a_2, -a_1] \\ &= 0 \quad \text{otherwise.} \end{aligned} \tag{1.5}$$

Fig. 1.1 shows the amplitude of the ideal bandpass filter with cut-off wave lengths 18 and 96 months and the absolute value of the discrete Fourier transform of Latvia's gross domestic product (interpolated to monthly frequency).

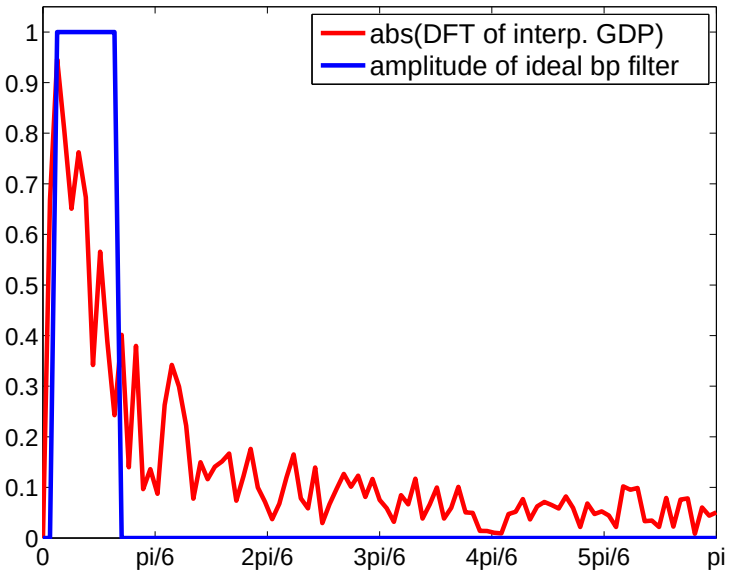

Fig. 1.1. The absolute value of the discrete Fourier transform of Latvia's gross domestic product (interpolated to monthly frequency) and the amplitude of the ideal bandpass filter with cut-off wave lengths 18 and 96 months.

Baxter and King (1999) have proposed to obtain a symmetric, fixed length approximation to the ideal filter, (1.3) and (1.4), by minimizing

$$\begin{aligned} Q &= \int_{-\pi}^{\pi} \delta(\omega)\delta(-\omega)d\omega \\ &\text{s.t.} \\ \hat{\beta}(0) &= \sum_{k=-K}^{K} \hat{b}_k = 0 \\ \hat{b}_k &= \hat{b}_{-k}, \end{aligned} \tag{1.6}$$

where $\delta(\omega) = \beta(\omega) - \hat{\beta}(\omega)$ is the discrepancy between the exact and the approximate filter amplitudes at frequency ω, and the constraint $\hat{\beta}(0) = 0$ is to ensure zero weight on the trend frequency, in line with the assumption $a_1 > 0$. The solution to (1.6) is a truncation of the ideal filter symmetrically at length K, and addition of a constant $(-\sum_{k=-K}^{K} b_k)/(2K+1)$ to all filter weights to ensure $\hat{\beta}(0) = 0$. Baxter and King (1999)

suggest the value of K to be about 3 years, i.e, K=12 for quarterly data, and K=36 for monthly data. The symmetry of the filter together with the condition $\hat{\beta}(0) = 0$ implies that the filter renders stationary time series that is integrated of oder 2 (I(2)) or less. Thus, the symmetric BK filter has trend-reduction property and, therefore, it can be applied to nonstationary, up to I(2) series.

Since the symmetric BK filter can not be used to extract the desired frequencies at the very end (for the first and the last K observations) of the input series, a natural extension of the Baxter and King (1999) filter is to allow the approximate filter to be asymmetric, to be able to use the filter in real time. In order to optimally approximate an ideal symmetric linear filter in a Baxter-King sense, the problem is to minimize

$$Q = \int_{-\pi}^{\pi} \delta(\omega)\delta(-\omega)d\omega$$

s.t.

$$\hat{\beta}(0) = \sum_{h=-p}^{f} \hat{b}_h = 0. \tag{1.7}$$

The condition $\hat{\beta}(0)$ ensures zero weight on zero frequency, thus this asymmetric filter also has a trend-reduction property, however, it alone, without symmetry, is not sufficient to render I(2) process stationary. Thus, the ability of the asymmetric BK filter of real time signal extraction comes at a cost of losing the power to eliminate two unit roots from the input series.

To solve (1.7), form the Lagrangian

$$\mathcal{L} = Q - \lambda\hat{\beta}(0) \tag{1.8}$$

with first order conditions (FOCs):

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \hat{b}_h} &= \frac{\partial Q}{\partial \hat{b}_h} - \lambda = 0 \\ \frac{\partial \mathcal{L}}{\partial \lambda} &= -\hat{\beta}(0) = 0. \end{aligned} \tag{1.9}$$

Since

$$\frac{\partial}{\partial \hat{b}_h}[\delta(\omega)\delta(-\omega)] = \frac{\partial \delta(\omega)}{\partial \hat{b}_h}\delta(-\omega) + \delta(\omega)\frac{\partial \delta(-\omega)}{\partial \hat{b}_h}, \tag{1.10}$$

and since the frequency response function of the approximating filter is $\hat{\beta}(\omega) = \sum_{h=-p}^{f} \hat{b}_h e^{-i\omega h}$, it follows that

$$\frac{\partial \delta(\omega)}{\partial \hat{b}_h} = -e^{-i\omega h}. \tag{1.11}$$

(1.7), (1.10) and (1.11) imply

$$\frac{\partial Q}{\partial \hat{b}_h} = -\int_{-\pi}^{\pi} [e^{-i\omega h}\delta(-\omega) + \delta(\omega)e^{i\omega h}]d\omega. \tag{1.12}$$

By the property $\int_{-\pi}^{\pi}[f(\omega)+f(-\omega)]d\omega = 2\int_{-\pi}^{\pi} f(\omega)d\omega$ (since $\int_{-\pi}^{\pi} f(\omega)d\omega = \int_{0}^{\pi} f(\omega)d\omega + \int_{-\pi}^{0} f(\omega)d\omega = \int_{0}^{\pi}[f(\omega)+f(-\omega)]d\omega$ is real, then $\int_{-\pi}^{\pi} f(\omega)d\omega = \int_{-\pi}^{\pi} f(-\omega)d\omega$, and the property follows), (1.12) becomes

$$\frac{\partial Q}{\partial \hat{b}_h} = -2\int_{-\pi}^{\pi} \delta(\omega)e^{i\omega h}d\omega. \tag{1.13}$$

By the property

$$\begin{aligned}\int_{-\pi}^{\pi} e^{i\omega n}e^{-i\omega m}d\omega = \int_{-\pi}^{\pi} e^{-i\omega(m-n)}d\omega = 0 \quad &\text{for } n \neq m \\ = 2\pi \quad &\text{for } n = m,\end{aligned} \tag{1.14}$$

obtain

$$\int_{-\pi}^{\pi} \delta(\omega)e^{i\omega h}d\omega = \int_{-\pi}^{\pi}\left[\sum_{k=-\infty}^{\infty} b_k e^{-i\omega k} - \sum_{j=-p}^{f}\hat{b}_j e^{-i\omega j}\right]e^{i\omega h}d\omega = 2\pi[b_h - \hat{b}_h]. \tag{1.15}$$

Given (1.15), the FOCs are

$$-4\pi[b_h - \hat{b}_h] - \lambda = 0. \tag{1.16}$$

If there is no constraint on $\hat{\beta}(0)$, the optimal approximate (in Baxter-King sense) filter is simply derived by truncation of the ideal filter's weights. If there is a constraint on $\hat{\beta}(0)$, then λ must be chosen so that the constraint is satisfied. For this purpose, rewrite (1.16) as

$$\hat{b}_h = b_h + \theta,$$

where $\theta = \lambda/(4\pi)$. In order to have $\hat{\beta}(0) = \sum_{h=-p}^{f}\hat{b}_h = 0$, the required adjustment is

$$\theta = \frac{-\sum_{h=-p}^{f} b_h}{p+f+1}, \tag{1.17}$$

which yields the same optimal weight adjustment scheme as in the symmetric Baxter-King filter case.

Fig. 1.2 illustrates coefficients of 51-observation long symmetric BK filter and one-sided asymmetric BK filter targeting business cycle frequencies. The outputs of BK and

ABK filters applied on a sample data are shown in Fig. 1.3. Clearly, only the ABK filter can be used at the end point of time series.

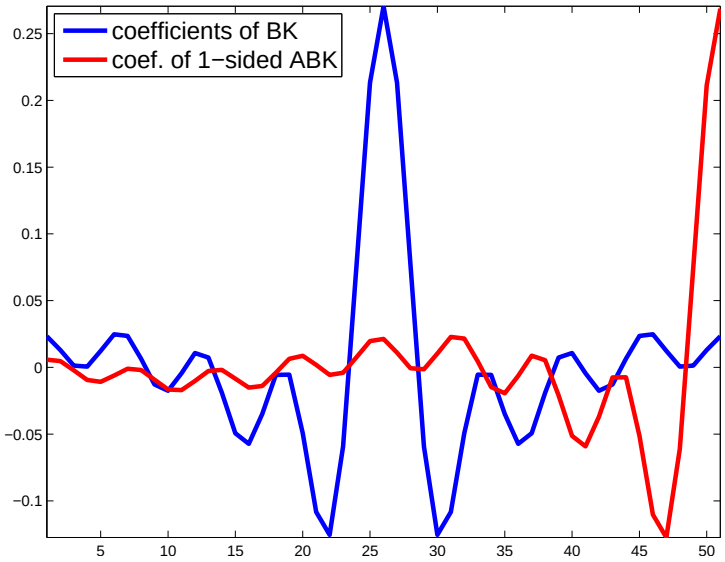

Fig. 1.2. Coefficients of 51-observation long symmetric BK filter and one-sided asymmetric BK filter targeting business cycle frequencies.

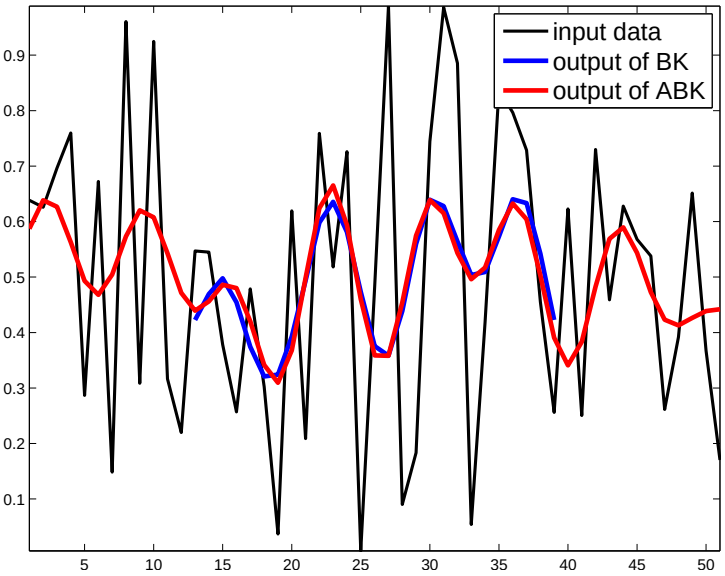

Fig. 1.3. The outputs of 51-observation long symmetric BK filter and asymmetric BK filter applied on a sample data.

Fig. 1.4 shows a flowchart for using the Baxter-King filter. In order to use the filter, the user has to choose a one-dimensional input series and the upper and lower bounds on the length of the cycle one wants to extract. In macroeconomics, these bounds typically are defined by the length of business cycle, i.e., between 1.5-8 years. In other fields, these might be different.

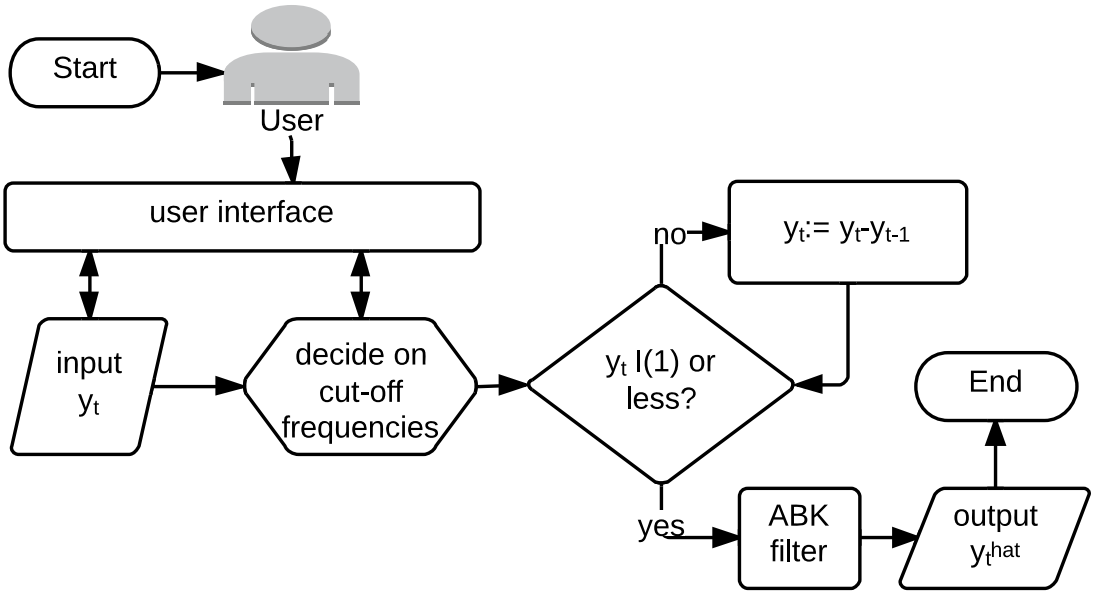

Fig. 1.4. A flowchart of the univariate asymmetric filter algorithm

The input series must be at most first-order integrated, otherwise differentiation is needed. Most seasonally adjusted macroeconomic series are up to I(1), so, typically, no data transformation is required.

The next section describes results from monte carlo simulation to assess the performance of the proposed filter.

1.3 Comparing filter performance by monte carlo simulation

This section assesses the performance of the proposed filter to extract business cycle frequencies (corresponding to wave length between 1.5 and 8 years) in comparison to the original BK filter, as well as symmetric and asymmetric Christiano-Fitzgerald (CF) filter which is optimized for an input signal following a random walk (RW) process (Christiano and Fitzgerald, 2003). Thus, the asymmetric CF filter assumes that the first difference of the input signal is mean-zero covariance stationary process. The symmetric CF filter allows for the input signal to follow RW with drift.

Consider the following data generating process (DGP):

$$y_t = \mu_t + c_t, \tag{1.18}$$

where

$$\mu_t = \mu_{t-1} + \epsilon_t \tag{1.19}$$

$$c_t = \phi_1 c_{t-1} + \phi_2 c_{t-2} + \eta_t \tag{1.20}$$

$$\epsilon_t \sim nid(0, \sigma_\epsilon^2),\ \eta_t \sim nid(0, \sigma_\eta^2). \tag{1.21}$$

Equation (1.18) defines a series, y_t, as the sum of a permanent component (stochastic trend), μ_t, and a cyclical component, c_t. The trend, μ_t, in this case is specified as a random walk process. The dynamics of the cyclical component, c_t, is specified as a second order autoregressive (AR(2)) process so that the peak of the spectrum of c_t could be at zero frequency or at business cycle frequencies. Disturbances, ϵ_t and η_t, are assumed to be uncorrelated.

The spectrum of an AR(2) process is

$$f_c(\omega) = \frac{\sigma_\eta^2}{1 + \phi_1^2 + \phi_2^2 - 2\phi_1(1-\phi_2)\cos\omega - 2\phi_2\cos(2\omega)} \tag{1.22}$$

with a peak at frequency other than zero for

$$\phi_2 < 0 \,\text{and}\, \left|\frac{\phi_1(1-\phi_2)}{4\phi_2}\right| < 1 \tag{1.23}$$

with the corresponding frequency $\omega = \cos^{-1}[-\phi_1(1-\phi_2)/(4\phi_2)]$ (Box, Jenkins and Reinsel, 1994; Priestley, 1981).

Data are generated from (1.18) with $\phi_1 = 1.2$ and different values for ϕ_2 to control the location of the peak in the spectrum of the cyclical component. I also vary the ratio of standard deviations of the disturbances, $\sigma_\epsilon/\sigma_\eta$, to change the relative importance of components of y_t. Such DGP can create series with spectral characteristics typical to macroeconomic variables, such as gross domestic product and inflation (Watson, 1986; Guay and St-Amant, 2005). The idea of such simulation is taken from Guay and St-Amant (2005).

Particularly, 10,000 samples of length 401 are created, with the first 200 observations of each sample dropped off as burn-in. The vector $[\phi_1, \phi_2]$ is set to five different values, as shown in Table 1.1.

Table 1.1.

Five different values of $[\phi_1, \phi_2]$ for the DGP

ϕ_1	ϕ_2	Fundamental period of the cycle (yrs)
1.2	-0.25	$\approx \infty$
1.2	-0.35	$\gg 8$
1.2	-0.44	8.2
1.2	-0.5	3.5
1.2	-0.8	1.9

The value of $\sigma_\epsilon/\sigma_\eta$ is set to change from 0 to 9.9 with step size 0.15 (Watson (1986) estimated this ratio for the U.S. GNP to be 0.75).

I compare four filters in their capability to extract business cycle frequencies: i) symmetric, fixed-length BK filter with $K = 12$ (see (1.6)), ii) asymmetric BK filter described

in Subsection 1.2, iii) symmetric, fixed-length CF filter with $K = 12$ for RW processes, and iv) default asymmetric specification of CF filter for RW processes.

The performance of filters is assessed by comparing the estimated correlation of the true cyclical component at time t with the estimated cyclical component at time t, $\hat{\rho}(c_t, \hat{c}_t)$, and by comparing the true AR(2) regression coefficients for the cycle with the fitted AR(2) regression coefficients. The distance measure $\hat{\rho}(c_t, \hat{c}_t)$ plays a key role from this point forward.

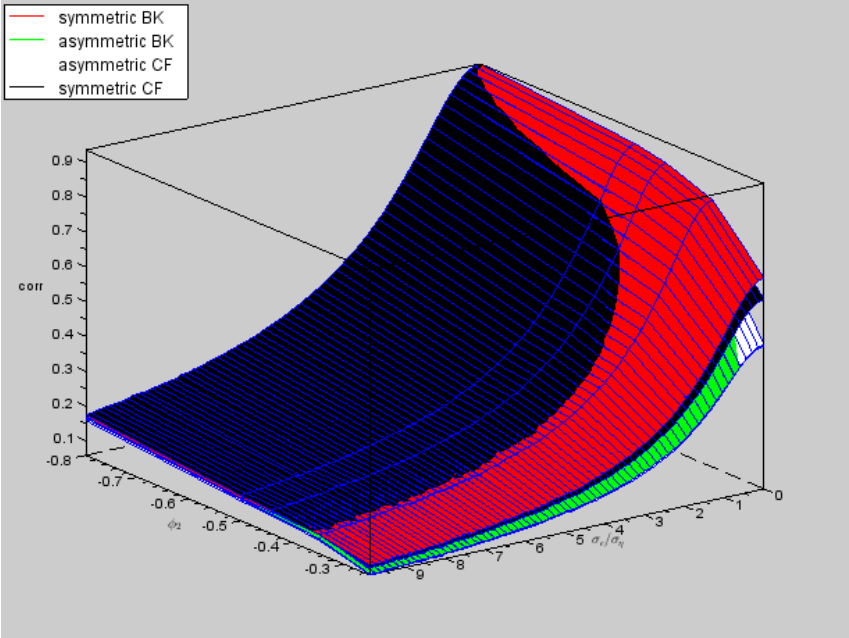

Fig. 1.5. Average estimated correlation between the true and estimated cyclical components at time t for given $[\phi_1, \phi_2]$ and $\sigma_\epsilon/\sigma_\eta$ values

Note: The correlation is estimated for the whole sample span except for the first and the last K=12 observations. The results show that, on average over the sample, symmetric filters are superior to asymmetric filters, and that asymmetric BK filter is superior to asymmetric CF filter.

Fig. 1.5 shows average estimated correlation between the true and estimated cyclical components, $\hat{\rho}(c_t, \hat{c}_t)$, for given values of $[\phi_1, \phi_2]$ and $\sigma_\epsilon/\sigma_\eta$. The correlation is estimated for the whole sample span except for the first and the last K=12 observations, since fixed-length symmetric filters do not produce the estimated cycle for those observations; these K=12 observations are deleted from the output of the asymmetric filters for a fair comparison between symmetric and asymmetric filters. Fig. 1.5 shows a similar behavior between the filters - their performance decreases with $\sigma_\epsilon/\sigma_\eta$, which is an expected result. When $\sigma_\epsilon/\sigma_\eta = 0$, the input signal is the true cycle, so the output signal (estimated cycle) correlates highly with the input. As $\sigma_\epsilon/\sigma_\eta$ increases, the influence of the permanent component in the input increases, thus making harder for filters to extract the cycle, thus the estimated correlation between the true and estimated cycles, $\hat{\rho}(c_t, \hat{c}_t)$, decreases.

Fig. 1.5 also shows that the performance of all filters decreases with an increasing ϕ_2. The value of $\phi_2 = -0.8$ together with $\phi_1 = 1.2$ corresponds to the length of the cycle 1.9

years, which is close to the usually defined minimum length of a business cycle, 1.5 years. The value of $\phi_2 = -0.44$ together with $\phi_1 = 1.2$ produces the cycle of length approximately 8.2 years, which is close to the usually defined maximum length of a business cycle, 8 years. With higher than $\phi_2 = -0.44$ values, the length of the true cycle rapidly increases. Although with $\phi_2 = -0.25$ the cycle still is considered stationary ($\phi_1 + \phi_2 < 1$, $\phi_2 - \phi_1 < 1$, and $|\phi_2| < 1$), it is a close approximation to a nonstationary process in a finite sample (Campbell and Perron, 1991). Thus, Fig. 1.5 shows expected deterioration in performance of BK filters as ϕ_2 increases. The similar deterioration in performance of the CF filters with an increasing length of the cycle was less expected. Another unexpected result is the inferior performance of asymmetric filters to their shorter symmetric counterparts.

Fig. 1.6 shows average correlation between the true and estimated cyclical components for given $[\phi_1, \phi_2]$ and $\sigma_\epsilon/\sigma_\eta$ values from fixed-length symmetric BK and CF filters.

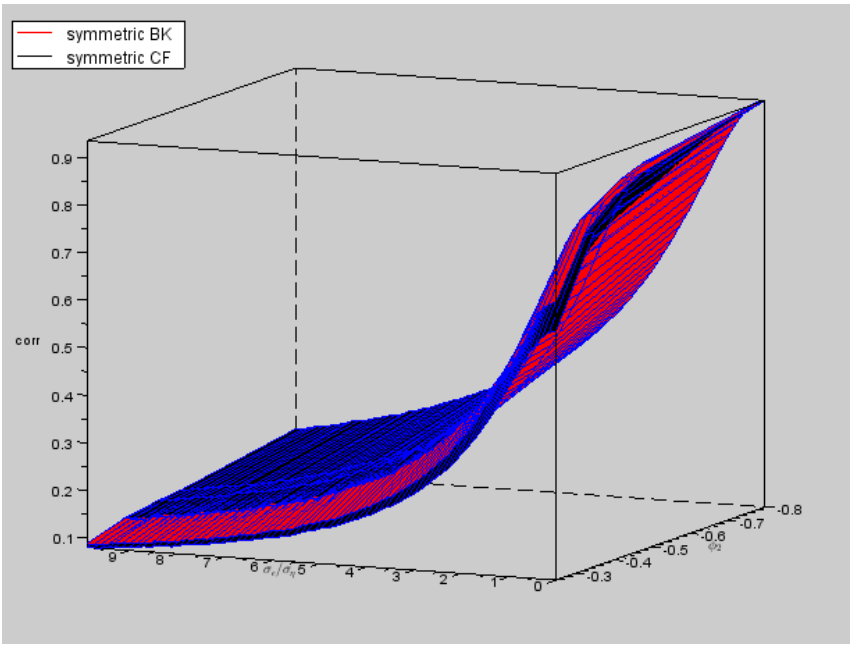

Fig. 1.6. Average correlation between the true and estimated cyclical components at time t for given $[\phi_1, \phi_2]$ and $\sigma_\epsilon/\sigma_\eta$ values from fixed-length symmetric BK and CF filters

Note: The correlation is estimated for the whole sample interval except for the first and the last K=12 observations. The performance of the filters is similar, regardless of wave length of the cycle or the influence of the permanent component.

This figure shows that the distance between the two correlation surfaces is not high at any point. The observation from Fig. 1.6 suggest that the performance of symmetric BK and CF filters is roughly the same regardless of the cycle length or the share of the permanent component.

Fig. 1.7 shows the view from the top of Fig. 1.6 to assess the regions of (although small, as seen in Fig. 1.6) relative superiority of the fixed-length symmetric BK and CF filters.

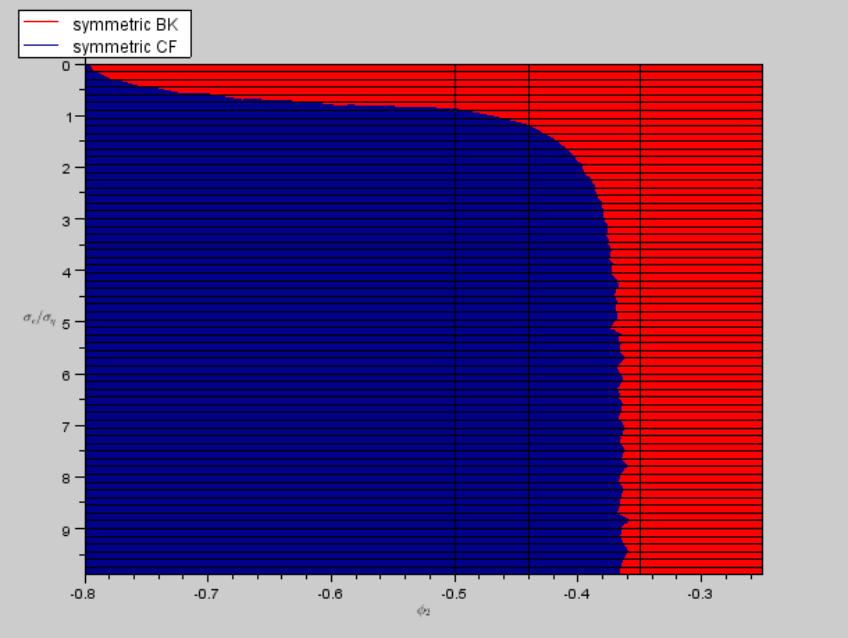

Fig. 1.7. The view from the top of Fig. 1.6, showing relative superiority of the fixed-length symmetric BK and CF filters

Note: The horizontal axis represent cycle length, while the vertical axis represent the importance of permanent component in the series. Fixed-length symmetric BK filter is superior to fixed-length symmetric CF filter for $0 \leq \sigma_\epsilon/\sigma_\eta < 0.5$ if cycle length is longer than 2 years. For most of the rest of the region, particularly - cycle length less than 8 years, given $\sigma_\epsilon/\sigma_\eta \geq 1$ - CF filter is slightly superior to the BK filter. For the remainder, i.e., $0.5 \leq \sigma_\epsilon/\sigma_\eta < 1$, CF filter shows superiority when cycle is relatively short (up to 3.5 years), and BK filter shows superiority when the cycle is longer than approximately 3.5 years.

Fig. 1.7 shows that fixed-length symmetric BK filter is superior to fixed-length symmetric CF filter for $0 \leq \sigma_\epsilon/\sigma_\eta < 0.5$ if cycle length is longer than 2 years. For most of the rest of the region, particularly - cycle length less than 8 years, given $\sigma_\epsilon/\sigma_\eta \geq 1$ - CF filter is slightly superior to the BK filter. For the remainder, i.e., $0.5 \leq \sigma_\epsilon/\sigma_\eta < 1$, CF filter shows superiority when cycle is relatively short (up to 3.5 years), and BK filter shows superiority when the cycle is longer than approximately 3.5 years.

Fig. 1.8 shows comparison of $\hat{\rho}(c_t, \hat{c}_t)$ between fixed-length symmetric and asymmetric BK (Fig. 1.8(a)) and CF (Fig. 1.8(b)) filters.

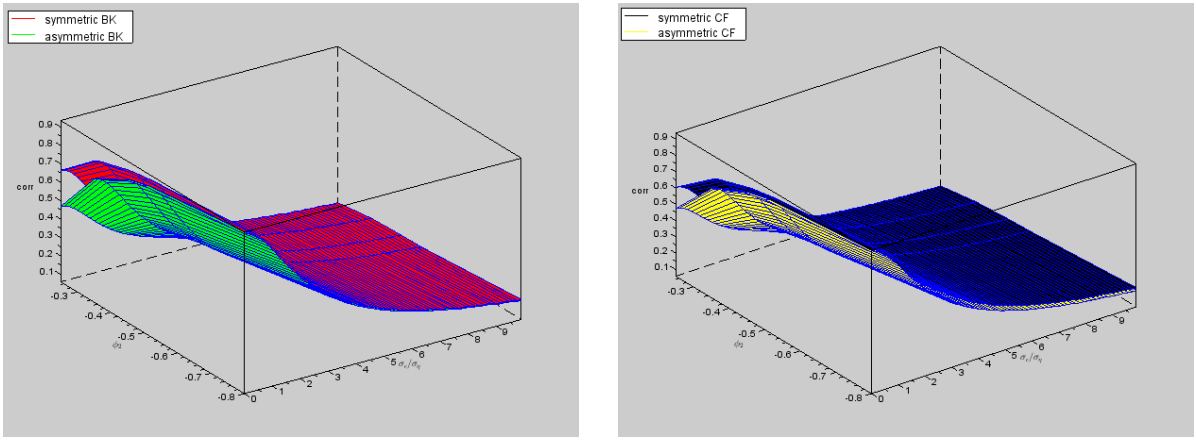

(a) BK (b) CF

Fig. 1.8. A comparison of performance of symmetric versus asymmetric filters

Note: In terms of average $\hat{\rho}(c_t, \hat{c}_t)$ over the sample. The left figure compares BK filters, and the right figure compares CF filters. The results show that symmetric filters are superior to their asymmetric counterparts.

Fig. 1.8 shows that symmetric filters are superior to asymmetric filters for all considered lengths of cycle and all proportions of permanent and cyclical components in the input series. When there is high influence of permanent component in the series, the performance of symmetric and asymmetric filters is close regardless of cycle length. As the influence of cyclical component rises, the performance of asymmetric filters (relative to symmetric ones) deteriorates with increasing length of cycle. Slightly more evident decrease of correlation between true and estimated cycles is for the asymmetric BK filter than for the CF filter. However, this is mainly due to the higher performance of symmetric BK filter compared to symmetric CF filter. The results in Fig. 1.8 are in contrast to those drawn by Christiano and Fitzgerald (2003), who conclude that filters using all the data, which, therefore, are asymmetric and time-varying, improve the extraction of the desired frequencies, compared to fixed-length symmetric filters.

Fig. 1.9 shows average correlation between the true and estimated cyclical components for given $[\phi_1, \phi_2]$ and $\sigma_\epsilon/\sigma_\eta$ values from asymmetric BK and CF filters.

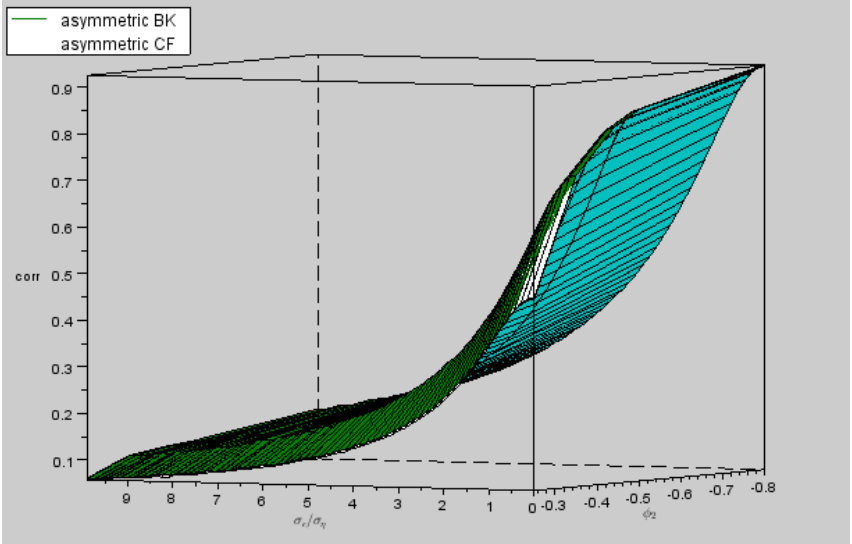

Fig. 1.9. Average correlation between the true and estimated cyclical components at time t for given $[\phi_1, \phi_2]$ and $\sigma_\epsilon/\sigma_\eta$ values from asymmetric filters

Note: The correlation is estimated for the whole sample interval except for the first and the last K=12 observations. Results show very similar performance of the filters - the correlation surfaces are almost identical.

This figure shows that the distance between the two correlation surfaces is practically nil at all points. The observation from Fig. 1.9 suggest that the performance, on average over the sample, of asymmetric BK and CF filters is almost the same regardless of the cycle length or the share of the permanent component.

Fig. 1.10 shows the view from the top of Fig. 1.9 to assess the regions of (small, as seen in Fig. 1.9) relative superiority of the asymmetric BK and CF filters.

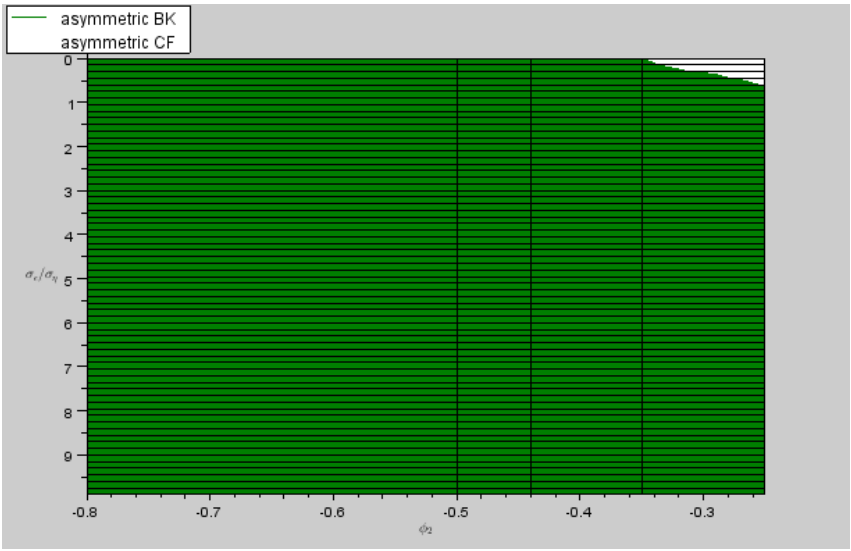

Fig. 1.10. The view from the top of Fig. 1.9, showing relative superiority of the asymmetric BK and CF filters

Note: The horizontal axis represent cycle length, while the vertical axis represent the importance of permanent component in the series. Even if the average performance of the asymmetric filters over the sample is very close, this figure shows that asymmetric BK filter is persistently superior to the asymmetric CF filter, regardless of wave length of the cycle, or the influence of permanent component in the input signal. This result is due to BK filter's superiority at the ends of a sample, see below.

Fig. 1.10 shows that the asymmetric BK filter is superior to asymmetric CF filter. The reason of the slight superiority of the BK filter will be evident below, when comparing the performances at the ends of a sample. A slightly surprising finding from Fig. 1.10 is the inability of asymmetric CF filter to perform better than the asymmetric BK filter in the region of high influence of the permanent component (corresponds to lower part of the graph).

Now, let us compare the performance of the asymmetric filters for the K=12 observations of the sample, where the fixed-length symmetric filters can not be applied. Fig. 1.11 to 1.14 show the estimated correlation of the true and estimated cycles at each of the K=12 observations, calculated across the 10,000 replications, and averaged over both symmetric ends.

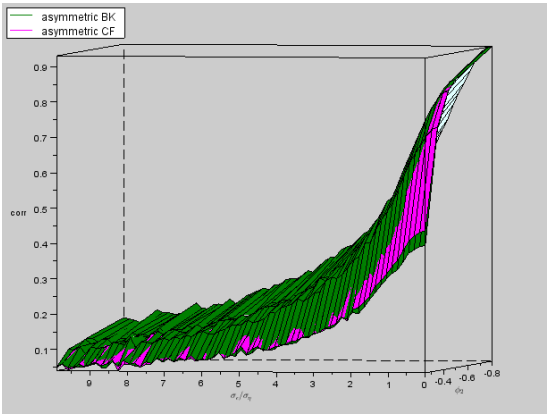

(a) Correlation at obs. 12

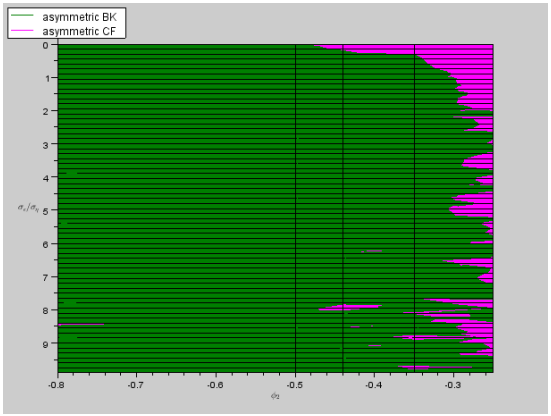

(b) view at 1.11(a) from the top

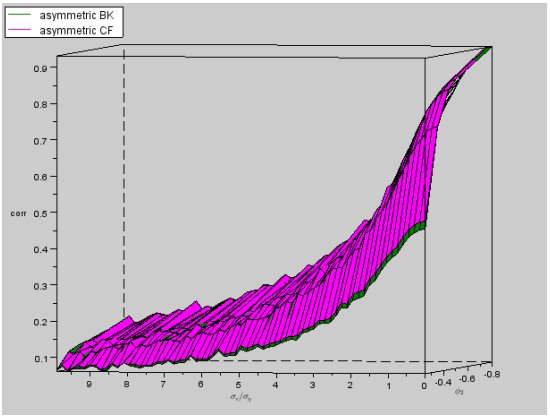

(c) Correlation at obs. 11

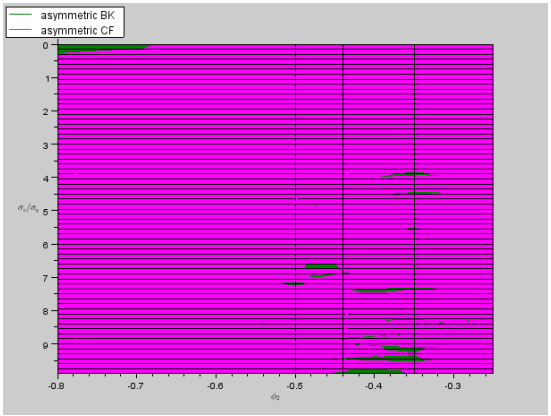

(d) view at 1.11(c) from the top

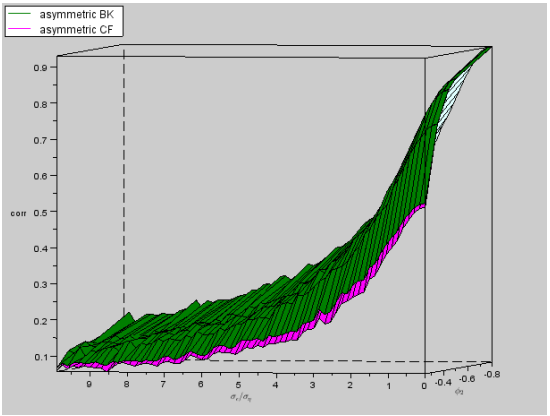

(e) Correlation at obs. 10

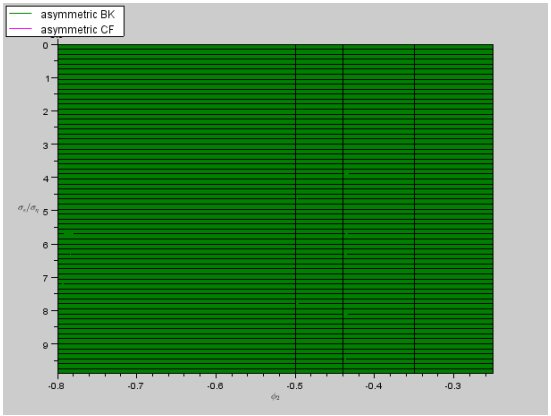

(f) view at 1.11(e) from the top

Fig. 1.11. Estimated correlation of the true and extracted cycles at time t, $\hat{\rho}(c_t, \hat{c}_t)$, by asymmetric BK and CF filters at observations number 12 to 10, counting from the end of the series

Note: The results show a close performance of the two filters, although the BK filter performs slightly better than the CF filter.

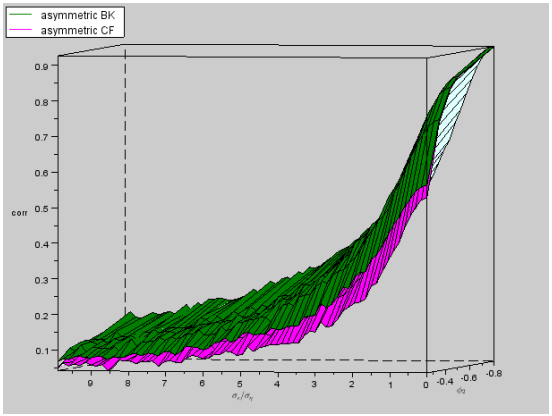

(a) Correlation at obs. 9

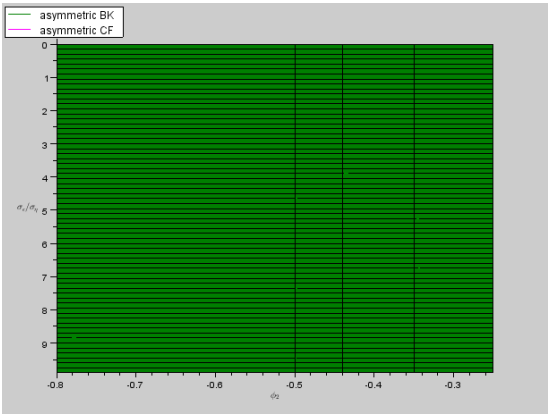

(b) view at 1.12(a) from the top

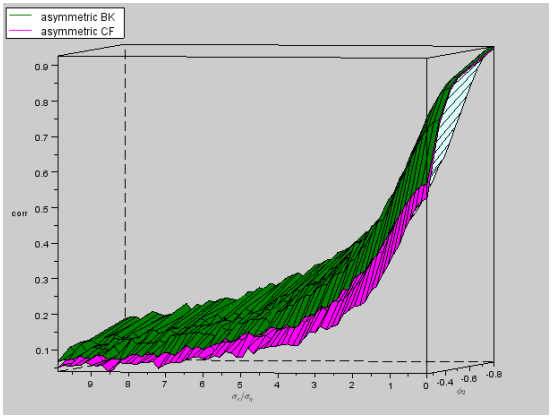

(c) Correlation at obs. 8

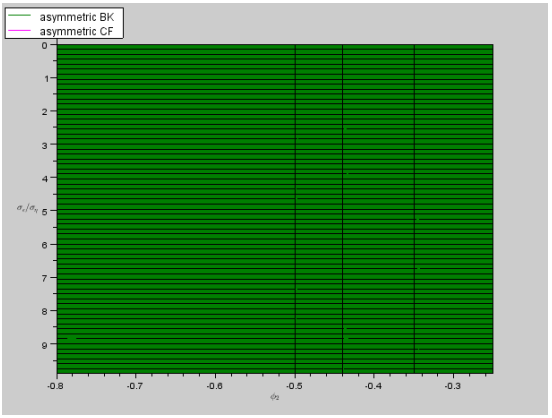

(d) view at 1.12(c) from the top

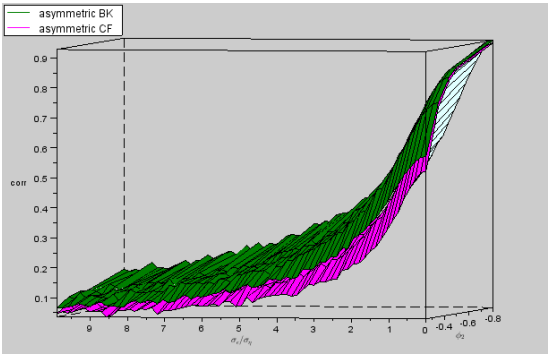

(e) Correlation at obs. 7

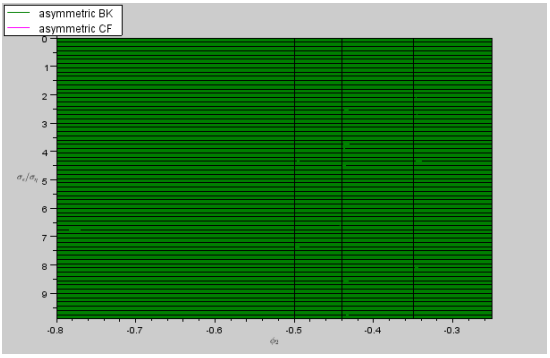

(f) view at 1.12(e) from the top

Fig. 1.12. Estimated correlation of the true and extracted cycles at time t, $\hat{\rho}(c_t, \hat{c}_t)$, by asymmetric BK and CF filters at observations number 9 to 7, counting from the end of the series

Note: Going further away from the center of the sample, the difference of the performance of the filters starts showing up, with the BK filter being persistently better.

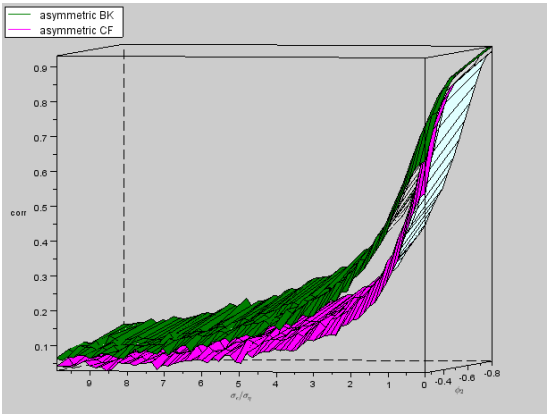

(a) Correlation at obs. 6

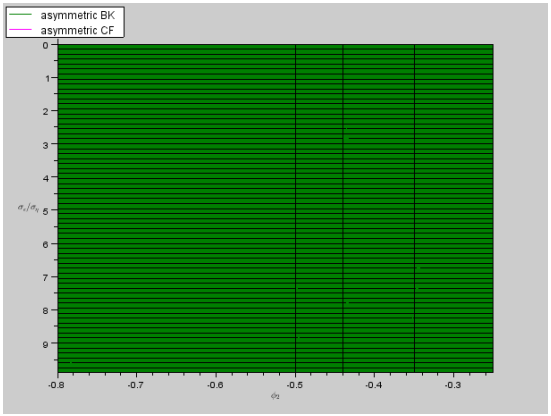

(b) view at 1.13(a) from the top

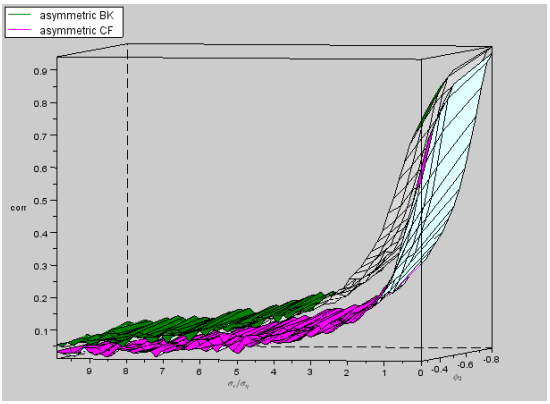

(c) Correlation at obs. 5

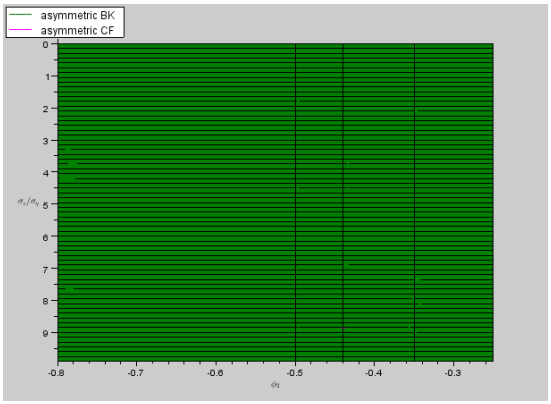

(d) view at 1.13(c) from the top

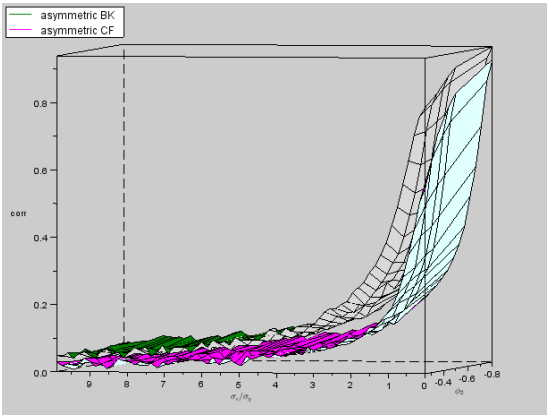

(e) Correlation at obs. 4

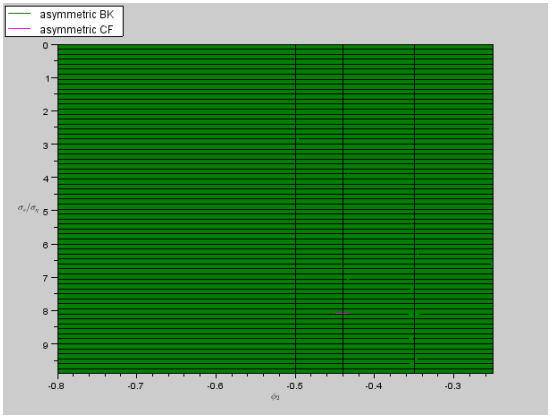

(f) view at 1.13(e) from the top

Fig. 1.13. Estimated correlation of the true and extracted cycles at time t, $\hat{\rho}(c_t, \hat{c}_t)$, by asymmetric BK and CF filters at observations number 6 to 4, counting from the end of the series

Note: Clearly, the BK filter is superior to the CF filter at observations close to the ends of the sample.

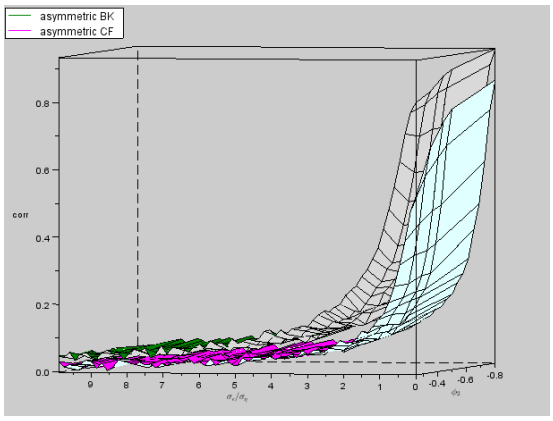

(a) Correlation at obs. 3

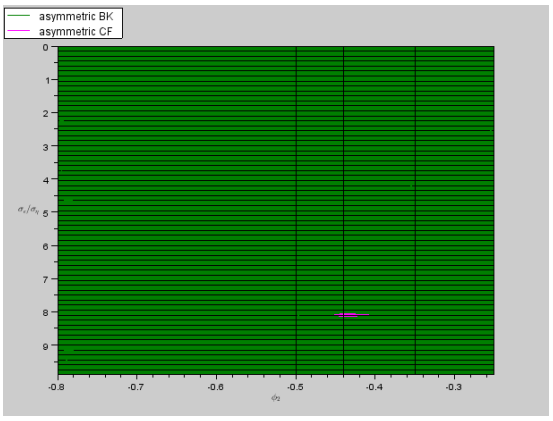

(b) view at 1.14(a) from the top

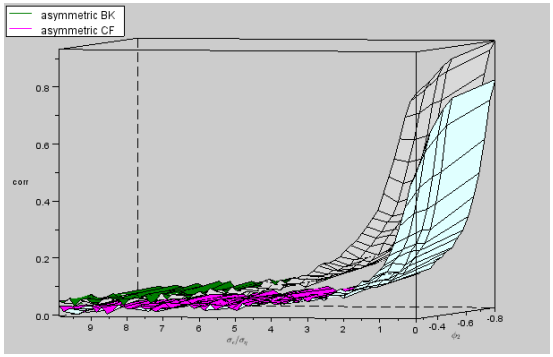

(c) Correlation at obs. 2

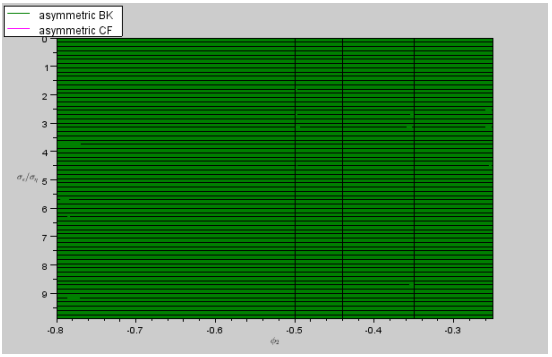

(d) view at 1.14(c) from the top

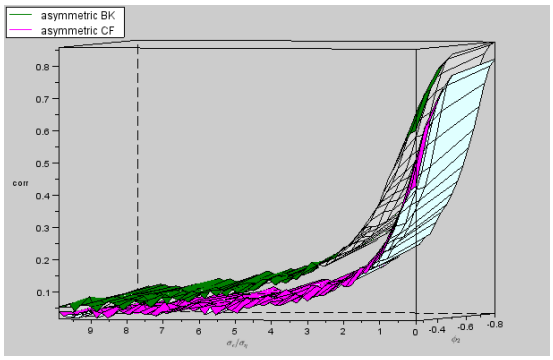

(e) Correlation at obs. 1

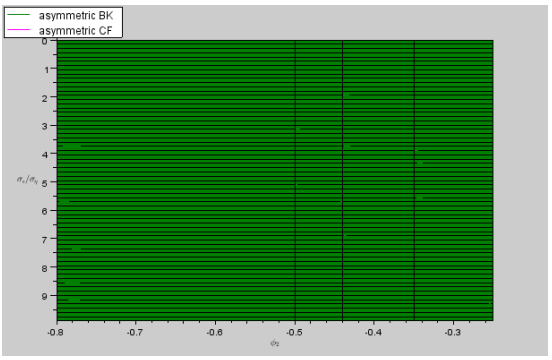

(f) view at 1.14(e) from the top

Fig. 1.14. Estimated correlation of the true and extracted cycles at time t, $\hat{\rho}(c_t, \hat{c}_t)$, by asymmetric BK and CF filters at observations number 3 to 1, counting from the end of the series

Note: If the considered DGP is relevant in practice, the asymmetric BK filter should be given preference over the asymmetric default specification of CF filter in real time signal extraction exercises.

Fig. 1.11 to 1.14 show the filters give close result at points closer to the center of

the sample. Indeed, at the center of the sample, where the asymmetric filters become symmetric, the correlation graph looks similar to Fig. 1.11(a), thus not shown here. As the estimation point approaches the end of the sample, filters become more asymmetric, and the difference in their performance becomes more obvious. Thus, the asymmetric filters perform roughly equally well at points that are at least about 3 years (for quarterly data) away from the end of the sample. Otherwise, the asymmetric BK filter becomes increasingly superior to the asymmetric CF filter for any cycle length and for any share of permanent component in the input signal considered in the simulation.

Specifically, and as a summary, for 'typical' macroeconomic series ($\sigma_\epsilon/\sigma_\eta = 3/4$, cycle length: 3.5 years), the ratio of correlations given by the ABK and CF filters is shown in Fig. 1.15, where the horizontal axis represents the point of signal extraction and T denotes the sample length.

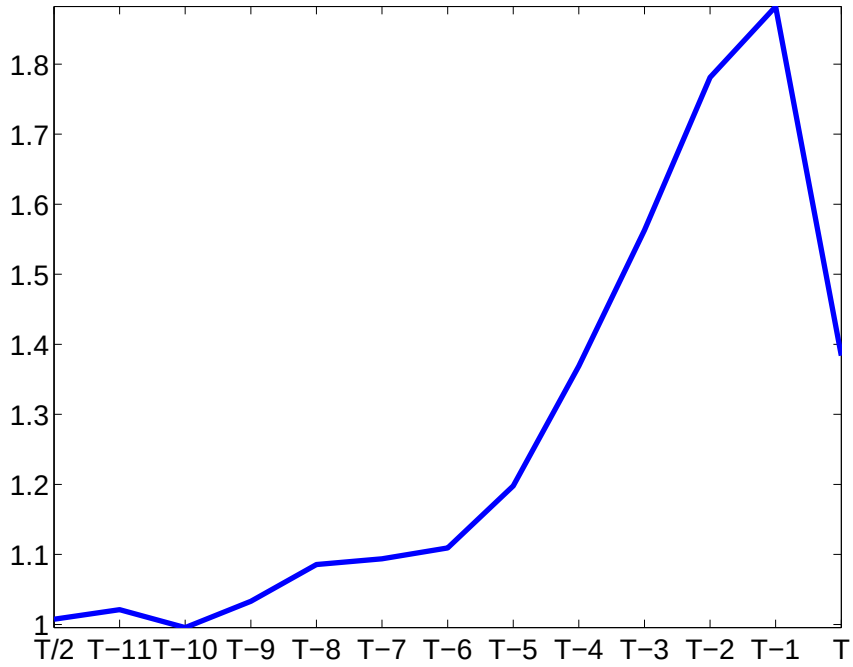

Fig. 1.15. The ratio of correlations given by the ABK and CF filters for 'typical' macroeconomic series ($\sigma_\epsilon/\sigma_\eta = 3/4$, cycle length: 3.5 years)

It shows that in the middle of the sample, $T/2$, both ABK and CF filters perform equally, i.e., the ratio of correlations is about unity. However, when the point of extraction moves to the end of the series, the ABK filter's relative performance increases, culminating at one observation from the end, where its performance (in terms of correlation with the target signal) is almost twice as high as that extracted by the CF filter. The relative performance decreases at the end-point of the sample to about 40% gain. (That decrease in the relative gain can be explained by the different mechanisms of making the filters extract nothing from the trend frequency - while all the coefficients of the ABK filter are adjusted by the same amount, the adjustment of the CF filter affects only the end-points of the filter.) The 40% gain at the end-point is still huge. Moreover, the ABK

filter's performance is smooth over the point of extraction, and the relative decrease of the performance at the end-point is due to the (rather huge) increase of the performance in the CF filter at the end-point. In practice, it is important that a method's performance is smooth over the point of extraction in order to avoid breaks in the extracted output; this is one more reason to prefer the ABK over the CF filter.

Thus, based on Fig. 1.11 - 1.14, 1.15, it is recommended to use the asymmetric BK filter rather than the asymmetric CF filter for the business-cycle frequency extraction in real time, i.e, at the very end of the sample, given that the considered DGP and metric are appropriate in practice.

Fig. 1.8 shows that, on average in the sample, the fixed-length symmetric filters show higher performance than their asymmetric counterparts. This result is somewhat unexpected, since Fig. 1.8 says the symmetry or the shorter length of the fixed-length filters are more important than the higher length of an asymmetric filter, for better performance. Since the asymmetric filters become symmetric in the middle of a sample, it is tempting to compare the performance of filters in the middlepoint of the sample. Thus, Fig. 1.16 shows correlation surfaces estimated at the middlepoint of the sample, where all filters are symmetric, but fixed-length filters are shorter than the asymmetric filters, with the latter spanning the whole sample length.

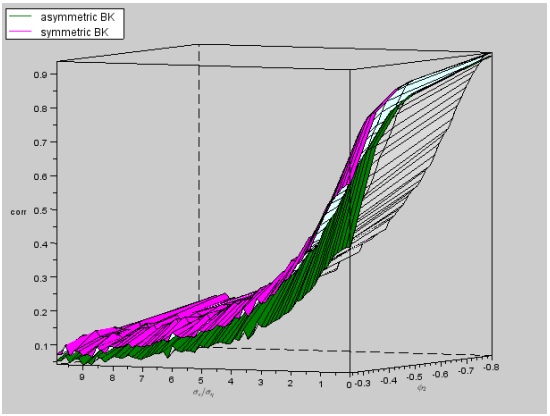

(a) fixed-length symmetric and asymmetric BK filters

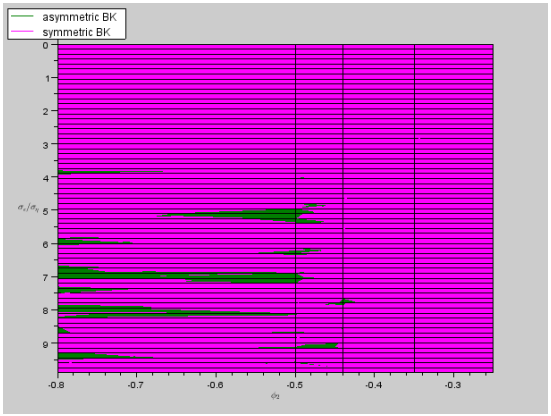

(b) view at 1.16(a) from the top

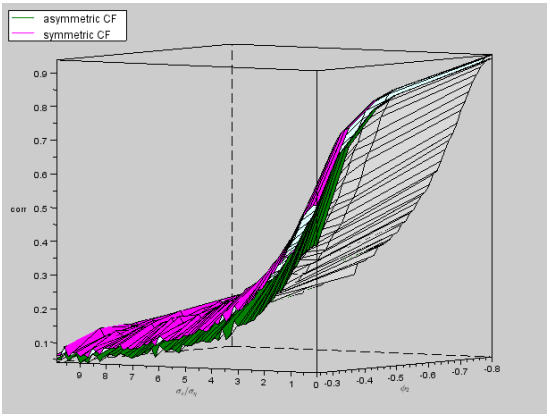

(c) fixed-length symmetric and asymmetric CF filters

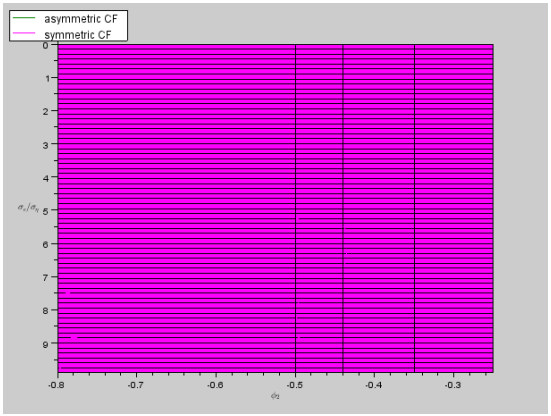

(d) view at 1.16(c) from the top

Fig. 1.16. Correlation surfaces estimated at the middlepoint of the sample, where all filters are symmetric, but fixed-length filters are shorter than the asymmetric filters, with the latter spanning the whole sample length

Note: The results show that the shorter filters outperform the longer ones. More simulations results indicate (not shown here) that, for BK filter, $K = 13$ gives the best results, outperforming the one with $K = 12$, regardless of wave length of the cycle, or the influence of permanent component in the input signal. For CF filter, $K = 12$ appears to be optimal, unless $\sigma_\epsilon/\sigma_\eta > 1$, for which $K = 13$ seems to yield higher $\hat{\rho}(c_t, \hat{c}_t)$.

The results show that the shorter filters outperform the longer ones, which is contrary to what most (know to me) literature takes as granted. More simulations results indicate (not shown here) that, for BK filter, $K = 13$ gives the best results, outperforming the one with $K = 12$, regardless of wave length of the cycle, or the influence of permanent component in the input signal. For CF filter, $K = 12$ appears to be optimal, unless $\sigma_\epsilon/\sigma_\eta > 1$, for which $K = 13$ seems to yield higher $\hat{\rho}(c_t, \hat{c}_t)$.

Fig. 1.17 and 1.18 show the estimated regression coefficients on the first and second lag, respectively, from fitting an AR(2) model on the cyclical component extracted by the four filters.

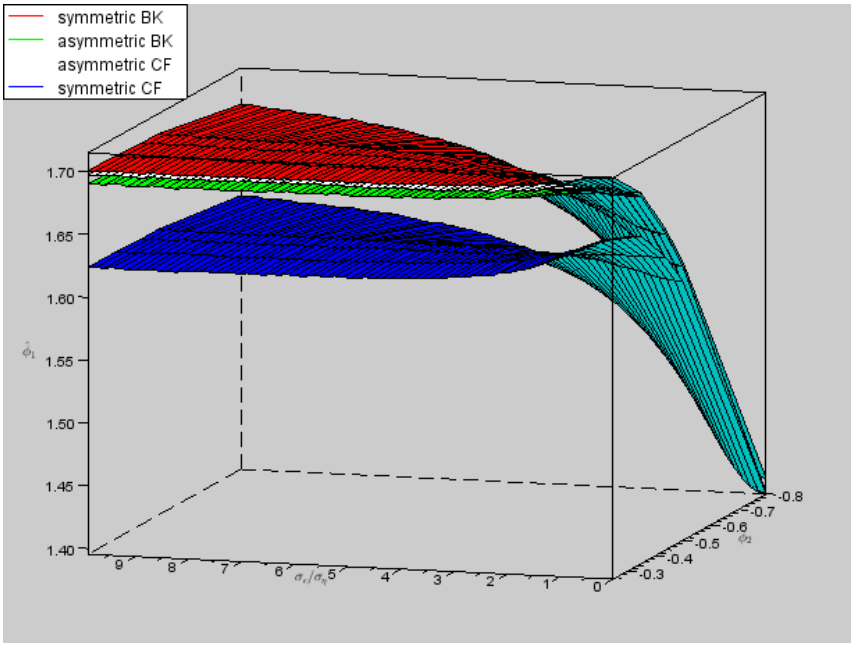

Fig. 1.17. Estimated ϕ_1 by the four filters for various ϕ_2 and $\sigma_\epsilon/\sigma_\eta$ values

Note: ϕ_1 is always overestimated.

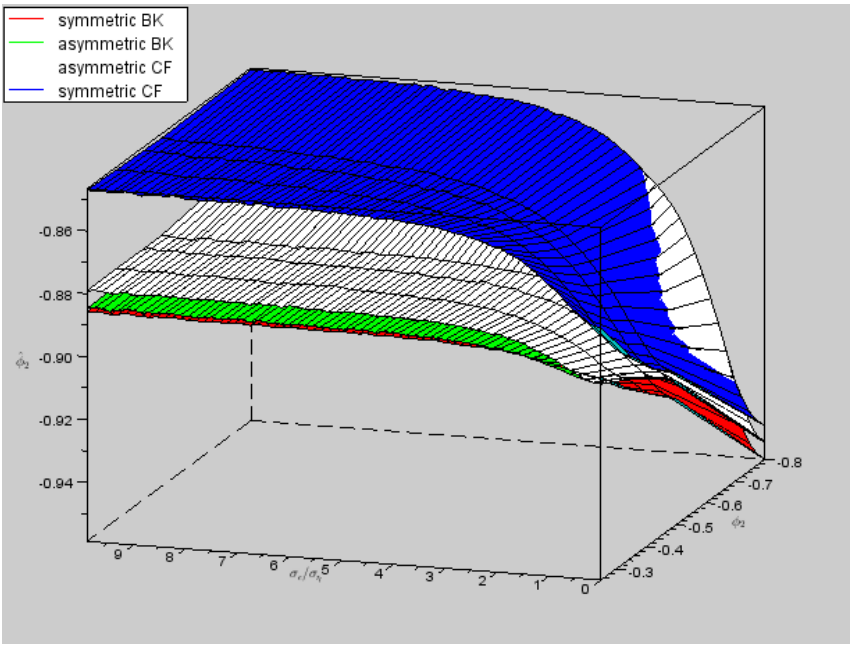

Fig. 1.18. Estimated ϕ_2 by the four filters for various ϕ_2 and $\sigma_\epsilon/\sigma_\eta$ values

Note: The estimated ϕ_2 converges to a constant as the influence of the permanent component on the input series increases, regardless of the true length of the cycle. While it is expected that $\hat{\phi}_1$ would be approximately constant, since ϕ_1 is always set to 1.2, it is not that plausible for $\hat{\phi}_2$ to converge to a constant, regardless of the value of ϕ_2. ϕ_2 is always underestimated.

Fig. 1.17 and 1.18 show both estimated regression coefficients converge to a constant as the influence of the permanent component on the input series increases, regardless of the true length of the cycle. While it is expected that $\hat{\phi}_1$ would be approximately constant,

since ϕ_1 is always set to 1.2, it is not that plausible for $\hat{\phi}_2$ to converge to a constant, regardless of the value of ϕ_2. The results also show that, regardless of the considered true length of the cycle and regardless of the influence of the permanent component in the input series, $\hat{\phi}_1 \in (1.4, 1.7)$ and $\hat{\phi}_2 \in (-0.96, -0.84)$, i.e, ϕ_1 is always overestimated and ϕ_2 is always underestimated.

Table 1.2 shows the length of the cycle extracted by the filters, when the influence of the permanent component in the input series is sufficiently high, i.e., about $\sigma_\epsilon/\sigma_\eta > 5$. In such case, the length of the extracted cycle is about constant, regardless of the true length of the cycle.

Table 1.2.
The true AR(2) parameters and cycle length, and the estimated AR(2) parameters and cycle length by the four filters, when the influence of the permanent component in the input series is sufficiently high, i.e., about $\sigma_\epsilon/\sigma_\eta > 5$

	ϕ_1	ϕ_2	Fundamental period of the cycle (yrs)
true	1.2	-0.25	$\approx \infty$
	1.2	-0.35	$\gg 8$
	1.2	-0.44	8.2
	1.2	-0.5	3.5
	1.2	-0.8	1.9
symmetric BK	1.699	-0.886	3.56
asymmetric BK	1.689	-0.884	3.48
asymmetric CF	1.696	-0.879	3.60
symmetric CF	1.623	-0.848	3.23

Note: In such case, the estimated AR(2) parameters and the length of the extracted cycle are about constant, regardless of the true length, or existence, of the cycle.

1.4 Conclusions

1. This chapter formally develops the asymmetric version of the Baxter-King filter for business-cycle frequency extraction at the end-points of univariate series.

2. Monte carlo simulation results show that the developed filter outperforms the most popular band-pass filter in macroeconomics - Christiano-Fitzgerald filter - within two years from the end-point; the correlation between the extracted and ideal signals at the end-point is increased by 40 per cent for typical macroeconomic series.

2 MULTIVARIATE FILTER FOR HIGH-DIMENSIONAL AND NOISY DATASETS

2.1 Unregularized multivariate direct filter approach: a new real-time indicator for the euro area GDP

This section proposes a new real-time unrevised indicator tracking medium-to-long-term component in the quarterly growth of the euro area GDP. The new indicator is based on recently developed real-time filtration methodology, the multivariate direct filter approach, applied to selected business and consumer survey and share price data. The new indicator is found to have led another established indicator, the Eurocoin, by about three months since mid-2009 and be about zero to one month ahead of and smoother than the Purchasing Managers Index. In addition to the euro area aggregate indicator, the section presents prototypical indicators for four biggest European Union economies - Germany, France, the United Kingdom and Italy. Overall, the described filter approach appears to have some merit in tracking business cycle developments.

2.1.1 Introduction

This section applies a recently developed real-time filtration methodology, i.e. multivariate direct filter approach (Wildi (2011)), to construct a new real-time unrevised indicator tracking medium-to-long-run component in the quarterly growth of euro area GDP.

The demand for real-time macroeconomic indicators exists mainly because many key macroeconomic variables are released with a considerable lag and subsequently revised in later releases. For example, the first rough estimate of the main macroeconomic aggregate, the gross domestic product (GDP), called 'flash GDP', is released only 45 days after the reference period in the European Union (EU), including the Euro Area (EA), and happens to be revised substantially. The first official release of the GDP in the EU is published only 65 days after the reference period, and even this first release is revised in subsequent releases. However, economic and financial agents are keen to have timely information on

the new developments in economy. Therefore, efforts are made in using timely information to capture the big picture on the overall economy in a more timely manner. For that matter, several indicators are available that try to map timely disaggregate information on the aggregate series, like the GDP. Particularly, there are several established real-time indicators tracking economic situation in the Euro Area, a few of which are subjectively selected below.

The (New) Eurocoin (Altissimo et al., 2010) is a replacement of its predecessor, (Old) Eurocoin (Altissimo et al., 2001), since May 2009. This indicator tracks medium-to-long-run component in quarterly growth of Euro Area GDP and is a result of a dynamic factor model applied on about 145 series.

The Euro Growth Indicator by EUROFRAME[1] is a monthly nowcast and a forecast of EA GDP two quarters ahead of official statistics. It is based on a bridge regression applied on timely survey and financial data.

A new project, Now-Casting[2], produces a nowcast of quarterly growth of the EA GDP on a weekly basis. Its real-time performance currently spans only the last three quarters, thus, it is premature to assess the quality of this indicator and, particularly, the added value from weekly - as opposed to monthly - updating frequency.

The Organization for economic cooperation and development, OECD, publishes a composite leading indicator for Euro Area (Arnaud and Hong, 2001) which is - at the time of writing this section - a revised monthly indicator targeting a lead in the business cycle (bandpass) of EA industrial production index. A drawback of a revised indicator, however, is that its revised values might wrongly suggest to its user about its real-time performance. Another nuance is that a bandpass excludes trend growth, which plays an important role in the existence and depth of recessions.

A quarterly IFO Economic Climate indicator of the Euro Area[3] is updated only quarterly which might be not as frequent as some economic agents would like.

The monthly Markit flash Euro Area Purchasing Managers Index[4] is one of the closely watched indices for the Euro Area economy because it is timely, unrevised, has straightforward methodology and is economically relevant. It is based on surveys of companies.

This section applies a recently developed multivariate real-time filtration technique on selected survey and stock price data to create a coincident, unrevised indicator for the Euro Area GDP and compares its performance to the one of the PMI and the Eurocoin. Various robustness checks are performed, as well as prototypical indicators for the biggest four European Union economies - Germany, France, the United Kingdom and Italy - are also presented. The main finding is that the new indicator is leading Eurocoin by about

[1] See www.euroframe.org

[2] See www.now-casting.com

[3] See http://www.cesifo-group.de/portal/page/portal/ifoHome/a-winfo/d1index/25indexweseuro

[4] See http://www.markiteconomics.com/Survey/Page.mvc/PressReleases

three months since mid 2009 and is about zero to one month ahead of, and smoother than, the PMI. Overall, the described filtration methodology is found to be able to provide somewhat better results in tracking business cycle developments than other widely used approaches.

The section is structured as follows. Subsection 2.1.2 reviews the filtration methodology used in the construction of the new indicator, Subsection 2.1.3 applies the filtration methodology to construct the new indicator, performs real-time comparison with Eurocoin and the PMI, as well as constructs prototypical indicators for the four biggest European Union economies, and Subsection 2.1.4 concludes.

2.1.2 Filtration methodology

This section is concerned with estimating a signal - a trendcycle or a business cycle - in real time. Denote y_T as the output of a symmetric, possibly bi-infinite filter, $\sum_{j=-\infty}^{\infty} \gamma_j L^j$, applied on input series x_T:

$$\begin{aligned} y_T &= \sum_{j=-\infty}^{\infty} \gamma_j L^j x_T \\ &= \sum_{j=-\infty}^{\infty} \gamma_j x_{T-j}, \end{aligned} \tag{2.1}$$

where L is called the lag or backshift operator. The filter in (2.1) is called the ideal filter and the filter output, y_T, is called the ideal filter output. Time series are finite in practice, therefore, the ideal filter is infeasible. A practitioner might use a finite symmetric filter as an approximation to (2.1) in the middle of the time series but even the symmetric approximation is infeasible at the very end of the sample, i.e., in real time. A real-time estimate of y_T, given a finite input $\{x_1, \ldots, x_T\}$, can be written as

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

It is a well-know fact that the filter in (2.2) generally possesses a nontrivial phase shift, i.e., it's output is lagging in time. In order to define the phase shift and, in general, the effect of a filter, 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. A generally complex number, $\Gamma(\omega)$, can be represented in polar coordinates as $\Gamma(\omega) = A(\omega)\exp(-i\Phi(\omega))$, where $A(\omega) = |\Gamma(\omega)|$ is called the amplitude, and $\Phi(\omega) = -\arg(\Gamma(\omega))$ is called the phase. The effect of a filter - as represented by its transfer function - can be summarized by the effect of amplitude and phase functions. The amplitude function, $A(\omega)$, can be interpreted as the weight - amplification or damping - attributed by the

filter to the input signal at frequency ω. The phase function, $\Phi(\omega)$, can be interpreted as a shift function of the amplified or damped signal at frequency ω. This section is concerned with obtaining a one-sided filter output which would be a good - as good as it can be in real time - approximation to the ideal output. Consider stationary processes because of easier exposition and because this section implements the filter on integrated of order zero, I(0), processes. Generalization to integrated processes is straightforward by using pseudo spectral estimates and filter constraints at frequency zero (Wildi, 2008). 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:

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

where $H(\omega)$ is the unknown spectral distribution of x_T. A finite sample approximation of the MSFE, (2.3), is

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

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

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

see Brockwell and Davis, 1987, Ch. 10 for the reason for w_k. $S(\omega_k)$ in (2.4) 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, (2.3). Therefore, this section uses a 'sufficient statistic' - periodogram, $I_{Tx}(\omega_k)$ - as $S(\omega_k)$ in (2.4):

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

Minimizing expression (2.4) yields the real-time filter output optimally approximated to the ideal output in mean squared error sense. As noted above, 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 feasible filter is assumed to be a bandpass and the spectral estimate, $S(\omega_k)$ be the spectrum of the random walk), TRAMO/SEATS seasonal adjustment methodology using Wiener-Kolmogorov filter (ARIMA-based spectral estimates) or a Hodrick-Prescott filter (which can be interpreted to be a Butterworth-type filter, assuming the input following an ARIMA(0,2,2) process, and its amplitude function being of a particular truncated bell-shaped form, see King and Rebelo (1993) and Maravall and Rio (2001)). Yet, Wildi (2008) proposes a customized version of (2.4) which this section finds to be useful in creating a real-time indicator, and which is described in the next subsection.

Univariate direct filter approach

Wildi (2008) proposes a customized version of (2.4), a part of which is implemented in this section. Rewrite discrete version MSFE, (2.4), as

$$\frac{2\pi}{T}\sum_{k=-[T/2]}^{[T/2]} w_k|\Gamma(\omega_k)-\hat{\Gamma}(\omega_k)|^2 I_{Tx}(\omega_k)W(\omega_k), \tag{2.7}$$

which is identical to (2.4) for $W(\omega_k) := 1$. However, a more general version of $W(\omega_k) := W(\omega_k, expw, cut)$ can be written as

$$W(\omega_k, expw, cut) = \begin{cases} 1 & \text{if } |\omega_k| < cut \\ (1+|\omega_k|-cut)^{expw} & \text{otherwise,} \end{cases} \tag{2.8}$$

which collapses to unity for $expw = 0$, in which case the classical mean squared optimization, (2.4), 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 (2.4) for positive *expw*.

Univariate filters can be useful in case the target variable of interest, x_t, is timely; however, it is usually not the case for macroeconomic series. For example, one of the key macroeconomic variables, the gross domestic product (GDP), is released with a delay and is revised in subsequent releases. Therefore, a practitioner might be interested in using other economic variables with shorter delays in creating benchmark indicators for the GDP of some release. This environment demands for the ability for a multiple-series analysis, which in the particular context relates to the next subsection.

Multivariate direct filter approach

The above univariate customized filter has been generalized to a multivariate filter in Wildi (2011). Rewrite univariate minimization problem, (2.7), with the discrete Fourier

transform (DFT), $\Xi_{Tx}(\omega_k)$:

$$\begin{aligned}
&\frac{2\pi}{T}\sum_{k=-[T/2]}^{[T/2]} w_k|\Gamma(\omega_k)-\hat{\Gamma}(\omega_k)|^2 I_{Tx}(\omega_k)W(\omega_k) \\
&=\frac{2\pi}{T}\sum_{k=-[T/2]}^{[T/2]} w_k|\Gamma(\omega_k)\Xi_{Tx}(\omega_k)-\hat{\Gamma}(\omega_k)\Xi_{Tx}(\omega_k)|^2 W(\omega_k), \qquad (2.9)
\end{aligned}$$

where

$$\Xi_{Tx}(\omega_k)=\sqrt{\frac{1}{2\pi T}}\sum_{t=1}^{T} x_t \exp(-it\omega_k). \qquad (2.10)$$

In addition to the filter output, y_t, and the corresponding input, 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 obtained with a univariate filter. Then, the second expression in the modulus on the second line of (2.9), $\hat{\Gamma}_X(\omega_k)\Xi_{Tx}(\omega_k)$, becomes

$$\hat{\Gamma}_X(\omega_k)\Xi_{Tx}(\omega_k)+\sum_{n=1}^{m}\hat{\Gamma}_{z_n}(\omega_k)\Xi_{Tz_n}(\omega_k), \qquad (2.11)$$

where

$$\hat{\Gamma}_X(\omega_k)=\left(\sum_{j=0}^{L} b_{xj}\exp(-ij\omega_k)\right)\Xi_{Tx}(\omega_k) \qquad (2.12)$$

$$\hat{\Gamma}_{z_n}(\omega_k)=\left(\sum_{j=0}^{L} b_{z_nj}\exp(-ij\omega_k)\right)\Xi_{Tz_n}(\omega_k) \qquad (2.13)$$

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 (2.9) can be written as

$$\frac{2\pi}{T}\sum_{k=-[T/2]}^{[T/2]} w_k\left|\left(\Gamma(\omega_k)-\hat{\Gamma}_x(\omega_k)\right)\Xi_{Tx}(\omega_k)-\sum_{n=1}^{m}\hat{\Gamma}_{z_n}(\omega_k)\Xi_{Tz_n}(\omega_k)\right|^2 W(\omega_k). \qquad (2.14)$$

The section uses the filter obtained by minimizing (2.14) subject to filter parameters and the first order constraint at frequency zero:

$$b_1^n+b_2^n+\ldots+b_L^n=w^n \qquad (2.15)$$

imposes amplitude constraint in frequency zero according to $\hat{\Gamma}_{z_n}(0)=w^n$ ($\hat{\Gamma}_{z_n}=\hat{\Gamma}_x$ if $m=0$), and w^n is set to unity.

2.1.3 Real-time indicator for the euro area GDP

The data

The target. The target is medium-to-long-run fluctuations in quarterly growth of seasonally adjusted Euro Area gross domestic product (GDP) in chain-linked prices, as published by Eurostat. Specifically, the filter is set to target an ideal lowpass of quarterly growth of seasonally adjusted EA GDP with cut-off wave length 12 months (Fig. 2.1).

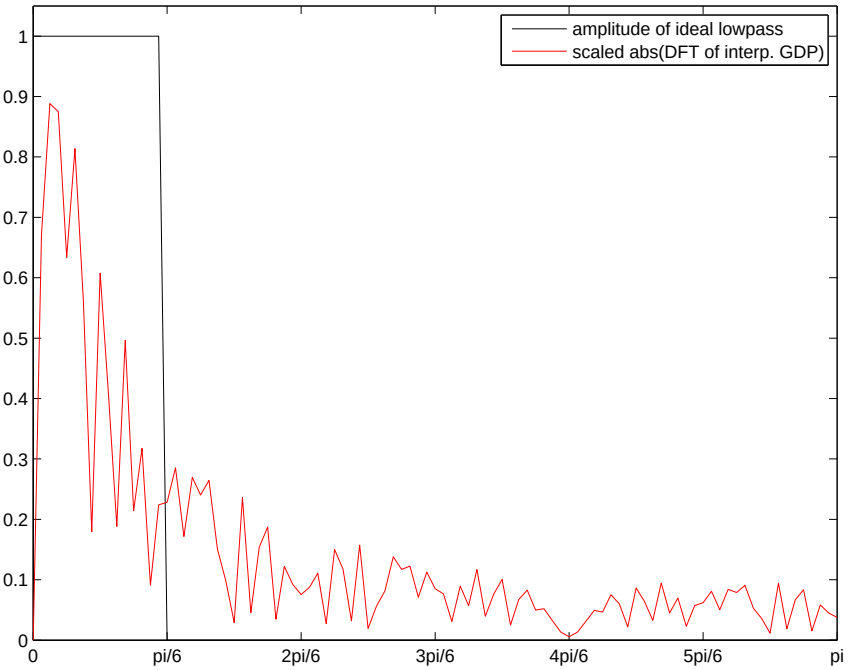

Fig. 2.1. Amplitude of ideal lowpass and a scaled absolute value of DFT of differenced interpolated GDP

Note: The latter is zero in frequency zero because GDP series is demeaned.

The usually defined minimum length of a business cycle is 18 months (Burns and Mitchell, 1946; Baxter and King, 1999; Christiano and Fitzgerald, 2003). However, passing higher-frequency content through a filter allows for both faster turning point detection and a closer amplitude fit. The quarterly GDP data are taken from 1995Q1 onwards. The data are linearly interpolated to monthly frequency, logged, regularly differenced and demeaned before their spectral content enters the filter. Alternatively, the interpolation and differencing can be performed in frequency domain but the results show no obvious improvement with the latter compared to such data transformation in time domain.

Explanatory variables. Select monthly confidence indicators published by DG Ecfin and the US share price index published by Eurostat are used as explanatory variables for the Euro Area GDP. DG Ecfin data are usually published at the end of reference

month, except for December for which data are published in early January. 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[5]. The above-mentioned considerations make Ecfin data comfortable for real-time filtration. The selected business and consumers confidence data are: production trend observed in recent months (industry), assessment of order-book levels (industry), assessment of stocks of finished products (industry), production expectations for the months ahead (industry), employment expectations for the months ahead (industry), confidence indicator in construction, confidence indicator in retail trade, and consumers confidence indicator. The choice of the indicators is based on economic relevance and data availability.

Other Ecfin survey data have been tested and found not to add substantial quality to the indicator. Seasonally adjusted Ecfin data are used (Fig. 2.2).

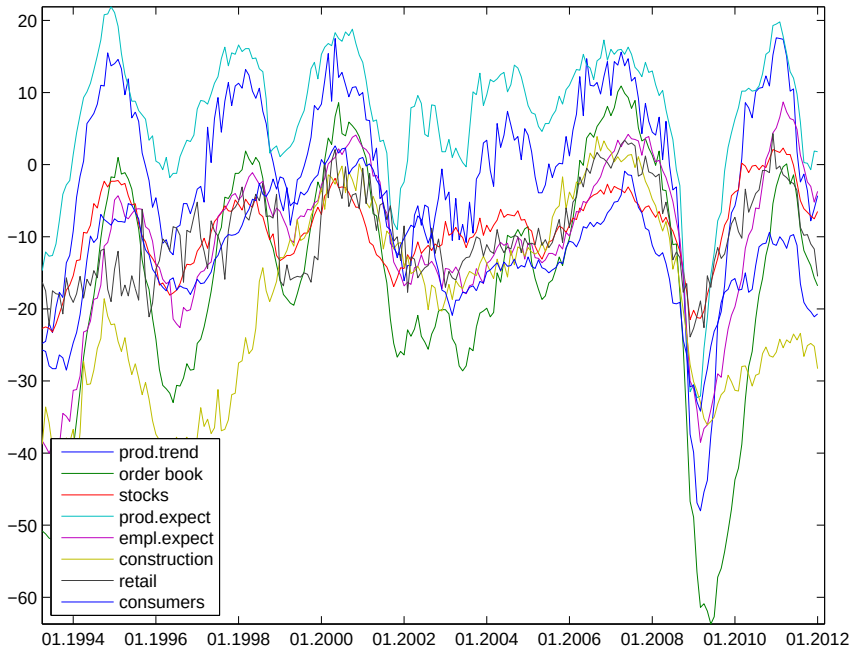

Fig. 2.2. DG Ecfin survey data

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

[5]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

Table 2.1.

The list of data used for the construction of the indicators

Variable	Source	EA	Germany	France	the United Kingdom	Italy	Transformation
Gross domestic product, chain-linked, SA	Eurostat	x	x	x	x	x	$\Delta \log$,lin.interp.
Production trend observed in recent month (industry), SA	DG Ecfin	x	x	x	x	x	Δ
Assessment of order-book levels (industry), SA	DG Ecfin	x	x	x	x	x	Δ
Assessment of stocks of finished products (industry), SA	DG Ecfin	x	x	x	x	x	Δ
Assessment of stocks of finished products (industry), SA	DG Ecfin	x	x	x	x	x	Δ
Production expectations for the months ahead (industry), SA	DG Ecfin	x	x	x	x	x	Δ
Confidence indicator in construction, SA	DG Ecfin	x	x	x	x	x	Δ
Confidence indicator in retail trade, SA	DG Ecfin	x	x	x	x	x	Δ
Consumers confidence indicator, SA	DG Ecfin	x	x	x	x	x	Δ
Confidence indicator in services, SA	DG Ecfin		x				Δ
The EA share price index	Eurostat		x		x	x	$\Delta \log$
The US share price index	Eurostat	x	x			x	$\Delta \log$

Note: The confidence indicators and the GDP data corresponds to the one of the particular region.

All explanatory variables are taken from 1995M1 onwards, regularly differenced and standardized to zero mean and unit variance. Table 2.1 lists the data and their transformations.

Filter specification

Filter dimension. The filter is tested for various dimension sizes. Estimated mean squared filter error (EMSFE), expression (2.14), is the implemented statistic for the goodness of fit, according to which low-dimensional filters often match the target DFT the best. 1-variable and 2-variable filters show about the same EMSFE, while 3-variable filter and 4-variable filter EMSFEs are about 3 and 9 times higher, respectively. Thus, 2-variable filter is chosen. Although the GDP series itself could be filtered, it is chosen not to do so as the GDP series is substantially lagged and revised, and appears not to be useful in prompt real-time filtration. Therefore, only the spectral content of the GDP series is fed to the filter and the monthly variables are set to target that spectral content. Thus, a 2-variable filter means that two monthly variables are filtered to target the latent output of the ideal filter that would be applied to the GDP series.

Filter length is set to depend on the number of input variables to be filtered to control for overparameterization. The particular 2-dimensional filter is set to be 38 observations long.

Filter constraint. First-order integration constraint is imposed by setting filter amplitude to unity at frequency zero. Since all the input series are demeaned and are I(0), the first order integration constraint ensures a close tracking of the scale of the filter output to that of the target signal.

Noise suppression. Extra suppression of high-frequency content in the stopband is implemented with positive *expw* parameter as described in the previous section. As a result, the filter output is smoother than the one of the classical mean squared filter error problem.

Cross-sectional aggregation. Given that the considered set of explanatory variables contains more variables than the selected filter dimension, all possible combinations of variables are entered in the filter, such that the final filter output is combined 2-variable filter outputs weighted by weights proportional to inverse EMSFE. Alternatively, principal components (Stock and Watson, 1998) could be used to shrink the dimension of the data set to a few factors but the results show no obvious improvement in this regard.

Output re-scaling and adding the mean back. Filter output obtained thus far is then rescaled to the variance of the output of 31-months long finite symmetric filter

for the time span available at the particular real-time estimation moment. Although the first order filter restriction ensures to some extent that the scale of the filter output be the same as the target signal, such rescaling, nonetheless, is harmless and is proved to be useful. Finally, the mean of the output of the finite symmetric filter is added to the filter output. One could use the mean of the GDP series, instead, but since the latter is more volatile in real-time setting than the output of the finite symmetric filter, the latter is preferred.

Averaging along time dimension. At this point, one can produce the real-time indicator. However, a couple of issues emerge. If filter coefficients are updated every time new data become available but the historic values of the indicator are not updated, then the noise generated by the changing estimated latent level of the target signal can suppress the estimated signal. Fig. 2.3 shows an indicator resulted from such filter coefficient re-estimation every quarter - the noise induced by frequent coefficient update suppresses the estimated target signal.

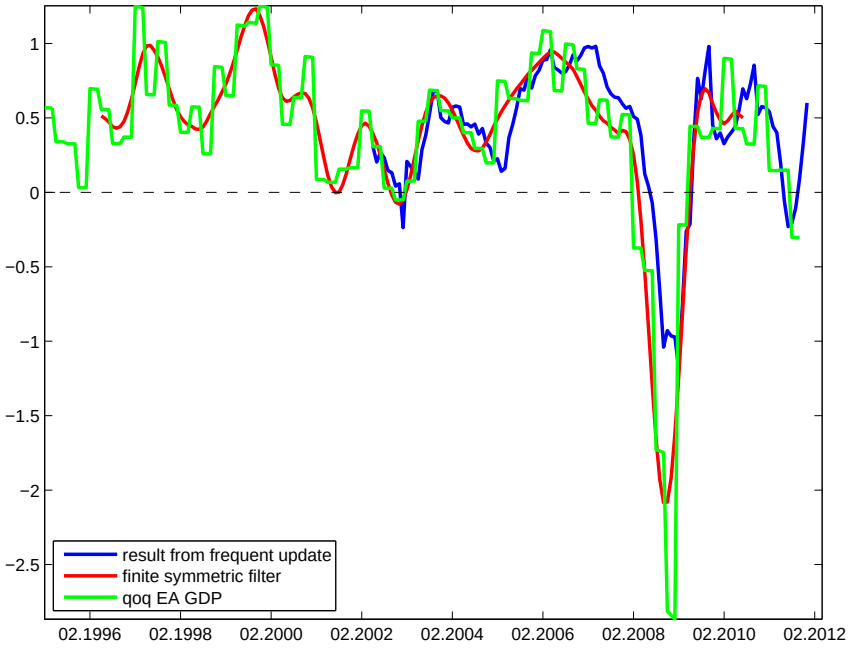

Fig. 2.3. Filter output resulted from frequent filter coefficient update

Note: Also shown EA GDP quarterly growth and the output of a 31-months long finite symmetric trendcycle filter of qoq GDP.

In this light, one might choose an in-sample span yielding most satisfactory out-of-sample performance and fix coefficients. In this case, the real-time estimate is smooth. A real-time recession indicator for the US economy, USRI (Wildi, 2009), is based on the latter approach. Yet, another issue emerges - sooner or later the filter coefficients might need an update due to possible structural changes in the economy.

The above issues are solved with the following filter averaging along the time dimension. Update the filter coefficients every time new GDP data become available. The final indicator is an average of all filter outputs created up to the estimation period. Furthermore, due to revisions of GDP series, some of its observations at the end of series - call them 'gap' - are left out of sample. That is, denote the whole sample length by 'len'. The maximum in-sample period spans observations $1 : len - gap$. Iteratively estimate filter coefficients for reduced in-sample length, $1 : len - gap - 3$, $1 : len - gap - 6$, etc. until some minimum in-sample length, but estimate filter outputs for the whole sample span, 'len'. Take an arithmetic average[6] of all historically estimated indicators. More indicators are averaged as more data come in but the maximum aggregation time span is set to five years in order to take into account possible structural changes. This updating mechanism overcomes the problem induced by the level of target series and at the same time robustifies signal detection, and incorporates new information/possible structural changes as new data become available. Setting *gap* to 9 or 15 months does not make much difference but insures against GDP revisions at the end of series and ensures that the target is about the 'final' GDP as opposed to a 'first-release'. The presented indicator is a result of setting 'gap' to nine months.

Such mechanism is performed for the last ten years. Out-of-sample time span is limited by the filter and sample lengths. The resulting filter output (blue) is shown in Fig. 2.4, along with quarterly growth of EA GDP (green) and the output of a 31-months long finite symmetric trendcycle filter (red).

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

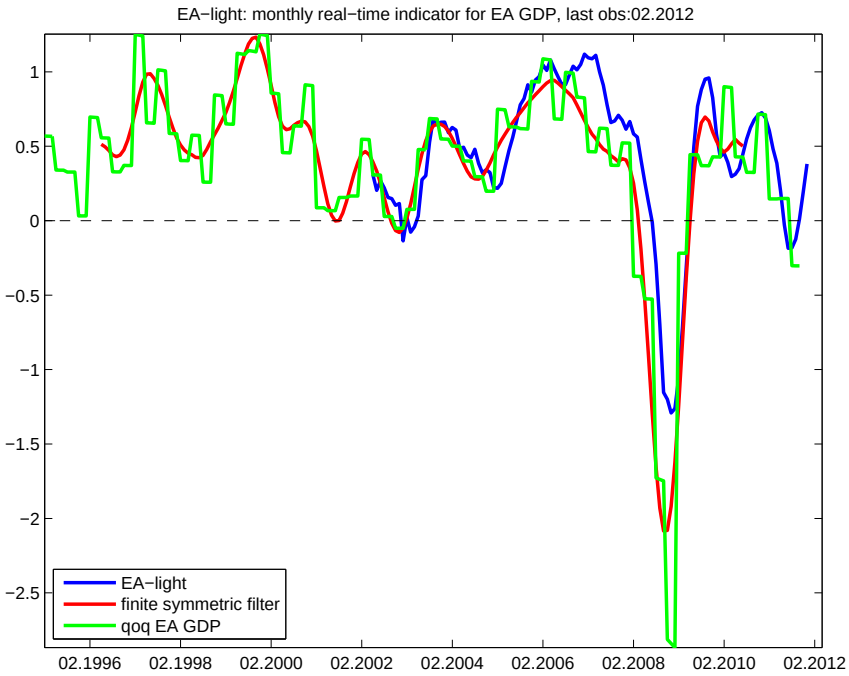

Fig. 2.4. The filter output

Note: Also shown EA q-o-q GDP and the output of a 31-months long finite symmetric filter.

Fig. 2.4 shows that the new indicator is slightly lagging in a historic perspective, and that its amplitude is comparable to that of the output of the finite symmetric filter except for the great recession episode. A visual observation suggests that the indicator gets slightly faster after the great recession episode, when it appears to be coincident with the GDP growth. Methodologically, the indicator is constructed as a coincident indicator. Furthermore, the indicator gets slightly smoother over time. The improvement in both the speed and smoothness over time could be explained by the increasing data length available to the filter. Moreover, the smoothness performance is partly due to the filter averaging along the time dimension - the first observation of the indicator in early 2002 is a result of a single cross-sectionally aggregated filter output; the next indicator values are averages of increasing number of filter outputs along the time dimension until five years pass, when only the filter outputs over the last five years are averaged to account for the effect of different phases of a business cycle, as well as possible structural changes, on the indicator's performance.

Robustness checks

Adjustment for changing volatility. The implemented filter, like its special variants - e.g., Baxter-King (Baxter and King, 1999), Christiano-Fitzgerald (Christiano and Fitzgerald, 2003), Hodrick-Prescott (Hodrick and Prescott, 1997), Wiener-Kolmogorov (Kolmogorov, 1941; Wiener, 1949) - assumes inputs with constant volatility. However,

economic data might be subject to changing volatility. Filter assuming inputs with constant volatility but applied to inputs with changing volatility might yield spurious output.

In this light, GARCH(1,1) (Bollerslev, 1986) is applied to Ecfin data as a simple check for changing volatility. Variance at period t, h_t, of a demeaned input variable x_t is modeled with GARCH(1,1) process:

$$h_t = \alpha_0 + \alpha_1 x_{t-1}^2 + \beta_1 h_{t-1}, \tag{2.16}$$

where $\alpha_0 > 0$, $\alpha_1 \geq 0$, $\beta_1 \geq 0$ and assumed $x_t \sim N(0, h_t)$. Kolmogorov-Smirnov (KS) test for normality (see, e.g., Massey, 1951) fails to reject the null hypothesis of (standard) normal distribution for all input variables at 5% significance. However, it is well known that KS test has small power for small sample sizes. An alternative, Jarque-Bera test (Jarque and Bera, 1987) rejects the null of normality for 5 of 8 input series, at 5% significance. Therefore, the assumption $x_t \sim N(0, h_t)$ appears not to hold. Nonetheless, GARCH(1,1) estimated by quasi-maximum likelihood is found to be generally consistent under mild assumptions (Bollerslev and Wooldridge, 1992).

Indeed, Ecfin data are found to be subject to heteroskedasticity (see Fig. 2.5 for estimated changing standard deviations).

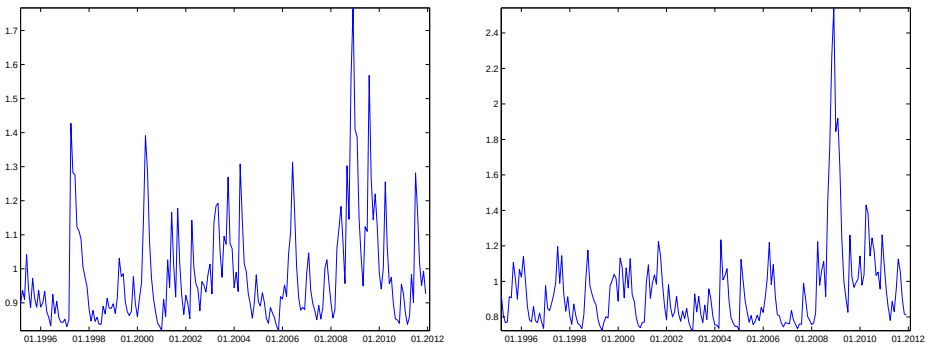

(a) prod.trend

(b) order book

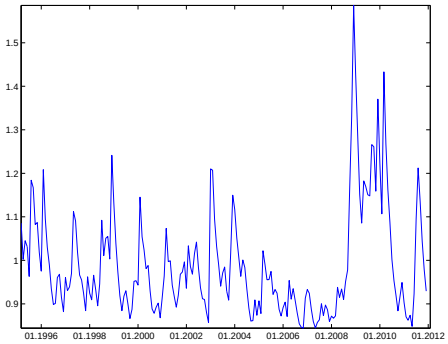

(c) stocks

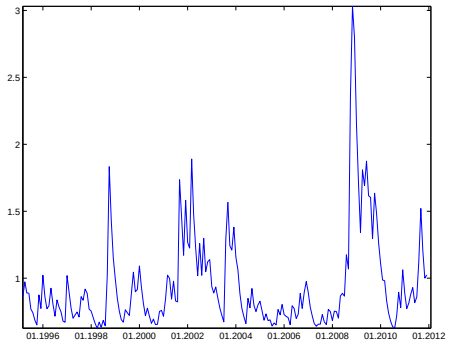

(d) prod.expect

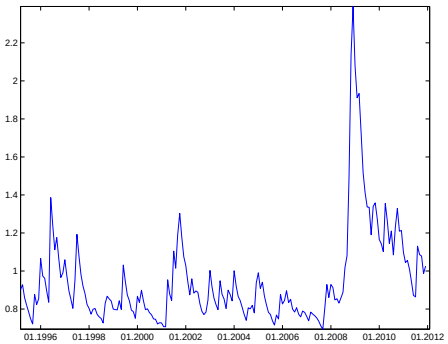

(e) empl.expect

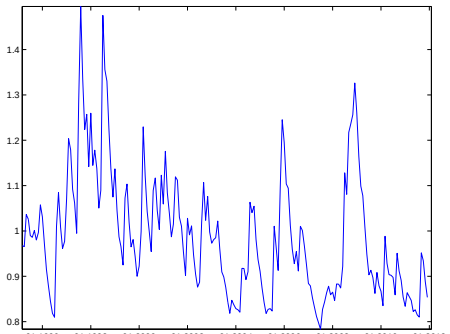

(f) construction

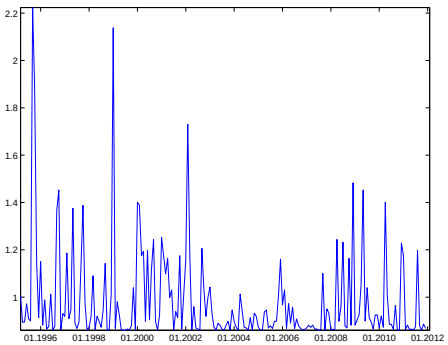

(g) retail

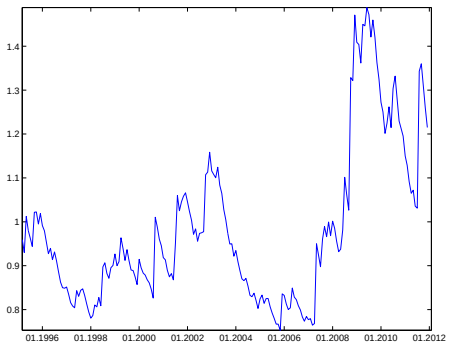

(h) consumers

Fig. 2.5. GARCH(1,1)-estimated standard deviations of Ecfin survey data

Therefore, GARCH(1,1)-normalized data are fed to the filter. However, GARCH-

augmented filter slightly looses to the non-GARCH one with respect to the estimated filter error as well as with respect to the amplitude fit during the great recession episode. Thus, GARCH normalization of input data is not implemented in this section, while the issue of the effect of the heteroskedasticity on the filter output remains.

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

Both seasonally adjusted and unadjusted survey data were tested as inputs. Moreover, seasonal adjustment that excluded a stable seasonal component, exclusively, was tested. Preliminary results have shown that the use of seasonally adjusted survey data - by Dainties - for filter inputs gives the most satisfactory results. Another nuance is whether to use seasonally adjusted or unadjusted GDP for spectral estimates. If the target is lowpass with cut-off to the left of $\pi/6$, both could be used without much difference in outcome. However, since the cut-off is set to the seasonal frequency, $\pi/6$, seasonally adjusted GDP is used for the target spectral estimate.

First order constraint. The above specification implements the first-order integration constraint that ensures that the scale of output is compatible with the scale of the target. This is because all series are demeaned, so the I(1) constraint tackles scale/variance fitting. An alternative - no-integration - specification was tested as well. Although the alternative works well in-sample, it yields inferior real-time performance with respect to scale fitting compared to the I(1) specification. Specifically, the amplitude fit during the great recession episode is clearly worse than that for the I(1)-constrained filter.

Data revisions. Seasonal adjustment (SA) method - Dainties - applied to Ecfin survey data does not induce revisions, as opposed to the results of the most popular SA methods - X12[7] and TRAMO/SEATS[8]. However, the last month of the non-adjusted data might be slightly revised thus inducing revisions also in SA data. These revisions have magnitude of about 0.1 in absolute value for unadjusted data, if they occur (some series have been found to be not revised). The greatest revisions are observed for the 'construction confidence indicator' series - of about 0.5 in absolute value. DG Ecfin does not provide vintage data publically, however, DG Ecfin was very kind to provide the vintages of non-adjusted data upon request. Results show that the output of the implemented filtering technique is proved to be robust against these minor revisions.

[7]See http://www.census.gov/srd/www/x12a/

[8]See http://www.bde.es/webbde/es/secciones/servicio/software/econom.html

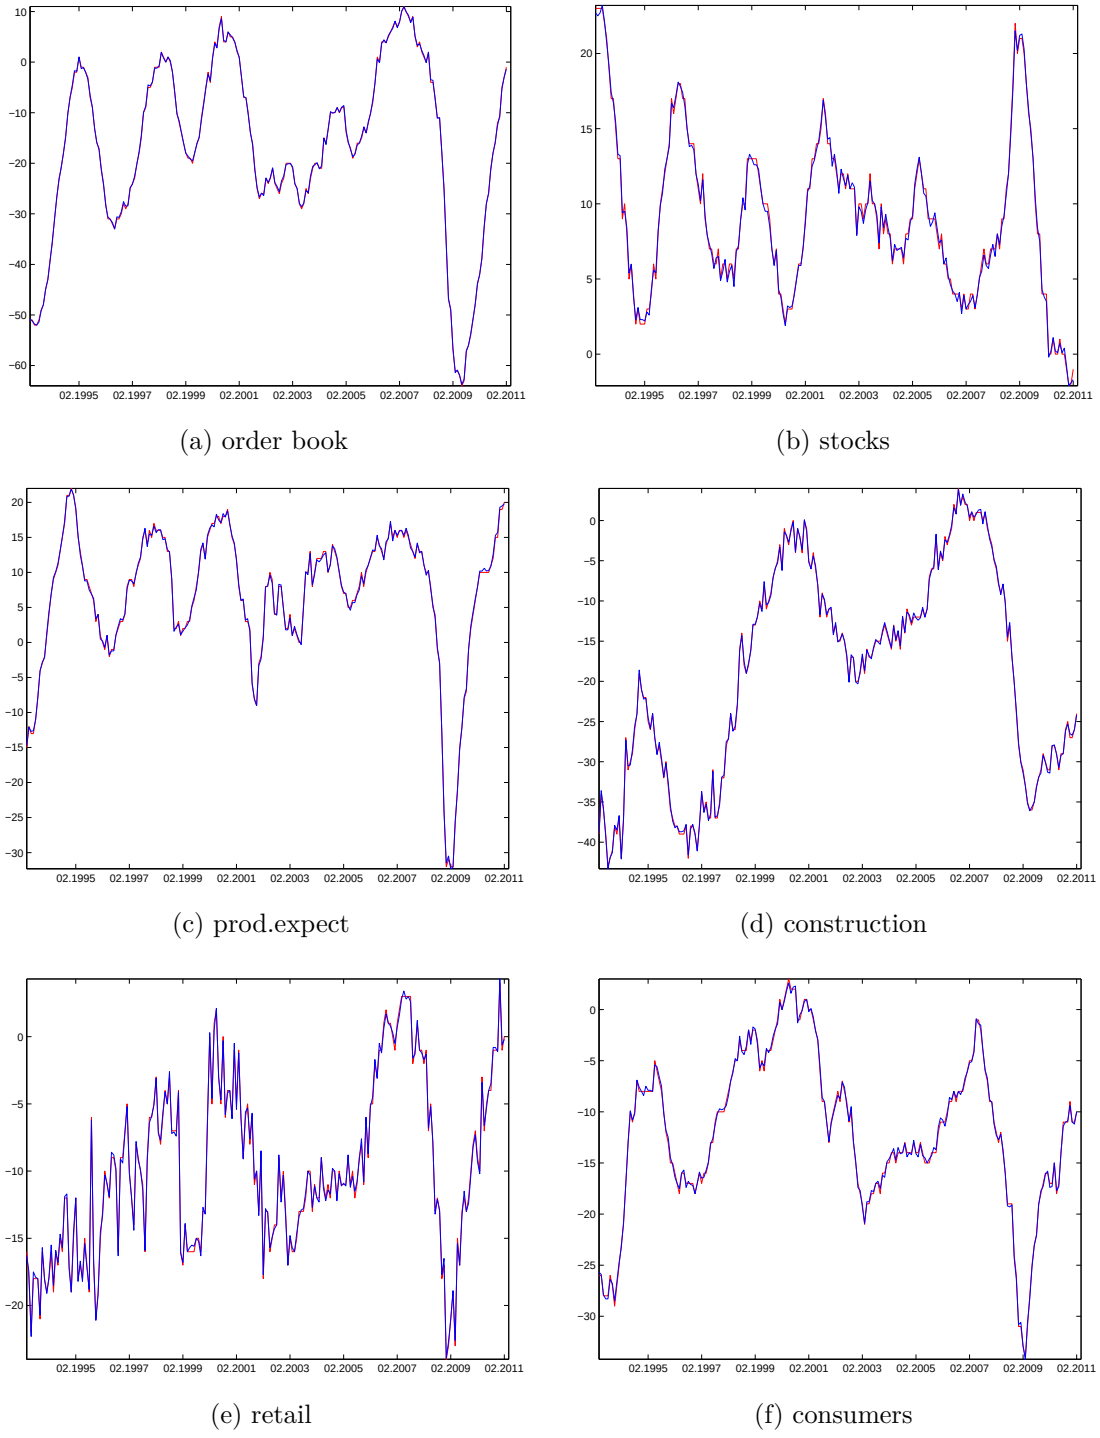

(a) order book (b) stocks

(c) prod.expect (d) construction

(e) retail (f) consumers

Fig. 2.6. Rounded first-release (red) vs final vintage (blue) of six Ecfin survey data series

As another vintage data source, the real time data base[9] (RTDB) maintained by the ECB provides vintage data for six out of the eight seasonally adjusted variables used in this section until Feb 2011. Unfortunately, the provided data are rounded to zero digits after comma. Nonetheless, comparing 1st release vintages with the latest vintage results in estimated mean deviation being zero, mean absolute deviation being about 1/4, and

[9]See http://sdw.ecb.europa.eu/browse.do?node=4843526

maximum absolute deviation ranging from 0.5 to 0.7, which basically reflects the effects of rounding (see Fig. 2.6). The filter output from the first release inputs is indistinguishable - with a naked eye - from the one with final release data, thus not presented.

Another source of potential data revision is revisions in GDP data. However, the solution to a revised GDP series is straightforward - GDP series itself is not filtered; instead, only its spectral content is fed to the filter; therefore, a robust filter output is obtained by feeding the filter with estimated spectra of the final/revised GDP by dropping several observations from the end of the series. That is, the described method does not use uncertain 1st releases, and dropping last 3 or 5 quarters from the estimation routine gives similar results.

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

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

Passband specification. The defined passband is lowpass with a cut-off frequency $\pi/6$ which corresponds to yearly frequency. Setting cut-off to the commonly defined upper bound of business cycle frequency - 18 months - makes the indicator slightly slower, and amplitude fit slightly worse. Therefore, keeping cut-off low allows for higher-frequency content passing through that speeds up turning point detection and improves amplitude fit.

The effect of high-frequency noise suppression. Fig. 2.7 shows the effect of the high-frequency noise suppression in the rightmost stopband by the customized filter.

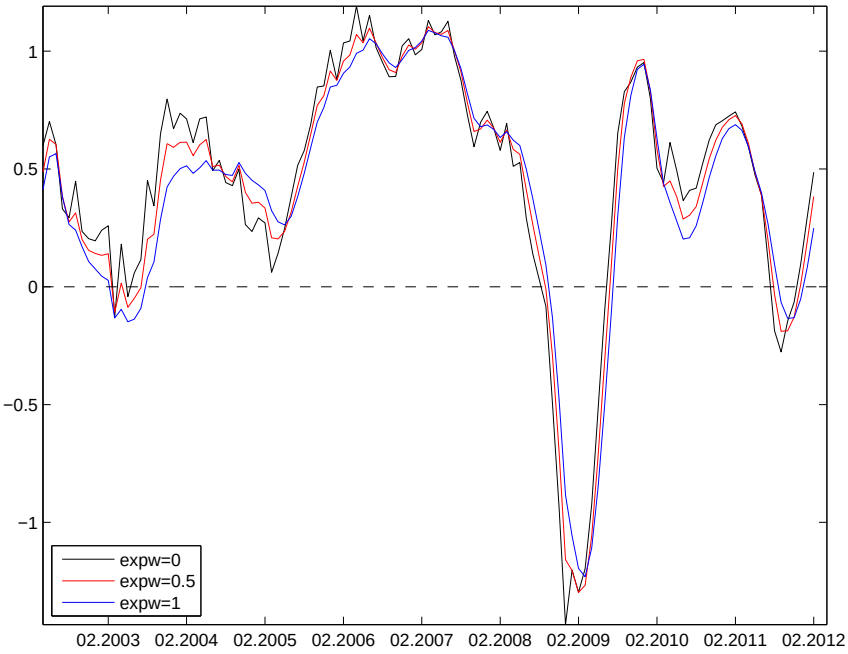

Fig. 2.7. Three outputs from the customized filter with different degrees of noise suppression in the rightmost stopband: $expw = 0$, $expw = 0.5$ and $expw = 1$

Note: The black line corresponds to the output of the classical mean squared error problem solution. The final indicator presented in this section is the middle case with $expw = 0.5$.

The black line corresponds to the classical mean squared optimization problem with $W(\omega_k) = 1$, i.e., there is no extra noise suppression. The other two lines lines correspond to the customized filter output with $expw = 0.5$ and $expw = 1$. It can be seen that the indicator without extra noise suppression is the fastest but also the noisiest of the three. A user of a real-time indicator might want to trade speed for extra smoothness such that the real-time signal would be more reliable. In that case, the filter customization kicks in by suppressing high-frequency content in the stopband. Ideally, that extra noise suppression would not alter the phase function in the passband but in practice, it does to some extent, i.e., the extra noise suppression in the stopband slows down the filter output somewhat. The presented final indicator is the middle line with $expw = 0.5$.

Comparison with Eurocoin and Markit EA PMI

Eurocoin. Eurocoin (Altissimo et al., 2010) is an established coincident indicator for EA GDP. According to its website[10], it is a smooth real time estimate of quarterly growth of EA GDP, released at the end of reference month, and is not revised since May 2009. It is a product of generalized dynamic factor model (Forni et al, 2000; Forni and Lippi, 2001) applied to - according to Altissimo et al., 2010 - about 145 series. In contrast, the new indicator is a result of a multivariate filter applied on only selected nine explanatory

[10]See http://eurocoin.bancaditalia.it/

variables. Eurocoin and the new indicator have a few features in common, as well - both target a trendcycle in the quarterly growth of Euro Area GDP as defined by a lowpass with cut-off frequency $\pi/6$, and both are designed to be coincident indicators. As such, Eurocoin is a decent benchmark for the new indicator.

Fig. 2.8 shows Eurocoin (black), latest vintage of quarterly growth of EA GDP (green), and the new indicator (red) for the last ten years.

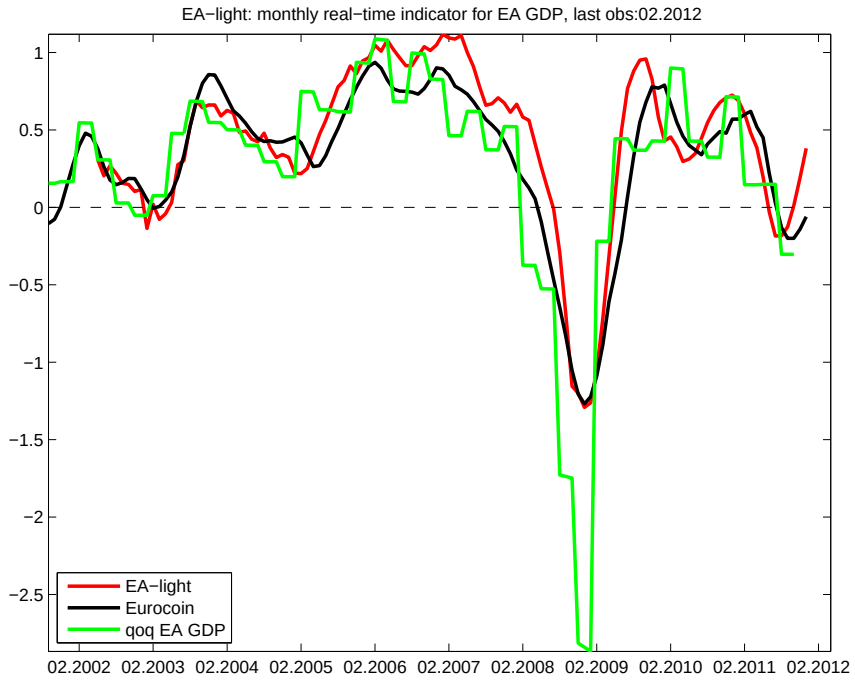

Fig. 2.8. Comparison with Eurocoin.

Fig. 2.8 shows that the amplitudes of both indicators are similar over the considered period. The new indicator is about coincident with Eurocoin during 2002-2004 but less smooth due to the small in-sample period. The new indicator appears to be faster than Eurocoin during 2005, and coincident but with worse level fit during 2007-2008. Eurocoin gets behind on average by about 3 months since 2009. Recall that Fig. 2.4 showed that the new indicator is slightly lagging the output of a finite symmetric filter for the whole episode of 2002-2008 and gets about coincident since 2009. Therefore, Fig. 2.8 indicates that Eurocoin is lagging with respect to the output of a finite symmetric filter for the whole comparison time span, 2002-2012. Given that the true out of sample period for Eurocoin is only since May 2009, and that its pseudo real-time values are calculated on the last vintage data for the period before May 2009[11], it might indicate that the real-time performance of Eurocoin is slightly worse than its pseudo real-time estimates suggest, and that the new indicator potentially slightly leads Eurocoin by a couple of months during true real time performance of both indicators.

[11]See the note in http://eurocoin.cepr.org/files/file/Ecoin_realtime_99Feb12.xls

Markit EA PMI. Markit euro area purchasing managers index is revealed a few days after the end of reference month and is advertised as being closely correlated with quarterly growth of EA GDP. Purchasing managers index is one of the closely watched indices by economic and financial agents due to its early release, simple design and economic relevance. The PMI data for the last five years are collected from Bloomberg and plotted against the new indicator. Since the PMI is of different magnitude than GDP growth, all variables are normalized to zero mean and unit variance for an easier comparison. Fig. 2.9 shows that the new indicator is about coincident with the PMI, while having a one month lead during the troughs of the last two recessions, and being smoother.

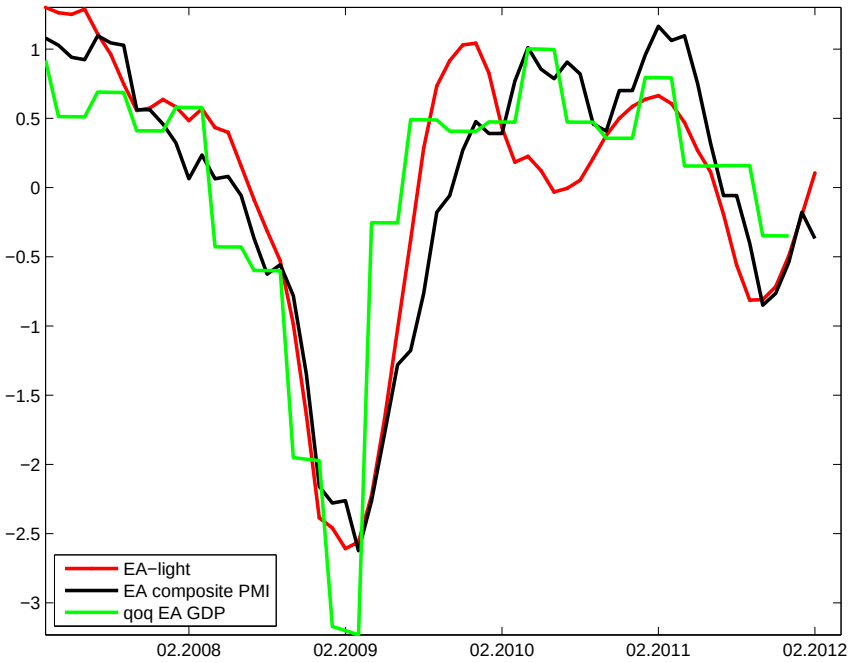

Fig. 2.9. Comparison with Markit EA PMI

Note: All series are normalized to zero mean and unit variance.

Additional robustness check: indicators for the four biggest European Union economies

Fig. 2.10 to 2.13 show prototypical indicators for German, French, the UK and Italian GDPs whose designs are a copy/paste design from the EA indicator with slight idiosyncratic modifications for each individual economy with respect to selected input data. The precise list of input data for each country is listed in 2.1.

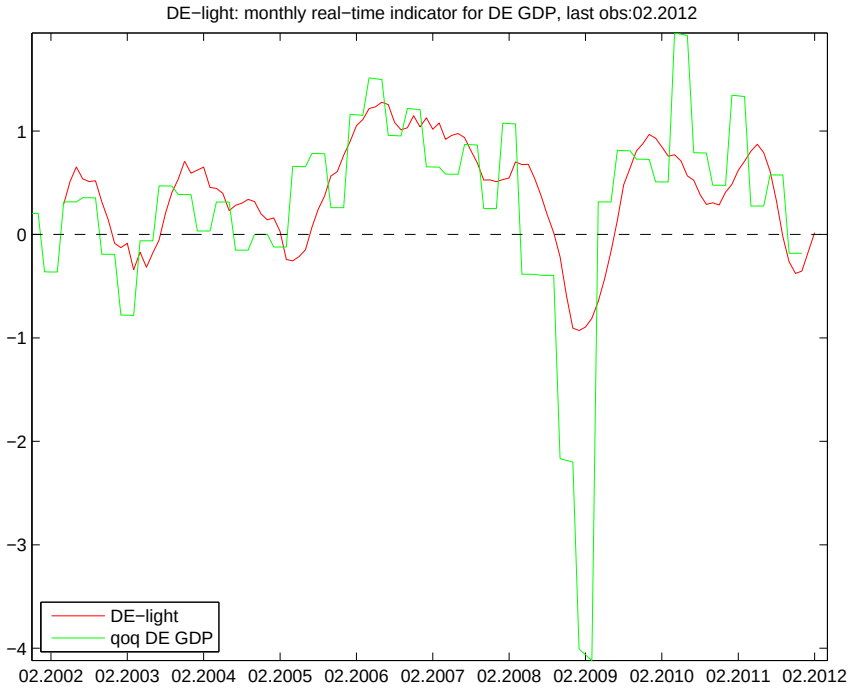

Fig. 2.10. A prototypical indicator for German GDP

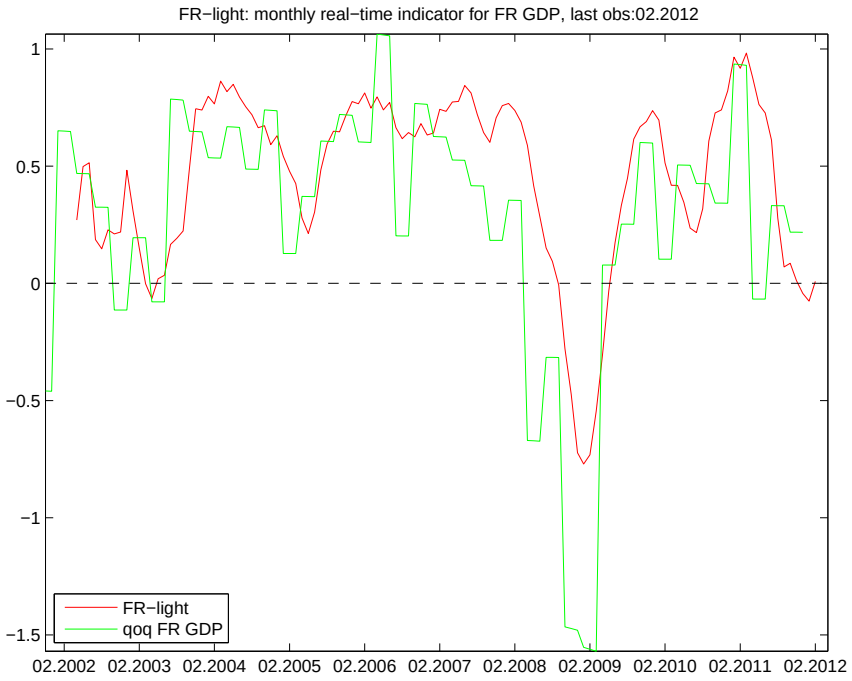

Fig. 2.11. A prototypical indicator for French GDP

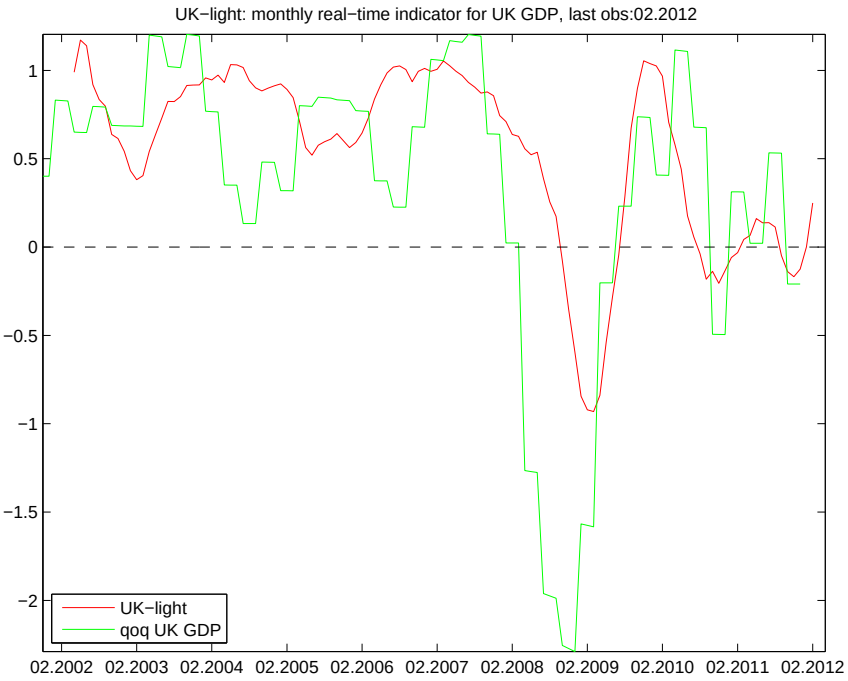

Fig. 2.12. A prototypical indicator for the UK GDP

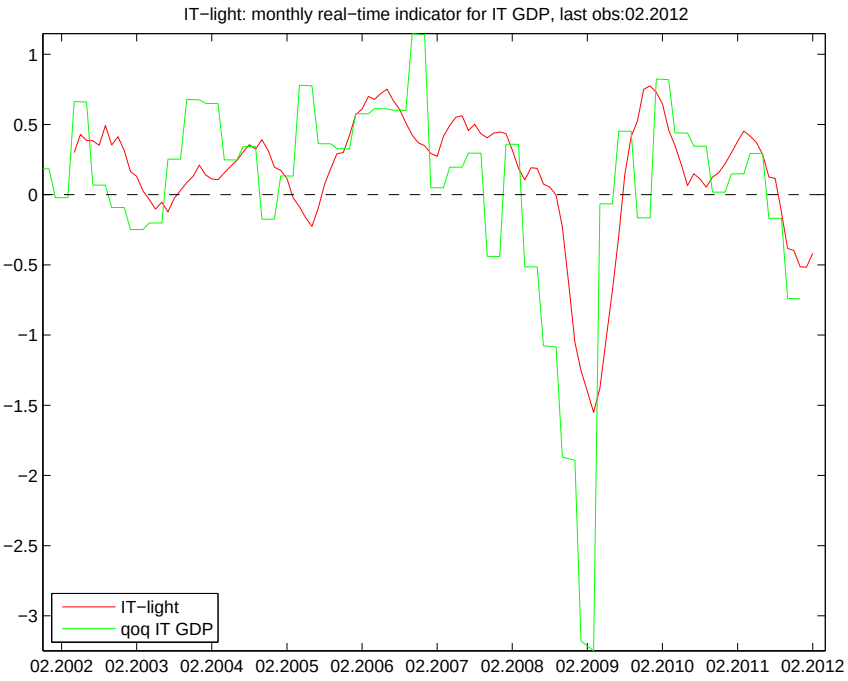

Fig. 2.13. A prototypical indicator for Italian GDP

Fig. 2.10 to 2.13 show that the filter design works quite well for German and French GDPs, slightly less so for Italian GDP, and its quality is even worse for the UK. These results show that the filter design has some merit but also that one size doesn't fit all - every region has its own specificities to be accounted for by differentiating the input data.

2.1.4 Conclusions

1. This section builds a new real-time indicator tracking medium-to-long-rum component of quarterly growth of euro area gross domestic product using new filtration methodology - the multivariate direct filter approach.

2. The new indicator behaves similarly to another established indicator - Eurocoin - but leads it by about 3 months after mid 2009.

3. The multivariate direct filter approach is found to be a valuable tool for multivariate filtering but it cannot be applied on high-dimensional datasets due to overparameterization.

2.2 Forecasting and signal extraction with regularized multivariate direct filter approach

This section studies 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 section illustrates the features of the filter by tracking the medium-to-long-run component in GDP growth for euro area, including the replication of an established indicator's 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 a 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, and a decent alternative to dynamic factor methodology.

2.2.1 Introduction

Nowadays, gathering rich datasets is relatively easy. A more difficult exercise is effectively using them for a particular problem at hand. This section adds to the forecasting/regularization/ shrinkage/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), principal components (Stock and Watson, 2002)) by exploring the properties an abilities of a regularized direct filter approach (Wildi, 2012) in signal extraction and forecasting using 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. Thus, this section is the first paper known to the author that studies and implements the RMDFA on data. Furthermore, this section studies how RMDFA could be implemented on high-dimensional real-world data sets. This section finds that filter regularization helps in real-time filter extraction since it controls 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 show in this section, a particular regularization feature used in the section might remind about the 'lag decay' term in Minnessota prior (see, e.g. Doan, Litterman and Sims, 1984) in Bayesian econometrics. Forcing more distant filter coefficients to zero both saves degrees of freedom and effectively shortens the filter, thus making it more responsive to changing environment. Another regularization feature studied in the

section is cross-sectional shrinkage that makes filter coefficients to behave similarly for similar 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 have 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 section 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, and is proved to stand the test.

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 section is structured as follows. Subsection 2.2.2 introduces new regularization features in the direct filter approach. Subsection 2.2.3 illustrates the new features of the filter by creating indicators for yearly and quarterly growth rates of euro area GDP. Each subsection starts with indicator creation for the euro area GDP. Subsection 2.2.4 performs

a robustness check on a less homogeneous Latvia's dataset. Subsection 2.2.8 concludes. Section 2.2.7 lists the data and their transformations.

2.2.2 Regularized multivariate direct filter approach

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 (see, e.g. 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 accordingly 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 too short filter would result in deteriorating quality of its output. Therefore, similar to standard econometric practices in parameter shrinkage (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)), 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 appraoch (Wildi, 2011) and that control for cross-sectional shrinkage, shrinkage along time dimension, and that impose 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 section builds on the methodology of Section 2.1.2, and rewrites the filtration problem in a least-squares form.

Rewriting filtration problem in a least squares form

In order to conveniently define the regularized filter problem, the multivariate filtration problem in Section 2.1.2 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:

$$X'_k = (1 + I_{k>0}) \operatorname{Vec} \begin{pmatrix} \Xi_{Tx}(\omega_k) & \exp(-i\omega_k)\Xi_{Tx}(\omega_k) & \cdots & \exp(-iL\omega_k)\Xi_{Tx}(\omega_k) \\ \Xi_{Tz_1}(\omega_k) & \exp(-i\omega_k)\Xi_{Tz_1}(\omega_k) & \cdots & \exp(-iL\omega_k)\Xi_{Tz_1}(\omega_k) \\ \Xi_{Tz_2}(\omega_k) & \exp(-i\omega_k)\Xi_{Tz_2}(\omega_k) & \cdots & \exp(-iL\omega_k)\Xi_{Tz_2}(\omega_k) \\ \vdots & \vdots & \vdots & \vdots \\ \Xi_{Tz_m}(\omega_k) & \exp(-i\omega_k)\Xi_{Tz_m}(\omega_k) & \cdots & \exp(-iL\omega_k)\Xi_{Tz_m}(\omega_k) \end{pmatrix}, \tag{2.17}$$

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

$$b = \operatorname{Vec} \begin{pmatrix} b_{x0} & b_{z_1 0} & b_{z_2 0} & \cdots & b_{z_m 0} \\ b_{x1} & b_{z_1 1} & b_{z_2 1} & \cdots & b_{z_m 1} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ b_{xL} & b_{z_1 L} & b_{z_2 L} & \cdots & b_{z_m L} \end{pmatrix}, \qquad Y = \begin{pmatrix} \Gamma(\omega_0)\Xi_{Tx}(\omega_0) \\ 2\Gamma(\omega_1)\Xi_{Tx}(\omega_1) \\ 2\Gamma(\omega_2)\Xi_{Tx}(\omega_2) \\ \vdots \\ 2\Gamma(\omega_{[T/2]})\Xi_{Tx}(\omega_{[T/2]}) \end{pmatrix}. \tag{2.18}$$

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

$$(Y - Xb)'(Y - Xb) \to \min_b \tag{2.19}$$

Since X and Y are complex-valued, the solution to (2.19) 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 (2.19) is unaffected:

$$\begin{aligned} X_{k,rot} &= X_k \exp\left(-i \arg(\Gamma(\omega_k)\Xi_{Tx}(\omega_k)) + ih\omega_k\right) \\ Y_{rot} &= |Y|, \end{aligned} \tag{2.20}$$

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

$$(Y_{rot} - X_{rot}b)'(Y_{rot} - X_{rot}b) \to \min_b . \tag{2.21}$$

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

$$X_{k,rot}^{cust} = X_{k,rot}\sqrt{W(\omega_k, expw, cut)} \tag{2.22}$$

$$Y_{rot}^{cust} = \begin{pmatrix} |\Gamma(\omega_0)\Xi_{Tx}(\omega_0)|\sqrt{W(\omega_0, expw, cut)} \\ 2|\Gamma(\omega_1)\Xi_{Tx}(\omega_1)|\sqrt{W(\omega_1, expw, cut)} \\ \vdots \\ 2|\Gamma(\omega_{[T/2]})\Xi_{Tx}(\omega_{[T/2]})|\sqrt{W(\omega_{[T/2]}, expw, cut)} \end{pmatrix}, \tag{2.23}$$

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

$$(Y_{rot}^{cust} - X_{rot}^{cust}b)'(Y_{rot}^{cust} - X_{rot}^{cust}b) \to \min_b, \tag{2.24}$$

which collapses to (2.21) for $expw = 0$.

We are now ready to introduce the regularized filter problem.

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:

$$(Y_{rot}^{cust} - X_{rot}^{cust}b)'(Y_{rot}^{cust} - X_{rot}^{cust}b) + \lambda_s b'Q_s b + \lambda_c b'Q_c b + \lambda_d b'Q_d b \to \min_b, \tag{2.25}$$

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

$$b'Q_s b = \sum_{u=0}^{m}\sum_{l=2}^{L}\left((1-L)^2 b_l^u\right)^2, \tag{2.26}$$

where $(1-L)^2 b_l^u = b_l^u - 2b_{l-1}^u + b_{l-2}^u$ is the second order difference of b_L^u, $l = 0, \ldots, L$, and $u = 0, \ldots, m$. Therefore, the term in (2.26) 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

$$\sum_{u=0}^{m}\left(\left(b_0^u-\frac{1}{m+1}\sum_{u'=0}^{m}b_0^{u'}\right)^2+\left(b_1^u-\frac{1}{m+1}\sum_{u'=0}^{m}b_1^{u'}\right)^2+\ldots+\left(b_L^u-\frac{1}{m+1}\sum_{u'=0}^{m}b_L^{u'}\right)^2\right) \tag{2.27}$$

which yields a symmetric bilinear form with

$$Q_c=\begin{pmatrix} q_{c,1} \\ q_{c,2} \\ \vdots \\ q_{c,(m+1)*(L+1)} \end{pmatrix} \tag{2.28}$$

where

$$\begin{aligned}
q_{c,1} &= (1-\frac{1}{m+1},0,\ldots,0|-\frac{1}{m+1},0,\ldots,0|-\frac{1}{m+1},0,\ldots,0|\ldots) \\
q_{c,2} &= (0,1-\frac{1}{m+1},0,\ldots,0|0,-\frac{1}{m+1},0,\ldots,0|0-\frac{1}{m+1},0,\ldots,0|\ldots) \\
q_{c,3} &= (0,0,1-\frac{1}{m+1},0,\ldots,0|0,0,-\frac{1}{m+1},0,\ldots,0|0,0,-\frac{1}{m+1},0,\ldots,0|\ldots) \\
&\ldots \\
q_{c,(m+1)*(L+1)} &= (0,0,\ldots,-\frac{1}{m+1}|0,0,\ldots,-\frac{1}{m+1}|0,0,\ldots,-\frac{1}{m+1}|\ldots|0,0,\ldots,1-\frac{1}{m+1})
\end{aligned} \tag{2.29}$$

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 (2.27).

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. For a Bayesian econometrician this would remind of the lag decay in the Minnesota prior (e.g. Doan, Litterman and Sims, 1984). This shrinkage is implemented by setting Q_d such that

$$b'Q_d b=\sum_{u=0}^{m}\sum_{l=0}^{L}\tilde{q}_l(b_l^u)^2, \tag{2.30}$$

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

$$\tilde{q}=\left(q^{0\vee h},q^{|1-0\vee h|},q^{|2-0\vee h|},\ldots,q^{|L-0\vee h|}\right), \tag{2.31}$$

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, (2.31) 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 (2.25) has an analytic solution. Setting $\lambda_s = \lambda_c = \lambda_d = 0$ gives the unregularized filter problem in (2.24). Or, setting $expw = 0$ but letting some of the regularization lambdas positive gives the regularized classical multivariate filter problem. This section 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 (2.25) subject to two potential constraints - first- and/or second-order constraints which are explained in the following subsection.

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:

$$b^u_{-h} + b^u_{-(h-1)} + \ldots + b^u_{L-h} = w^u, \tag{2.32}$$

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 ($h = 0$ for a concurrent filter, $h > 0$ for a smoother, and $h < 0$ for forecasting the signal).

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:

$$-hb^u_{-h} + (1-h)b^u_{1-h} + (2-h)b^u_{2-h} + \ldots + b^u_1 + 2b^u_2 + \ldots + (L-h)b^u_{L-h} = 0, \quad (2.33)$$

where h is the targeted lag ($h = 0$ for a concurrent filter, $h > 0$ for a smoother, and $h < 0$ for forecasting the signal).

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, e.g. a lag decay agenda for h large enough.

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

$$b = Rb_f + c, \quad (2.34)$$

where b_f is the vector of freely determined filter coefficients, plugging (2.34) in (2.25), solving for b_f, and then plugging the estimate of b_f into (2.34) to get the estimate of b; see Wildi (2012) for details.

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:

$$\hat{Y} = SY = X(X'X)^{-1}X'Y = PY. \quad (2.35)$$

The degrees of freedom is trace of the projection matrix:

$$d.f. = \mathrm{tr}(P), \quad (2.36)$$

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

For a regularized problem as in expression (2.25),

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

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 $\tilde{S}$:

$$\tilde{S} = \mathrm{Re}(X_{rot}^{cust}) \left((X_{rot}^{cust})'X_{rot}^{cust} + \lambda_s Q_s + \lambda_c Q_c + \lambda_d Q_d\right)^{-1} \mathrm{Re}(X_{rot}^{cust})', \tag{2.37}$$

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 $\tilde{S}$:

$$e.d.f. = \mathrm{tr}(\tilde{S}), \tag{2.38}$$

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

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

2.2.3 Tracking economic activity in the euro area

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.

Target. The filter is set to target an ideal lowpass of yearly growth of real GDP with cut-off wave length 12 months. The quarterly GDP data are taken from 1995Q1 onwards, as published by Eurostat. The data are linearly interpolated to monthly frequency, logged, yearly differenced and demeaned before their spectral content enters the filter.

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. Table 2.3 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[12]. 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 onwards, standardized to zero mean and unit variance. Integrated data are made non-integrated by suitable transformations. Table 2.3 lists the data and their transformations.

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) 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. Fig. 2.14(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 Fig. 2.14(b).

[12]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

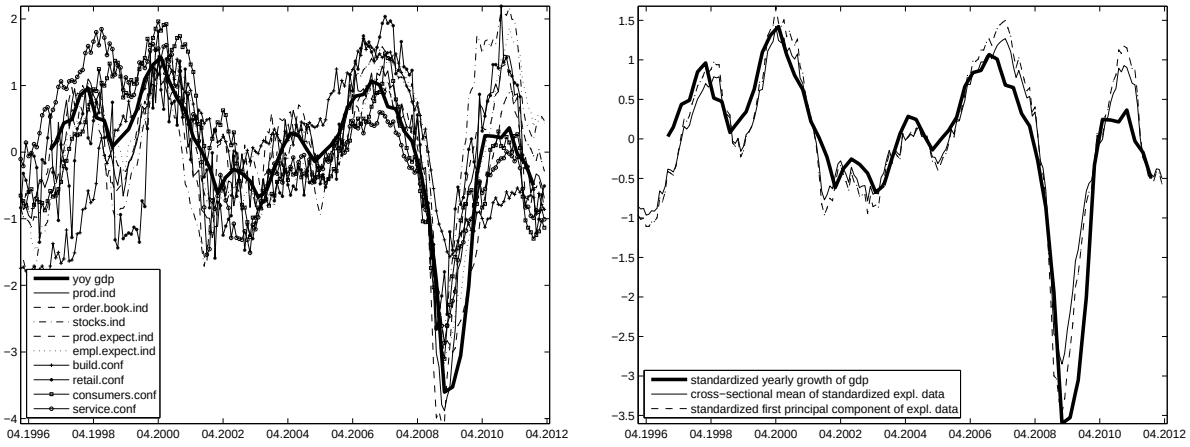

(a) (b)

Fig. 2.14. Yearly growth of GDP versus business and consumers data

Note: (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 section 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 estimation 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 having been shrunk to zero 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. Fig. 2.15(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. Fig. 2.15(b) shows the (rather chaotic) filter amplitudes corresponding to the coefficients in Fig. 2.15(a); it will be useful to analyze how the amplitudes change with various constraints and regularization restrictions.

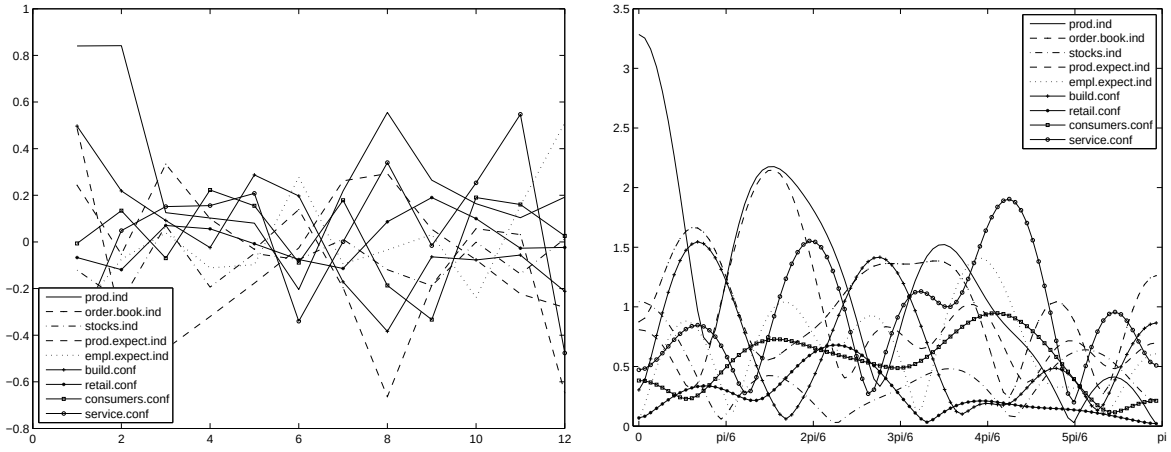

(a) (b)

Fig. 2.15. 9-variable filter without regularization and filter constraints

Note: (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 Fig. 2.15(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 saves one d.f. per input series, thus nine d.f. are saved for an unregularized nine-variable filter.

Fig. 2.16(a) and 2.16(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).

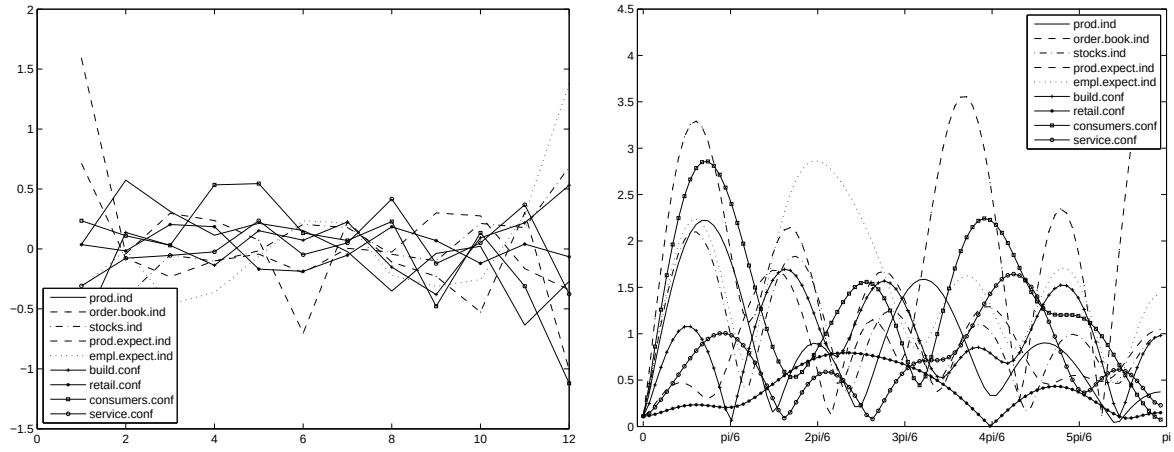

(a) (b)

Fig. 2.16. 1st-order constrained lowpass filter

Note: (a) Coefficients for a 1st-order constrained lowpass filter. (b) Filter amplitudes corresponding to the coefficients in Fig. 2.16(a).

The second order restriction imposes a vanishing phase shift at frequency zero for a targeted lead or lag, and also saves 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 Fig. 2.17(a) and 2.17(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 Fig. 2.14(a) and 2.14(b).) Still, the second order constrain is not a panacea since the amplitudes are still erratic and since the number of degrees of freedom

vanishes for samples smaller than 9*(12-1)=99 months which is 8 years of data.

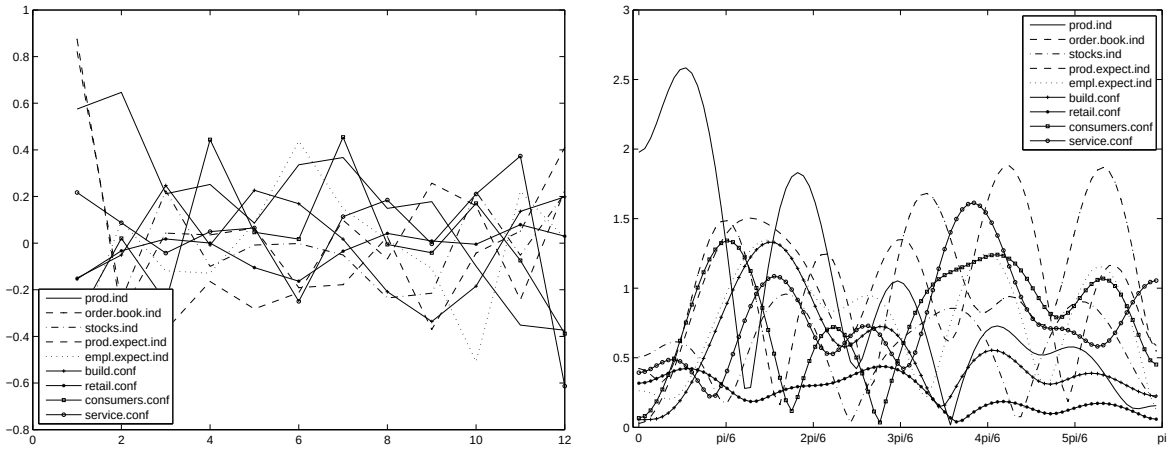

(a) (b)

Fig. 2.17. 2nd-order constrained concurrent filter

Note: (a) Coefficients for a 2nd-order constrained concurrent filter. (b) Filter amplitudes corresponding to the coefficients in Fig. 2.17(b).

Turning to the new regularization features, Fig. 2.18(a) to 2.20(b) 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.

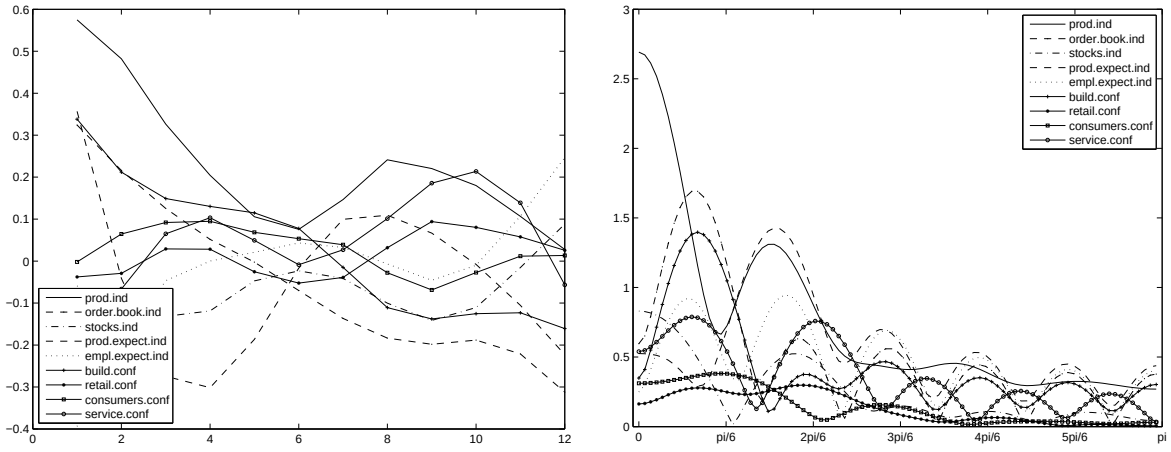

(a) (b)

Fig. 2.18. Unconstrained filter if $\lambda_s = 0.01$

Note: (a) Coefficients for an unconstrained filter if $\lambda_s = 0.01$. (b) Filter amplitudes corresponding to the coefficients in Fig. 2.18(a).

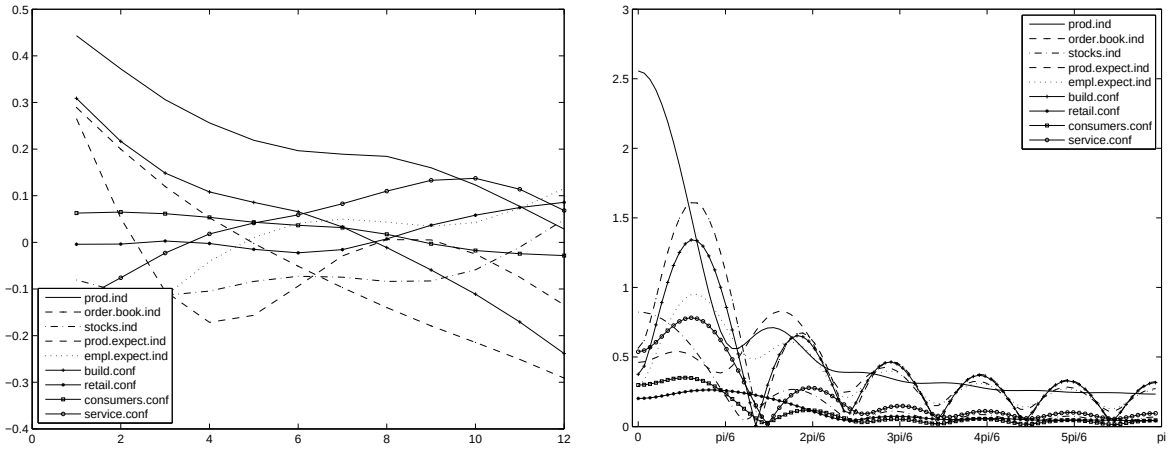

(a) (b)

Fig. 2.19. Unconstrained filter if $\lambda_s = 0.1$

Note: (a) Coefficients for an unconstrained filter if $\lambda_s = 0.1$. (b) Filter amplitudes corresponding to the coefficients in Fig. 2.19(a).

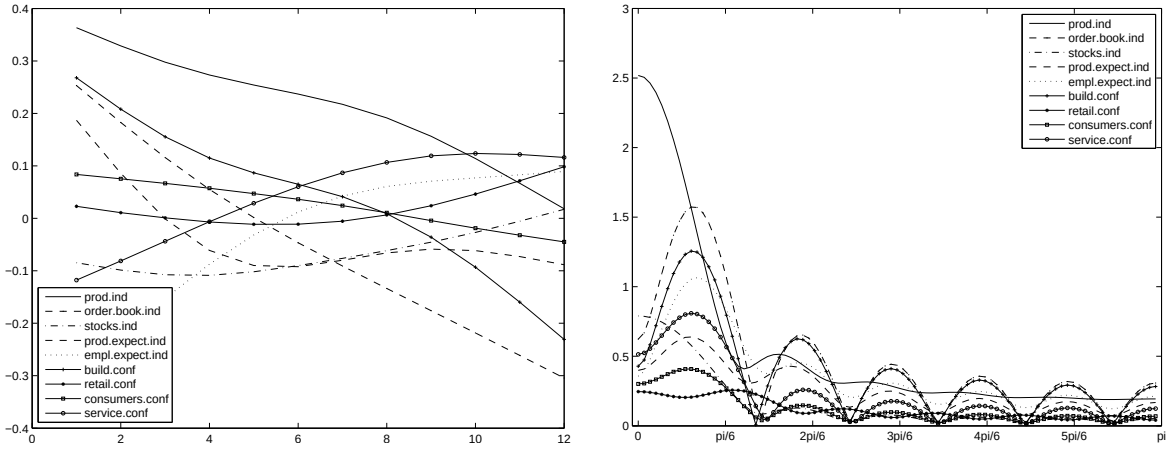

(a) (b)

Fig. 2.20. Unconstrained filter if $\lambda_s = 1$

Note: (a) Coefficients for an unconstrained filter if $\lambda_s = 1$. (b) Filter amplitudes corresponding to the coefficients in Fig. 2.20(a).

Fig. 2.18(a) to 2.20(b) 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.

Fig. 2.21(a) to 2.23(b) 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.

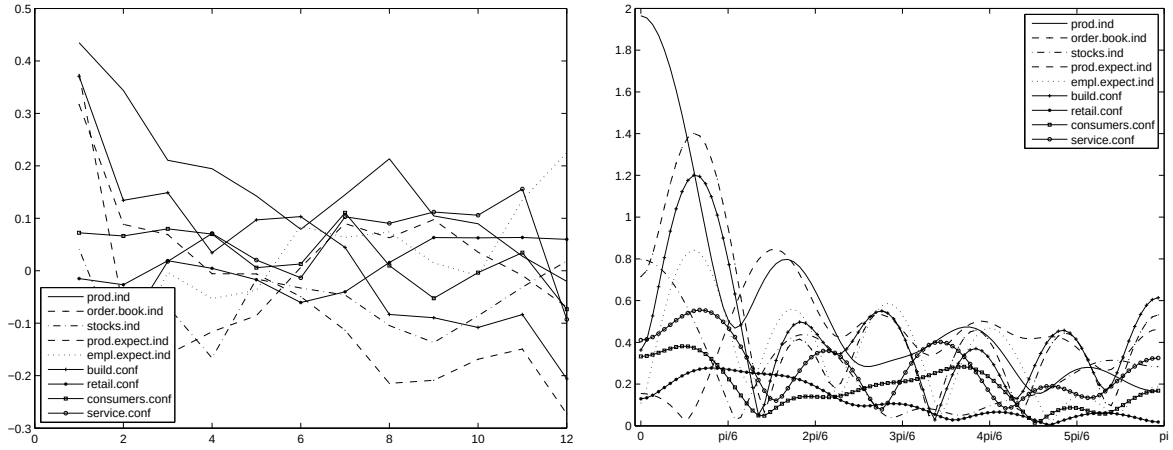

(a) (b)

Fig. 2.21. Unconstrained filter if $\lambda_c = 0.01$

Note: (a) Coefficients for an unconstrained filter if $\lambda_c = 0.01$. (b) Filter amplitudes corresponding to the coefficients in Fig. 2.21(a).

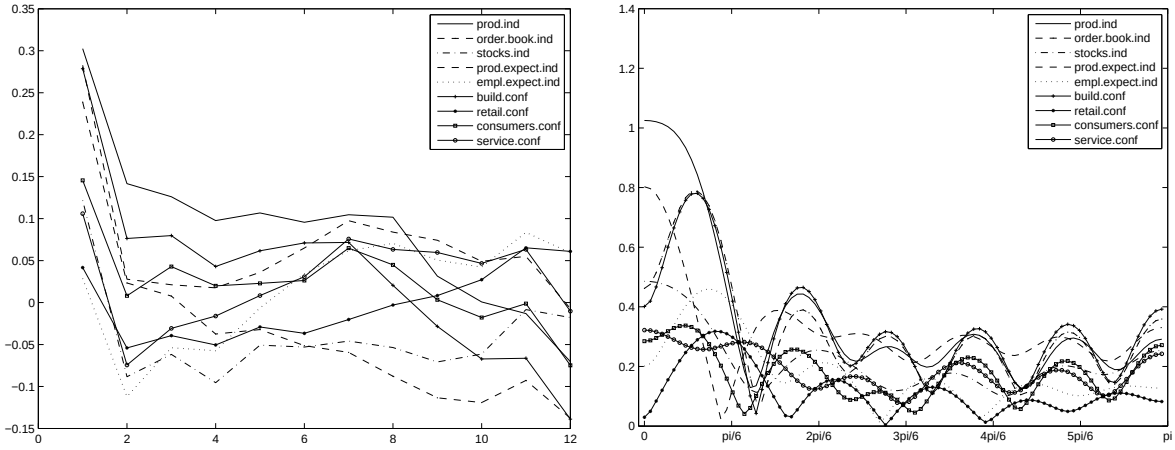

(a) (b)

Fig. 2.22. Unconstrained filter if $\lambda_c = 0.1$

Note: (a) Coefficients for an unconstrained filter if $\lambda_c = 0.1$. (b) Filter amplitudes corresponding to the coefficients in Fig. 2.22(a).

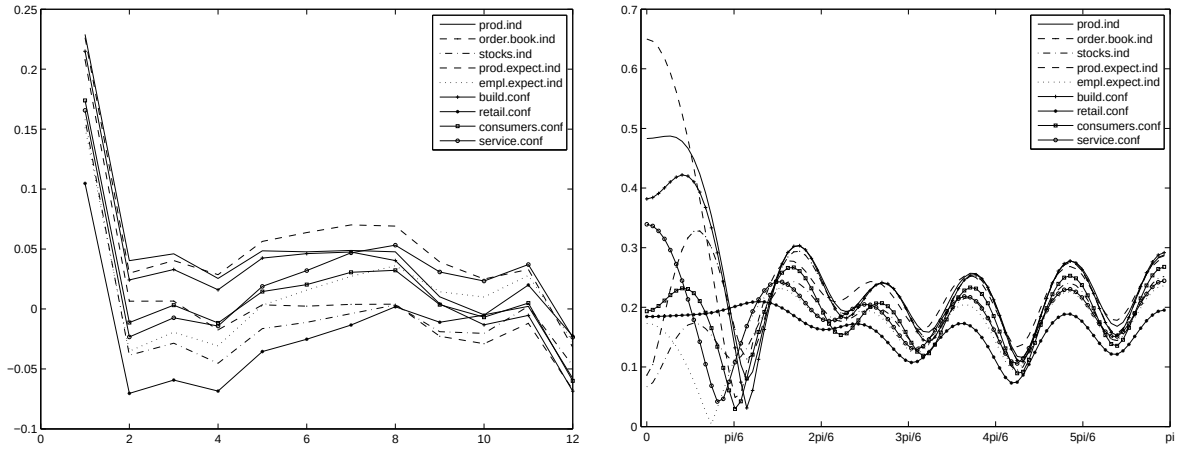

(a) (b)

Fig. 2.23. Unconstrained filter if $\lambda_c = 1$

Note: (a) Coefficients for an unconstrained filter if $\lambda_c = 1$. (b) Filter amplitudes corresponding to the coefficients in Fig. 2.23(a).

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 Fig. 2.21(a) and 2.21(b)) but further cross-sectional restriction can be harmful if applied alone (see amplitude behaviour in Fig. 2.23(b)). Such a cross-sectional restriction analysis might help understand which series or clasters of series are different from the others. In our exercise, no series clearly stands out from the rest.

As for the third regularization feature, Fig. 2.24(a) to 2.26(b) 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.

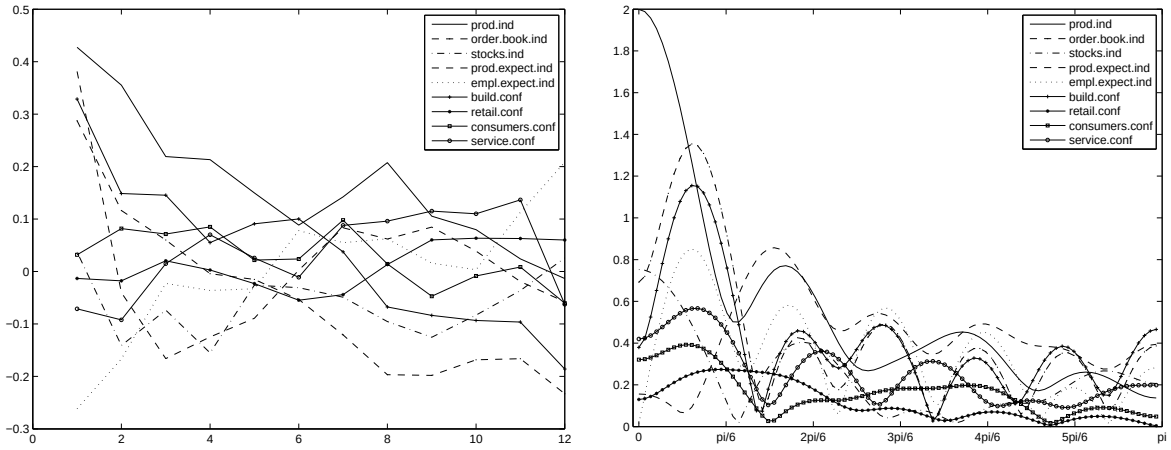

(a) (b)

Fig. 2.24. Unconstrained filter if $\lambda_d = 0.01$

Note: (a) Coefficients for an unconstrained filter if $\lambda_d = 0.01$. (b) Filter amplitudes corresponding to the coefficients in Fig. 2.24(a).

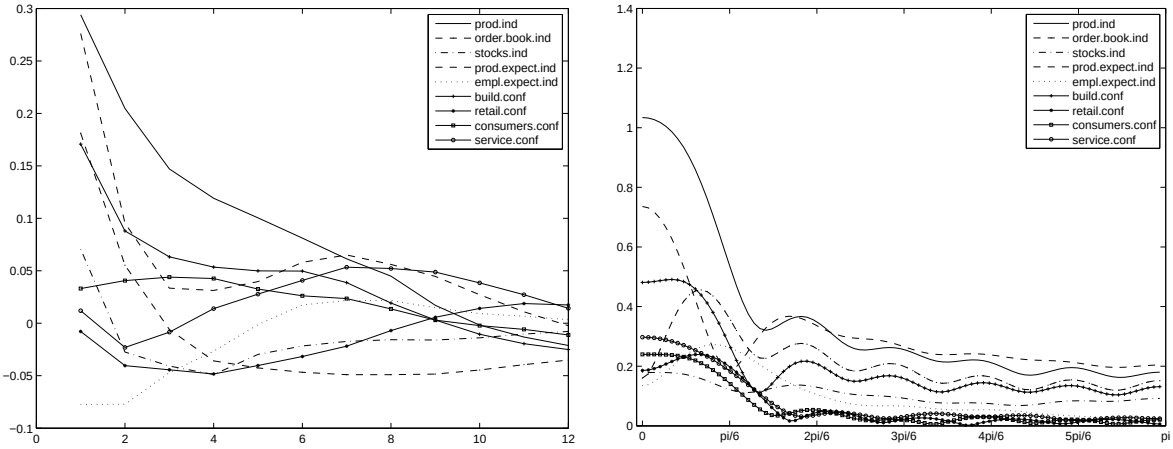

(a) (b)

Fig. 2.25. Unconstrained filter if $\lambda_d = 0.1$

Note: (a) Coefficients for an unconstrained filter if $\lambda_d = 0.1$. (b) Filter amplitudes corresponding to the coefficients in Fig. 2.25(a).

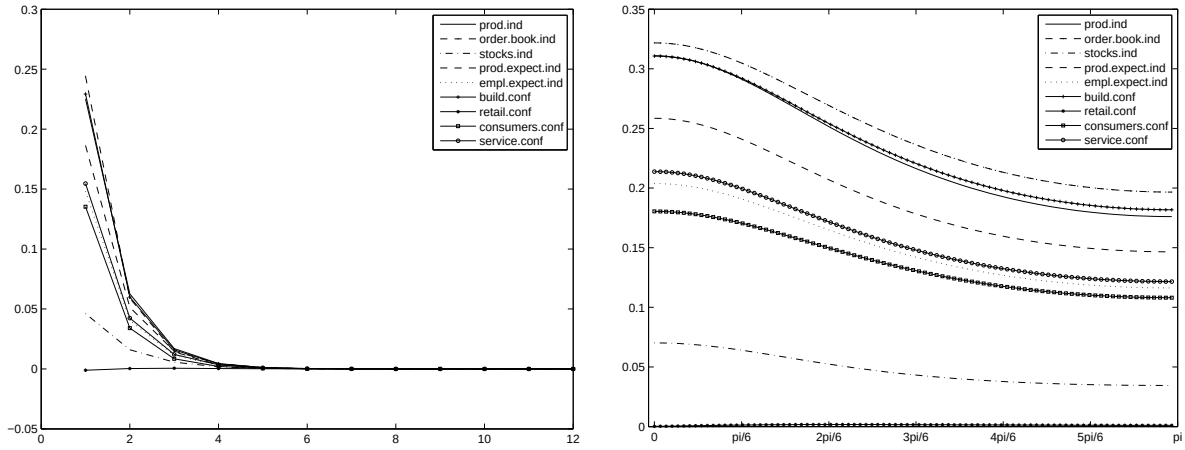

(a) (b)

Fig. 2.26. Unconstrained filter if $\lambda_d = 1$

Note: (a) Coefficients for an unconstrained filter if $\lambda_d = 1$. (b) Filter amplitudes corresponding to the coefficients in Fig. 2.26(a).

Fig. 2.24(a) to 2.26(b) 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. Fig. 2.26(b) shows that a sufficiently high longitudinal shrinkage forces filter amplitudes to shrink towards zero (see the scale of Fig. 2.26(b)) and flatten, resembling those of an allpass filter, which is an expected behaviour since a short filter cannot discriminate between frequencies effectively.

Coefficients in Fig. 2.24(a) and 2.25(a) are rather smooth which resembles the effect of parameter smoothness restriction. Also, Fig. 2.24(a) and 2.25(a) 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 Fig. 2.26(a). But, instead of using both longitudinal and parameter smoothness regularization features, one might just loosen the lag decay restriction.

The findings in this section indeed suggest that the longitudinal shrinkage might be the most useful of the three regularization features. Moreover, this section 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.

Fig. 2.27 shows a flowchart for using the regularized filter with many variables. First, the user selects the target variable and its any additional explanatory variables (if they exist), Next, one should decide on the target amplitude and the targeted lag or lead of the signal. Also, the user decides on the training sample length. The explanatory and the target variable should be cointegrated, otherwise all data should be stationarized. Then

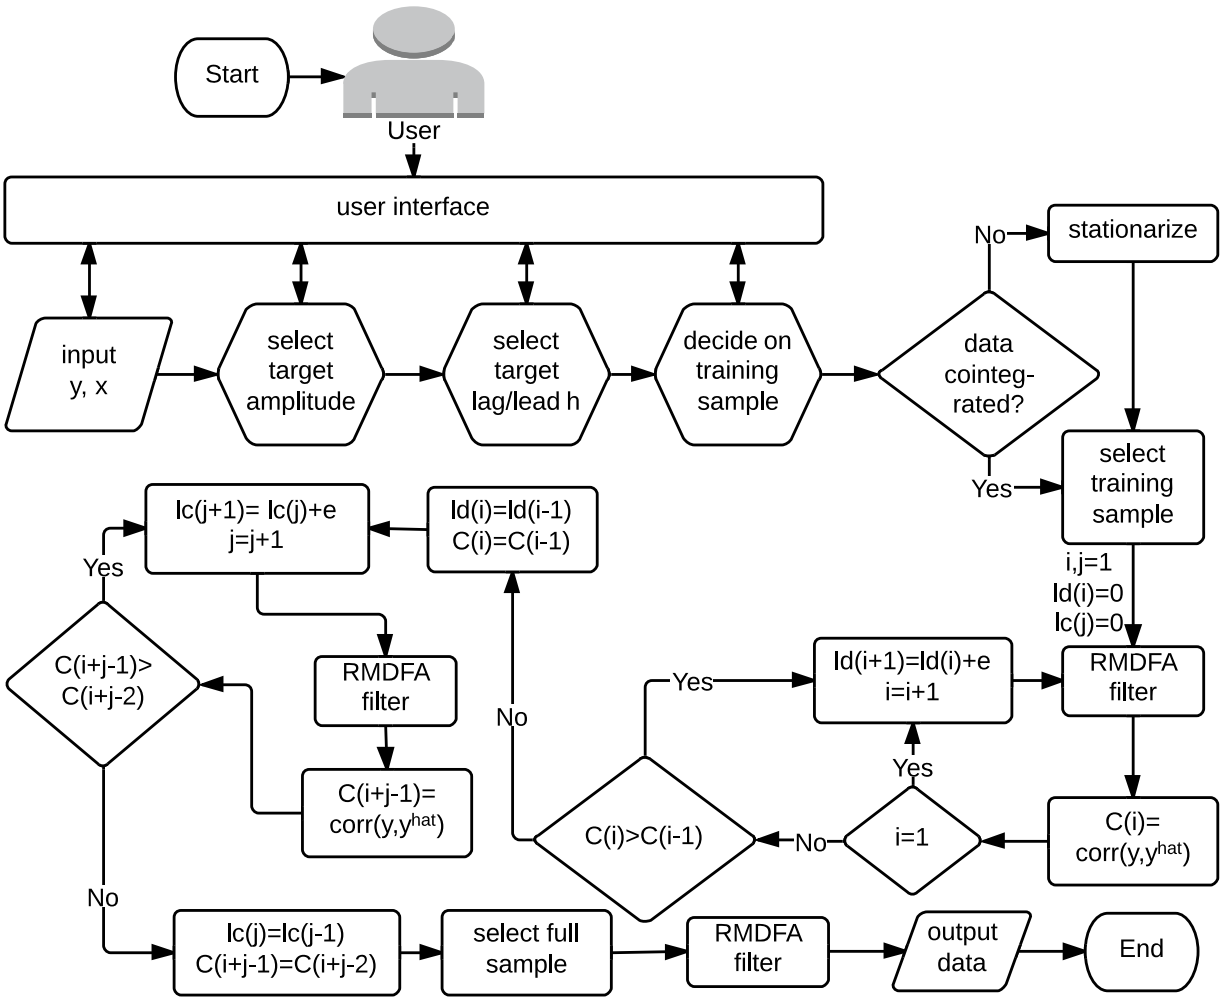

Fig. 2.27. A flowchart of the high-dimensional regularized filter algorithm

follows the most important part of deciding on the metric to be used as a measure of goodness-of-fit. Since the output of the ideal filter is unobserved, one cannot use a metric that uses the output of the ideal filter (i.e., a metric such as the mean squared error or the mean absolute error). Rather, one can use the correlation between the target variable and the output of the one-sided filter at the pre-specified targeted lead or lag, at least when the target amplitude is trendcycle. Thus, the algorithm follows by running the filter and incrementally increasing the longitudinal shrinkage parameter ('ld' in the diagram) until the maximum correlation between the target variable and the filter output at the prespecified lead or lag has been reached. Then, the algorithm selects the longitudinal shrinkage parameter that maximizes the correlation and then moves on to incrementally increasing the cross-sectional shrinkage parameter ('lc' in the diagram) to find the one that maximizes the chosen distance measure at the prespecified target lead or lag. Once the optimal shrinkage parameters have been found, the algorithm moves on by selecting the full sample and gives the output.

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.

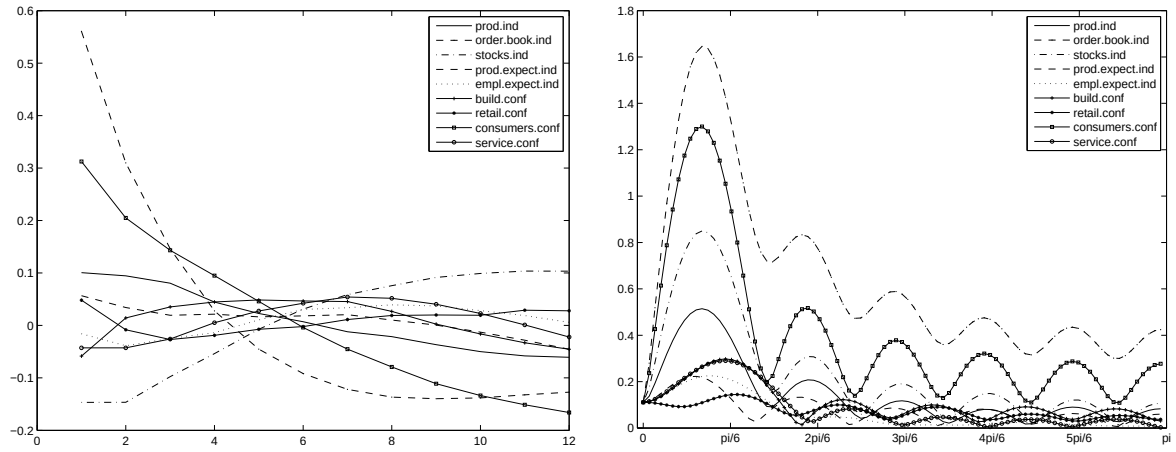

(a) (b)

Fig. 2.28. Filter if longitudinal regularization with $\lambda_d = 0.1$ and the 1st order constraint are implemented

Note: (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 Fig. 2.28(a).

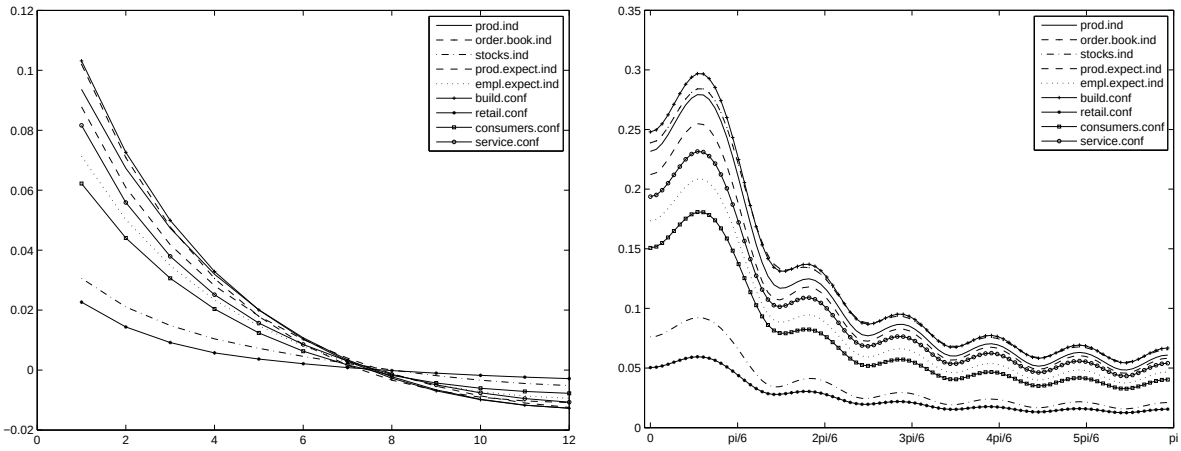

(a) (b)

Fig. 2.29. Filter if longitudinal regularization with $\lambda_d = 0.1$ and the 2nd order constraint are implemented

Note: (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 Fig. 2.29(a).

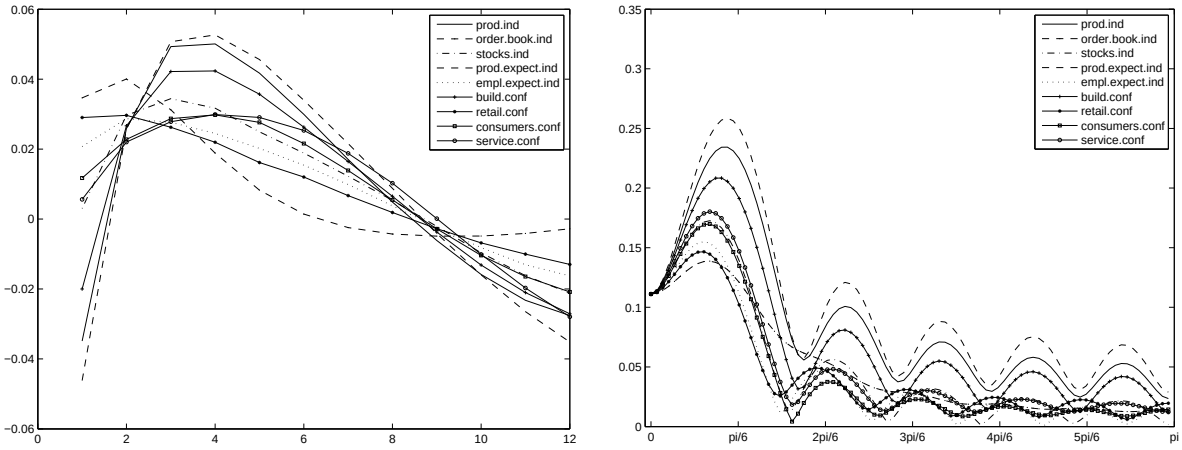

(a) (b)

Fig. 2.30. Filter if longitudinal regularization with $\lambda_d = 0.1$ and both the 1st- and the 2nd-order constraints are implemented

Note: (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 Fig. 2.30(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. Fig. 2.30(a) and 2.30(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.

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 Fig. 2.31(a) and 2.31(b), respectively. Figure labels are removed due to over-cluttering.

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 Fig. 2.32(a) along with real-time values of the appropriately scaled cross-sectional mean of the explanatory variables and the yearly growth of EA GDP. The particular parameter setting results in about 3 e.d.f. on average over the whole sample.

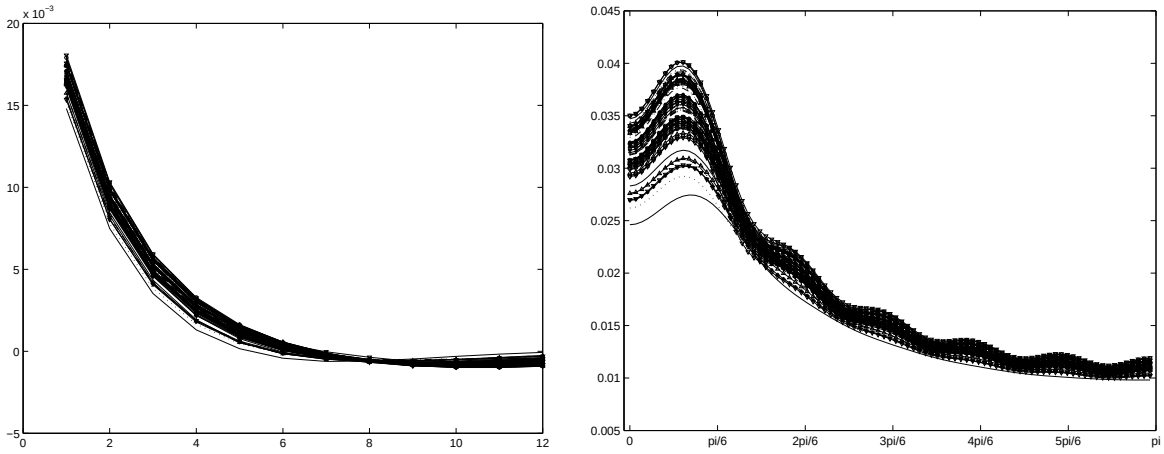

(a) (b)

Fig. 2.31. A concurrent 72-variables filter with 2nd-order constraint, $h = 0$, $\lambda_d = 0.2$, $\lambda_c = 5$

Note: (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 Fig. 2.31(a).

The filter output is smoother than the real-time cross-sectional mean of explanatory data but equally fast on turning point detection. It is coincident with yearly growth of EA GDP by construction due to the time-shift constraint.

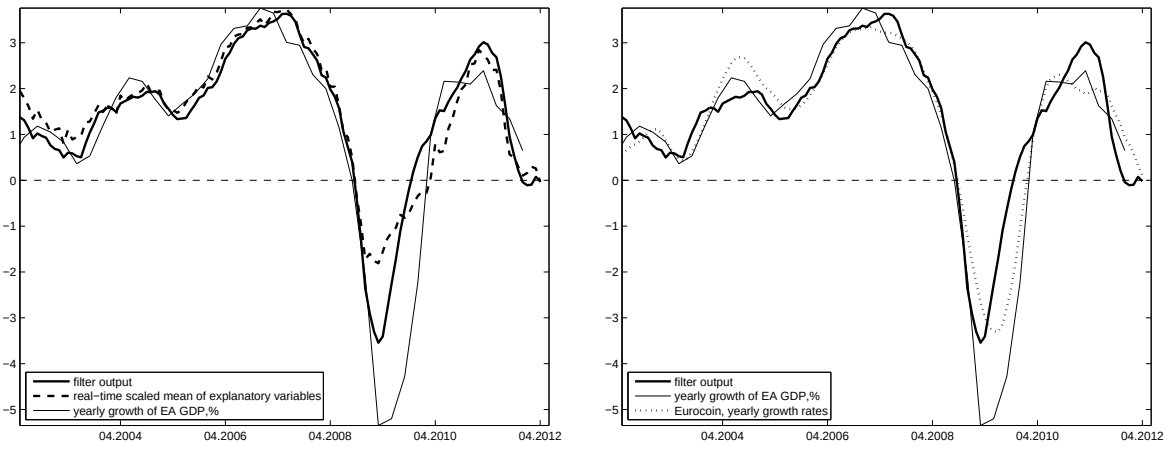

(a) (b)

Fig. 2.32. Output of a regularized 72-variable filter with $h = 0$, $\lambda_d = 0.2$, $\lambda_c = 5$ tracking a trendcycle in yoy EA GDP

Note: (a) 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 real-time values of scaled cross-sectional mean of explanatory data. (b) The same filter output as in (a) but plotted against Eurocoin transformed to represent yearly growth rates. The filter output is clearly faster than Eurocoin on several occasions.

Figure 2.32(a) shows that the filter output is not trivial, i.e., it not only does not lose w.r.t. timeliness and win w.r.t. smoothness, but also it clearly outperforms the real-

time cross-sectional mean of explanatory variables w.r.t. level fit. Fig. 2.32(b) compares the filter output with another established indicator Eurocoin (see Altissimo et al. 2010), the latter being transformed to yearly growth rates. Fig. 2.32(b) 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 Fig. 2.32(b), 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 2.2).

Forecasting. This paragraph shows that the regularized filter can be used not only for a concurrent signal extraction but also for forecasting. Fig. 2.33(a) and 2.33(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 (2.20), (2.31), (2.32), and (2.33). 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. Fig. 2.33(a) and 2.33(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 Fig. 2.34(a) and 2.34(b), which show even more dispersed coefficients and amplitudes.

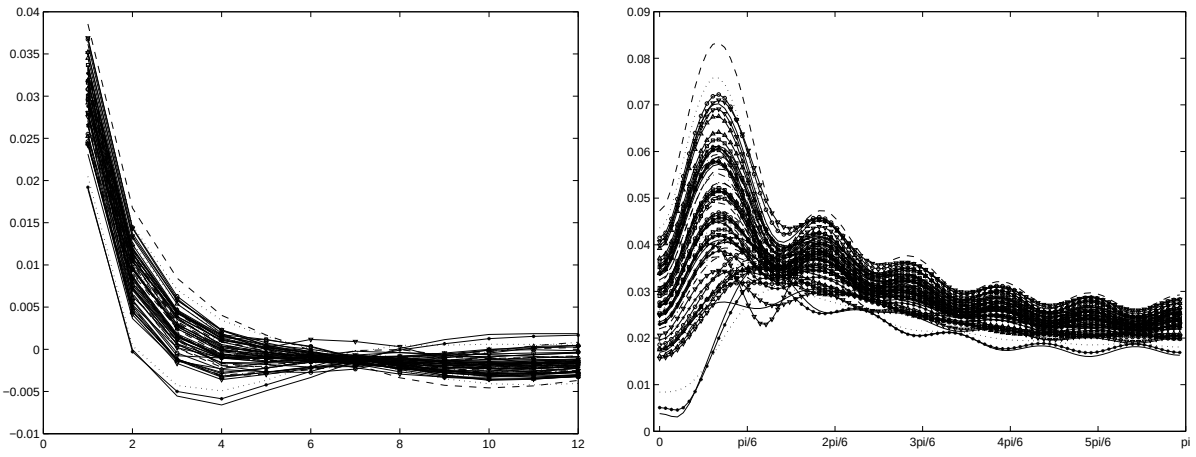

(a) (b)

Fig. 2.33. A 72-variables filter with 2nd-order constraint, $h = -3$, $\lambda_d = 0.2$, $\lambda_c = 5$

Note: (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 Fig. 2.33(a).

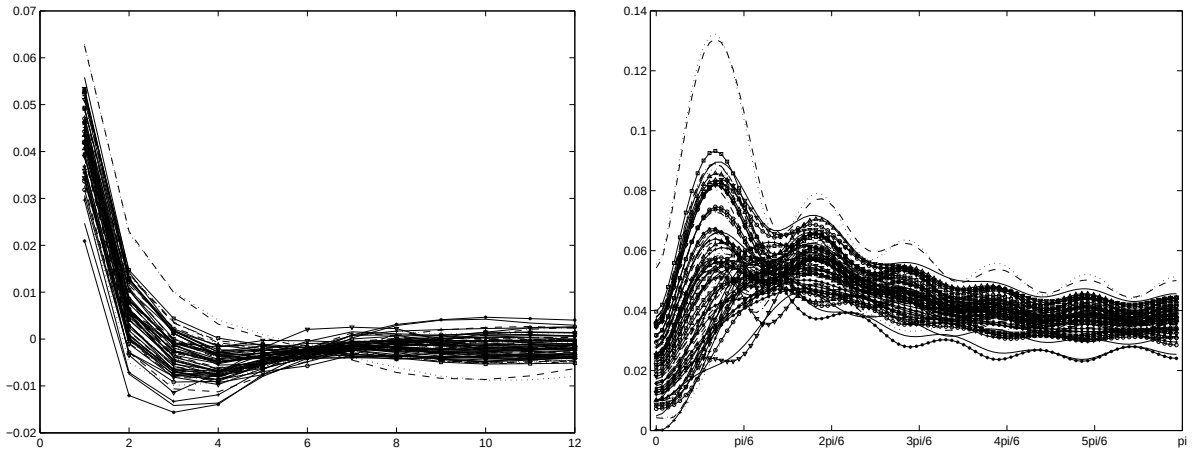

(a) (b)

Fig. 2.34. A 72-variables filter with 2nd-order constraint, $h = -6$, $\lambda_d = 0.2$, $\lambda_c = 5$

Note: (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 Fig. 2.34(a).

The corresponding effective degrees of freedom are 16 (for three-months lead) and 23 (for six-months lead), as opposed to 5 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 Fig. 2.35(a) and 2.35(b), respectively, together with Eurocoin.

Fig. 2.35(a) and 2.35(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 Fig. 2.35(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 Fig. 2.35(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 2.2).

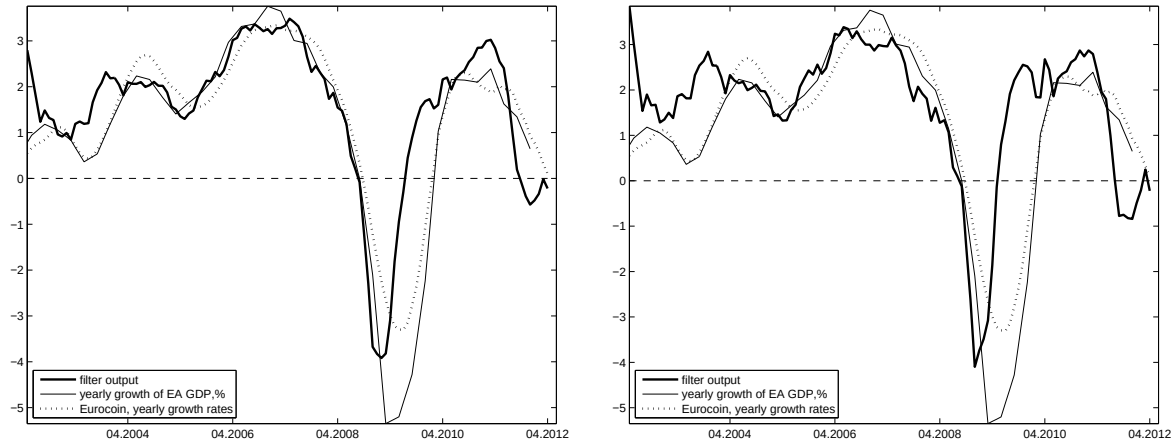

(a) (b)

Fig. 2.35. Filter outputs targeting a 3-months lead and a 6-months lead

Note: (a) Filter output corresponding to filter coefficients in Fig. 2.33(a) (targeting a 3-months lead) versus Eurocoin (yearly growth rates). (b) Filter output corresponding to filter coefficients in Fig. 2.34(a) (targeting a 6-months lead) versus Eurocoin (yearly growth rates).

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.

Tracking trendcycle in quarterly growth of GDP

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. Survey data are not transformed. A full list of data transformations is presented in Table 2.3.

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 non-transformed 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 3 e.d.f., thus more cross-sectional shrinkage is unnecessary. Filter coefficients and amplitudes are plotted in Fig. 2.36(a) and 2.36(b).

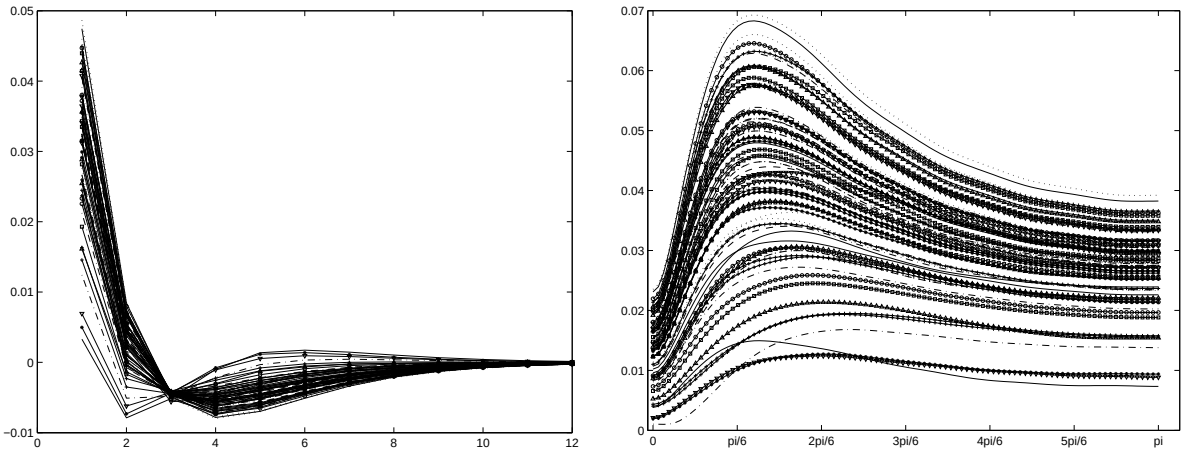

(a) (b)

Fig. 2.36. A 72-variables filter with both 1st- and 2nd-order constraints, $h = -6$, $\lambda_d = 0.4$, $\lambda_c = 1$

Note: (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 Fig. 2.36(a).

The resulting real-time filter output is plotted in Fig. 2.37 along with Eurocoin.

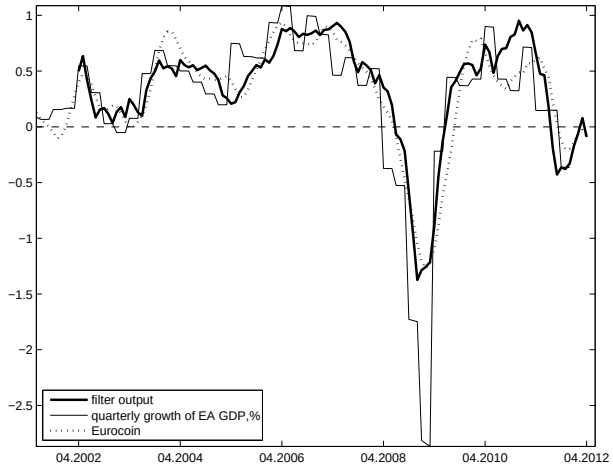

Fig. 2.37. Filter output corresponding to filter coefficients in Fig. 2.36(a) versus Eurocoin

Fig. 2.37 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 Fig. 2.37, the peak correlation is located at a one month lag w.r.t. GDP, with the second highest correlation being at a zero month lag w.r.t. GDP (Table 2.2).

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 Fig. 2.38(a) and 2.38(b).

Additional noise suppression can be achieved by suppressing amplitudes in the stop-band with positive *expw* parameter, see expression (3.6). 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 Fig. 2.39(a) and 2.39(b).

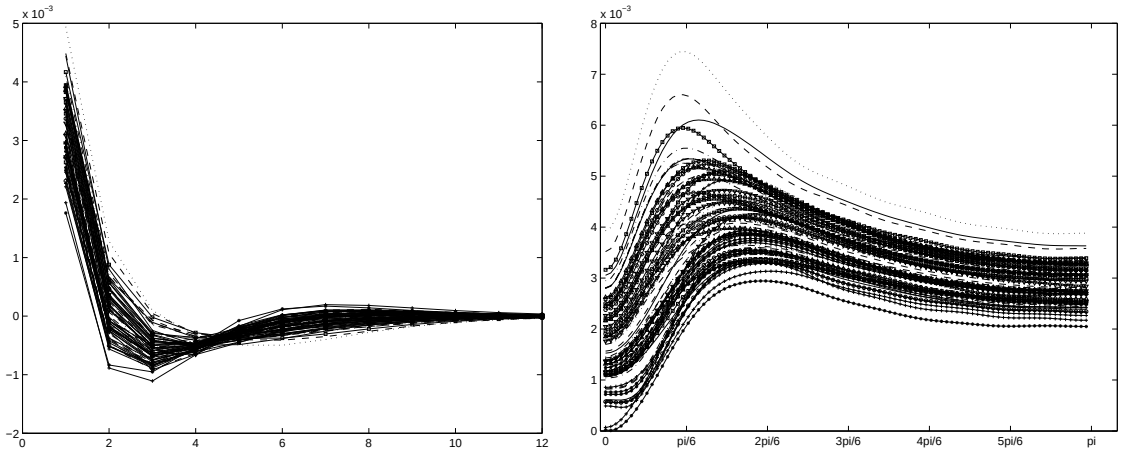

(a) (b)

Fig. 2.38. A 72-variables filter with 2nd-order constraint, $h = -3$, $\lambda_d = 0.4$, $\lambda_c = 5$

Note: (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 Fig. 2.38(a).

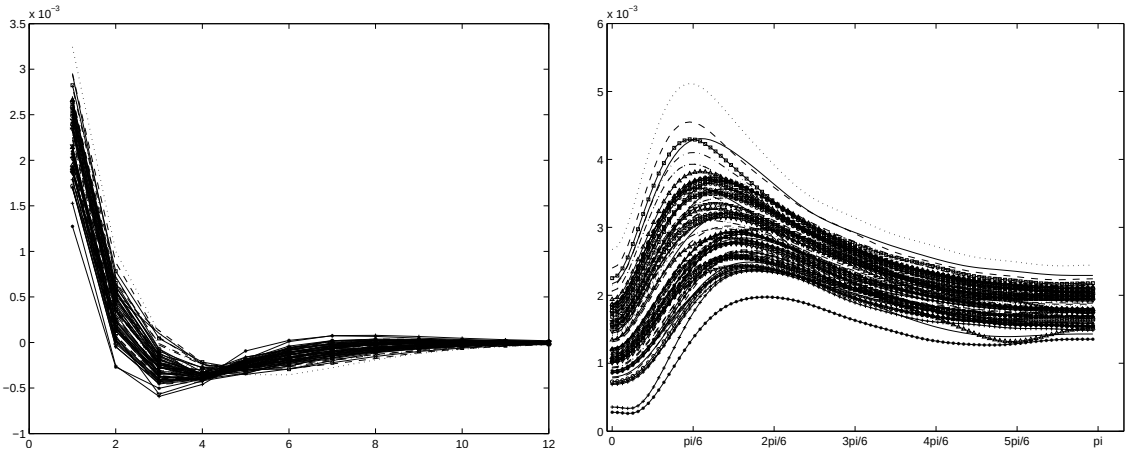

(a) (b)

Fig. 2.39. A 72-variables filter with 2nd-order constraint, $h = -3$, $\lambda_d = 0.4$, $\lambda_c = 5$, $expw = 0.5$

Note: (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 Fig. 2.39(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 Fig. 2.40(a) and 2.40(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 2.2).

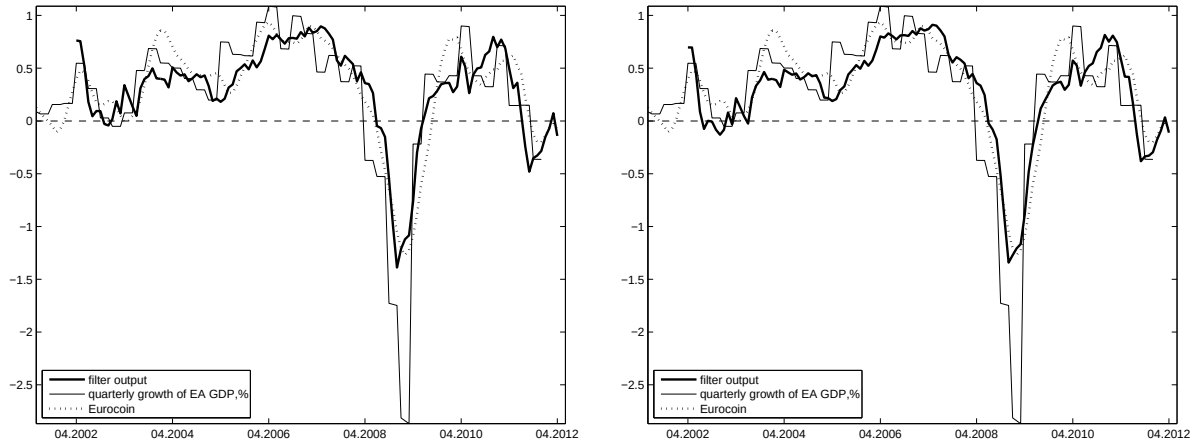

(a) (b)

Fig. 2.40. Filter output with and without noise suppression

Note: (a) Filter output corresponding to filter coefficients in Fig. 2.38(a) versus Eurocoin. (b) Filter output corresponding to filter coefficients in Fig. 2.39(a) versus Eurocoin.

Increasing targeted lead to $h = -6$ and noise suppression to $expw = 1$, yields filter coefficients and amplitudes shown in Fig. 2.41(a) and 2.41(b), respectively, and the real-time filter output in Fig. 2.42.

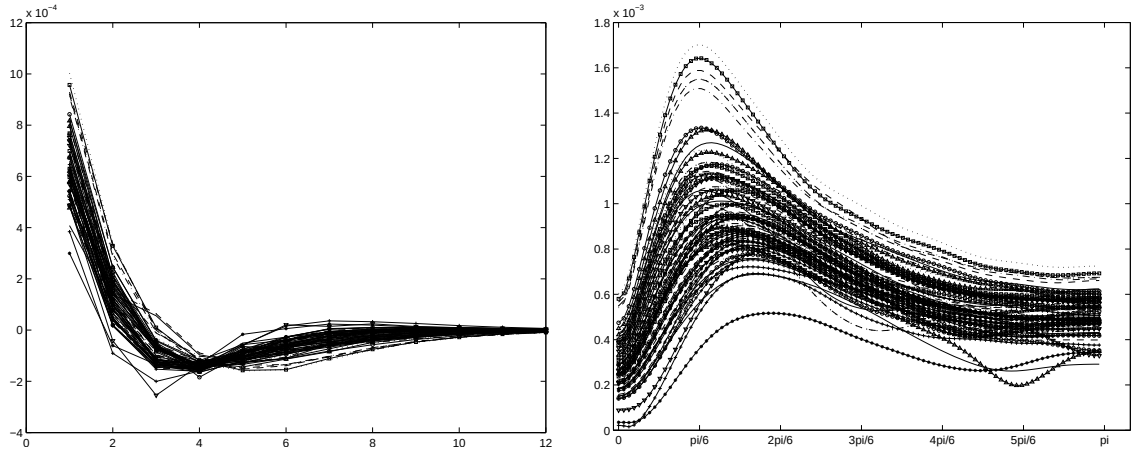

(a) (b)

Fig. 2.41. A 72-variables filter with 2nd-order constraint, $h = -6$, $\lambda_d = 0.4$, $\lambda_c = 5$, $expw = 1$

Note: (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 Fig. 2.41(a).

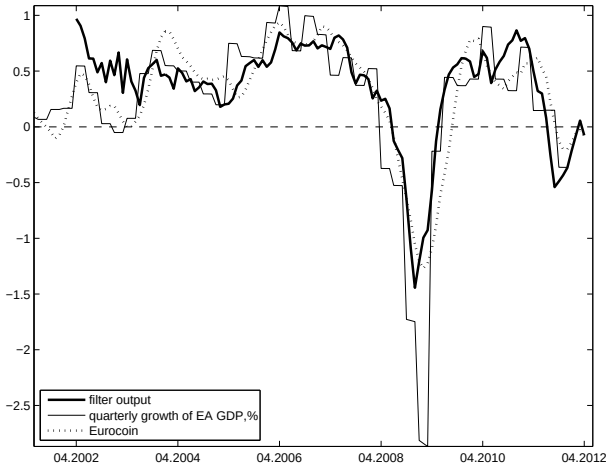

Fig. 2.42. Filter output corresponding to filter coefficients in Fig. 2.41(a) versus Eurocoin

The in-sample size at the beginning of estimation is evidently too short but soon after the filter output stabilizes and evidently 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 2.2).

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

2.2.4 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. Table 2.4 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 7 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 Fig. 2.43(a) and 2.43(b).

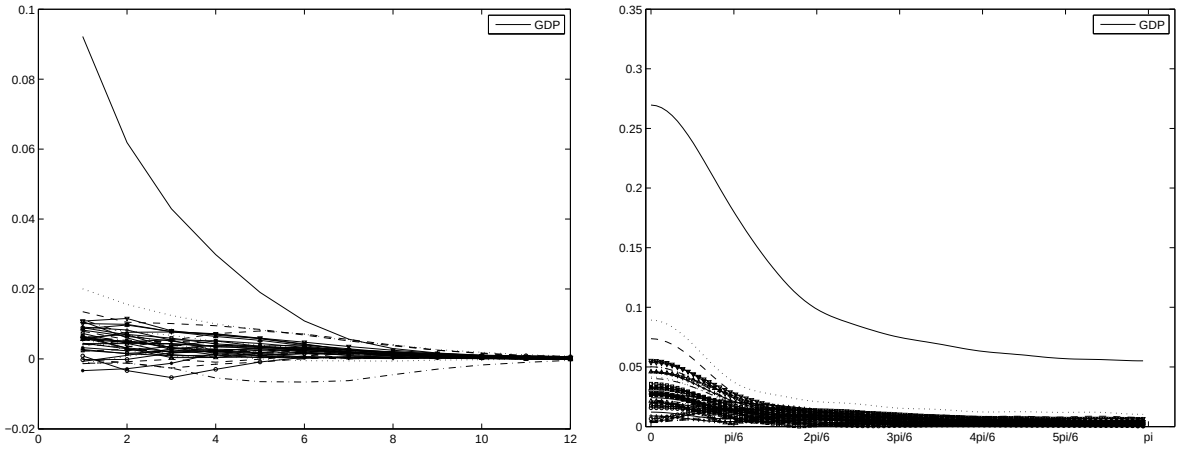

(a) (b)

Fig. 2.43. A 41-variables filter, $\lambda_d = 0.2$, $\lambda_c = 0.2$

Note: (a) Coefficients for a 41-variables filter, $\lambda_d = 0.2$, $\lambda_c = 0.2$. (b) Filter amplitudes corresponding to the coefficients in Fig. 2.43(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 Fig. 2.44, 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)).

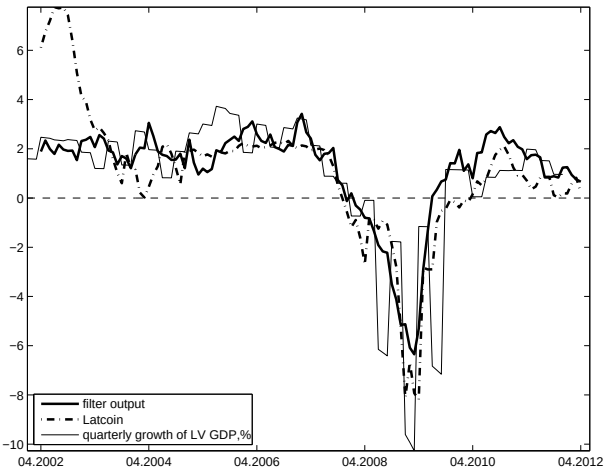

Fig. 2.44. Filter output corresponding to filter coefficients in Fig. 2.43(a) versus Latcoin

Fig. 2.44 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 2.2).

2.2.5 True real-time out-of-sample performance

The indicator tracking the trendcycle of quarterly growth of the euro area GDP has been tested and used at Bank of Latvia by producing its outputs each month since it has been developed at around April, 2012. Since it is May, 2013 when this text is written, it means there is a 13-month long sample of true real-time out-of-sample performance of this indicator. Its output, together with the realized final GDP and the alternative indicator - Eurocoin (produced by Banca d'Italia) are plotted in 2.45.

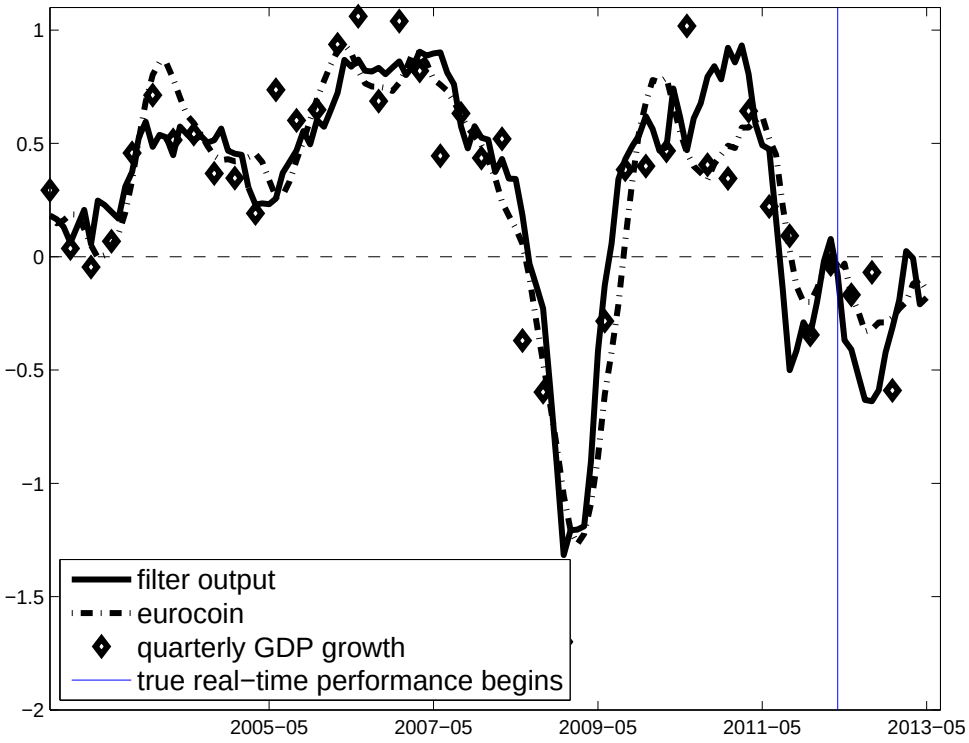

Fig. 2.45. True 13-month long real-time out-of-sample performance of the euro area GDP indicator versus Eurocoin from April, 2012 till May, 2013

Fig. 2.45 shows that the filter's true real-time performance does not seemingly deteriorate compared to the 'in-sample' period. Particularly, it does not get more volatile and captures the decline of the GDP growth well and in advance, and does not lose to Eurocoin with respect to timeliness.

This true real-time performance of the filter output reassures that the high-dimensional filtration methodology proposed in this section is well suited for practical application.

2.2.6 Tables

Table 2.2.

Dynamic correlations of indicators with GDP growth rates.

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. 2.32(b)	0.775	0.840	0.893	0.930	0.949**	0.946*	0.921	0.876	0.812	0.736
RMDFA in Fig. 2.35(a)	0.551	0.642	0.730	0.807	0.872	0.918	0.941**	0.938*	0.912	0.867
RMDFA in Fig. 2.35(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. 2.37	0.774	0.846	0.883**	0.870*	0.819	0.736				
RMDFA in Fig. 2.40(a)	0.751	0.824	0.869**	0.862*	0.817	0.740				
RMDFA in Fig. 2.40(b)	0.790	0.849	0.876**	0.856*	0.799	0.709				
RMDFA in Fig. 2.42	0.680	0.780	0.849	0.874**	0.858*	0.810				
LATCOIN in Fig. 2.44	0.685	0.716	0.747**	0.744*	0.713	0.668				
RMDFA in Fig. 2.44	0.664	0.723	0.769**	0.762*	0.755	0.742				

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

2.2.7 List of data

Table 2.3.

The EA dataset

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 2.4.

The LV 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

2.2.8 Conclusions

1. Nowadays, information is abundant. Statistical tools are being developed that are suitable to process large information for a particular problem at hand.

2. This section studies the regularized multivariate direct filter approach as a tool for signal extraction and forecasting using high-dimensional datasets. The section studies filter's properties by tracking the medium-to-long-run component of the euro area gross domestic product growth using up to a 72-variable filter.

3. 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.

4. 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.

5. The proposed high-dimensional filtering is found to be a promising tool for both signal extraction and forecasting using high-dimensional datasets and is a decent alternative to the dynamic factor methodology.

3 ROBUSTNESS OF TRADITIONAL METHODS AND FORECASTING SYSTEM OVERVIEW

3.1 Bayesian Minnessota prior

Bayesian inference requires an analyst to set a prior. Setting the right prior is crucial for precise forecasts. This section analyzes how optimal Litterman prior changes when there is a sudden change in dynamics of the target variable. By an 'optimal Litterman prior' in this section we define Litterman hyperparameters that minimize the root mean squared error from one-period ahead forecasts. For this task, an autoregressive distributed lag model (ARDL) is chosen. The prior is set up like in Litterman (1979). The model is solved by 'mixed estimation' set forth in Theil and Goldberger (1961). Latvia's gross domestic product (GDP) was found to be well suited for the analysis. The results show that a sharp economic slowdown changes the structure of the optimal weight prior by setting smaller weight on the lagged dependent variable compared to variables containing more recent information. Thus, Bayesian Minnessota prior is not robust against rapid change in the dynamics of the target variable.

3.1.1 Introduction

Bayesian inference requires an analyst to set a prior. Setting the right prior is crucial for precise forecasts. This section analyzes how optimal Litterman prior changes when an economy is hit by a recession. By an 'optimal Litterman prior' in this section we define Litterman hyperparameters that minimize the root mean squared error from one-period ahead forecasts.

Although the question about what hyperparameters to use has been addressed in a series of papers by, among others, Litterman and coauthors (Litterman (1979), Doan, Litterman and Sims (1984), Litterman (1986)) and LeSage and coauthors (LeSage and Magura (1991), LeSage and Pan (1995), LeSage and Krivelyova (1999), LeSage (1999)), the role of a business cycle on the optimal prior, to the best of our knowledge, has not been discussed. Thus, this section analyzes how (if any) prior hyperparameters should be altered for the best one-period ahead forecasting performance when there is a switch in a phase of a business cycle. For this task, an autoregressive distributed lag model

(ARDL) is chosen. The prior is set up like in Litterman (1979). The model is solved by 'mixed estimation' set forth in Theil and Goldberger (1961). Latvia's gross domestic product (GDP) was found to be well suited for the analysis. The results show that a sharp economic slowdown changes the optimal prior in two directions.

First, a lagged dependent variable loses its dominance as the key explanatory variable and, instead, more current information contained in leading indicator-type variables is of greater importance to improve forecasts. This changes the structure of the optimal weight prior, setting smaller weight on the lagged dependent variable compared to variables containing more recent information.

Second, greater uncertainty brought by a swift economic downturn requires more space for coefficient variation, which is set by the overall tightness parameter. Particularly, the results show that, in economic downturn, the optimal overall tightness parameter may increase to such an extent that Bayesian ARDL becomes equivalent to frequentist ARDL, which may imply that a greater uncertainty in an economy requires more skills from an analyst to set the right prior such that, during great economic uncertainty, one may become more comfortable using frequentist rather than Bayesian inference.

The section is organized as follows. Subsection 3.1.2 describes the model and its estimation procedure. Subsection 3.1.3 presents the results from a case study. Finally, Subsection 3.1.4 concludes.

3.1.2 Bayesian ARDL model

The Model

Consider an autoregressive distributed lag model (ARDL) of order (p, q):

$$y_t = \sum_{m=1}^{p} \beta_m y_{t-m} + \sum_{n=0}^{q} \gamma_n' x_{t-n} + \xi' z_t + \epsilon_t \tag{3.1}$$

where y_t is the dependent variable, x_t is a $d \times 1$ vector of key explanatory variables $x = [x_1\ x_2\ \ldots\ x_d]$, z_t is (a vector of) other explanatory variable(s) potentially containing a constant, a dummy variable for an outlying effect, etc., and $\epsilon_t \sim N(0, \sigma^2)$. The Bayesian prior is set to

$$\begin{aligned} \beta_m &\sim N(\mathbb{1}_{\{1\}}(m), \sigma_m^2) \\ \gamma_{in} &\sim N(0, \sigma_{in}^2) \end{aligned} \tag{3.2}$$

where $\mathbb{1}_{\{\}}()$ is an indicator function, $m = 1, 2, \ldots, p$, $i = 1, 2, \ldots, d$, and $n = 0, 1, \ldots, q$. The specification of the standard deviation of the prior is *à la* Doan, Litterman and Sims

(1984):

$$\begin{aligned}\sigma_m &= \theta k m^{-\phi} \\ \sigma_{in} &= \theta l (1+n)^{-\phi} \left(\frac{\hat{\sigma}_{u,i}}{\hat{\sigma}_{u,y}} \right)\end{aligned} \tag{3.3}$$

where $\hat{\sigma}_{u,y}$ and $\hat{\sigma}_{u,i}$ are the standard errors from a univariate autoregression involving y and x_i, respectively, so that $\hat{\sigma}_{u,i}/\hat{\sigma}_{u,y}$ is a scaling factor that adjusts for varying magnitudes of the involved variables. The parameter θ is referred as the overall tightness. The terms $m^{-\phi}$ and $(1+n)^{-\phi}$ are referred as lag decay functions for y and x_i, respectively, with $\phi \geq 0$ reflecting a shrinkage of the standard deviation with increasing lag length. The parameters k and l specify the relative tightness of the prior for variables y and x_i, respectively. Note that, for simplicity, we set l the same for all x_i.

Estimation

The model (3.1) to (3.3) can be estimated using the 'mixed estimation' method set forth in Theil and Goldberger (1961). For ease of exposition, drop z_t from (3.1) and rewrite it as

$$y = X\beta + \epsilon \tag{3.4}$$

where y is the $T \times 1$ vector of observations on the dependent variable, X the $T \times (p + (q+1)d)$ matrix of observations on the explanatory variables with rank $p + (q+1)d$, β the $(p + (q+1)d) \times 1$ vector of coefficients, and ϵ the $T \times 1$ vector of disturbances such that

$$E\epsilon = 0, \qquad \Sigma := E(\epsilon\epsilon') = \sigma^2 I_{T \times T}. \tag{3.5}$$

The Bayesian prior is included in

$$r = R\beta + \nu, \tag{3.6}$$

where r is a $(p + (q+1)d) \times 1$ vector $[1\ 0\ 0\ \ldots\ 0]'$, R is a $(p + (q+1)d) \times (p + (q+1)d)$ identity matrix, and ν is a $(p + (q+1)d) \times 1$ vector of disturbances such that

$$E\nu = 0 \tag{3.7}$$

and $E(\nu\nu')$ is a $(p+(q+1)d)\times(p+(q+1)d)$ diagonal matrix with diagonal elements being the variances specified in (3.3),

$$\Omega := E(\nu\nu') = \begin{bmatrix} \sigma_1^2 & 0 & & & \cdots & & & 0 \\ 0 & \sigma_2^2 & & & & & & \\ 0 & 0 & \ddots & & & & & \\ & & & \sigma_p^2 & & & & \\ \vdots & & & & \sigma_{10}^2 & & & \vdots \\ & & & & & \sigma_{11}^2 & & \\ & & & & & & \ddots & \\ & & & & & & \sigma_{d,q-1}^2 & 0 \\ 0 & & & & \cdots & & 0 & \sigma_{dq}^2 \end{bmatrix} \tag{3.8}$$

The sample and the independent extraneous information may be combined by writing

$$\begin{bmatrix} y \\ r \end{bmatrix} = \begin{bmatrix} X \\ R \end{bmatrix} \beta + \begin{bmatrix} u \\ \nu \end{bmatrix}; \quad E\begin{bmatrix} u \\ \nu \end{bmatrix} = 0; \quad E\left(\begin{bmatrix} u \\ \nu \end{bmatrix}\begin{bmatrix} u' & \nu' \end{bmatrix}\right) = \begin{bmatrix} \Sigma & 0 \\ 0 & \Omega \end{bmatrix}. \tag{3.9}$$

An application of generalized least squares (GLS) procedure leads to estimating β as

$$\hat{\beta} = \left(\begin{bmatrix} X' & R' \end{bmatrix}\begin{bmatrix} \Sigma & 0 \\ 0 & \Omega \end{bmatrix}^{-1}\begin{bmatrix} X \\ R \end{bmatrix}\right)^{-1}\begin{bmatrix} X' & R' \end{bmatrix}\begin{bmatrix} \Sigma & 0 \\ 0 & \Omega \end{bmatrix}^{-1}\begin{bmatrix} y \\ r \end{bmatrix} \tag{3.10}$$

or

$$\hat{\beta} = \left[X'\Sigma^{-1}X + R'\Omega^{-1}R\right]^{-1}\left[X'\Sigma^{-1}y + R'\Omega^{-1}r\right]. \tag{3.11}$$

Normalizing R:

$$\tilde{R} := \begin{bmatrix} \frac{\sigma}{\sigma_1} & 0 & & & \cdots & & & 0 \\ 0 & \frac{\sigma}{\sigma_2} & & & & & & \\ & & \ddots & & & & & \\ & & & \frac{\sigma}{\sigma_p} & & & & \\ \vdots & & & & \frac{\sigma}{\sigma_{10}} & & & \vdots \\ & & & & & \frac{\sigma}{\sigma_{11}} & & \\ & & & & & & \ddots & \\ & & & & & & \frac{\sigma}{\sigma_{d,q-1}} & 0 \\ 0 & & & & \cdots & & 0 & \frac{\sigma}{\sigma_{dq}} \end{bmatrix}$$

and r:

$$\tilde{r} := \begin{bmatrix} \frac{\sigma}{\sigma_1} & 0 & 0 & \ldots & 0 \end{bmatrix}$$

gives $E(\nu\nu') = \sigma^2 I$, and the GLS estimator in (3.11) reduces to an ordinary least squares estimator:

$$\hat{\beta} = \left[X'X + \tilde{R}'\tilde{R}\right]^{-1}\left[X'y + \tilde{R}'\tilde{r}\right]. \tag{3.12}$$

3.1.3 Results

The dependent variable of the model (3.1) is Latvia's quarterly GDP series from $1995Q1$ till $2009Q1$. The key explanatory variables x are two quarterly series, the output in manufacturing industry (according to Nace revision 1.1 subsequently called D) and output in electricity, gas and water supply industry (E). All three series are chained priced as of year 2000 and twice regularly and once seasonally differenced. The second regular difference is performed for better forecasting performance during the latter part of the GDP series due to a sharp economic downturn (see Buss, 2009 for a discussion). Series D and E are published before the GDP flash estimate is released, thus we can potentially use these series to forecast GDP before its other components are known. The model may contain a constant and other explanatory variables, all contained in z in (3.1). All calculations are performed in Scilab with the aid of its econometrics toolbox Grocer.

Warm-up

To start, Table 3.1 shows root mean squared forecast errors (RMSE) for the whole sample, the first half of the sample (RMSE1sthalf) and the second half of the sample (RMSE2ndhalf) from one-period ahead pseudo real-time forecasts beginning at sample size 17 from simple benchmark seasonal autoregressive moving average model (SARMA), autoregressive models (AR), and frequentist and Bayesian autoregressive distributed lag models (FARDL and BARDL, respectively) of order (p, q) with explanatory variable in parenthesis.

Notation (D+E) means the variables are summed to result in a single explanatory variable. The Bayesian counterpart of ARDL requires to specify the hyperparameters for (3.3), called Litterman prior consisting of four parameters, k, l, θ, and ϕ, with $w := [k\ l]$ for one-dimensional x. The forecasts are called *pseudo* real-time because they are made on the revised values of left-hand-side and right-hand-side variables in (3.1); although the revisions for the specific variables used in this analysis tend to be relatively small, they might underestimate RMSE. Nonetheless, this does not harm for our purpose.

Table 3.1.

A brief comparison of SARMA, AR, FARDL and BARDL

Model	RMSE	RMSE1sthalf	RMSE2ndhalf
SARMA(01)(01)	0.0328737	0.0160291	0.0436398
AR(1)	0.0275043	0.0194567	0.0336810
AR(2)	0.0263058	0.0203990	0.0311106
FARDL(1,0)(D)	0.0277540	0.2011203	0.0330832
FARDL(2,0)(D)	0.0289995	0.0272706	0.0306310
FARDL(2,1)(D)	0.0253833	0.0196827	0.0300202
FARDL(2,1)(E)	0.0257016	0.0216257	0.0292142
FARDL(2,1)(D+E)	0.0247125	0.0220415	0.0271218
FARDL(3,2)(D)	0.0260984	0.0216730	0.0298754
FARDL(3,2)(E)	0.0257382	0.0217008	0.0292230
FARDL(3,2)(D+E)	0.0253316	0.0251711	0.0254912
BARDL(2,1)(D+E)(.95,.1,.8,0)	0.0239113	0.0196482	0.0275217
BARDL(2,1)(D+E)(.05,1,2,0)	0.0264237	0.0258526	0.0269828
BARDL(3,2)(D+E)(1,.35,.2,0)	0.0223288	0.0171109	0.0265400
BARDL(3,2)(D+E)(.8,.25,.2,0)	0.0225414	0.0166686	0.0271732

Note: The two latter models are specified by their orders, (p, q), key exogenous variables, e.g. (D+E), and the Bayesian ARDL with a single key exogenous variable is specified by its prior, (k,l,θ,ϕ), where prior weight $w := [k\ l]$. The least RMSE in each sample space is framed.

The sample is split in halves because the first half contains a smooth growth whereas the second half contains rapid economic downturn (see the GDP series in Fig. 3.1), so we can analyze how the forecasting performance of the models changes with the business cycle and, especially, how Bayesian prior has to be altered for the best forecasting performance.

The least RMSE in each column is framed. It can be seen that Bayesian ARDL models compare well with other models. It can also be seen that the BARDL(3,2) models give the most precise one-period ahead forecasts for the whole sample as well as for the first half of the sample among all the ARDL models considered, but they are outperformed by FARDL for the second half of the series. This observation suggests that the optimal Bayesian prior might be different for the first half of the model (smooth positive growth) compared to the second half of the sample when there is a rapid economic downturn. We check this hypothesis further by employing grid search for the optimal prior.

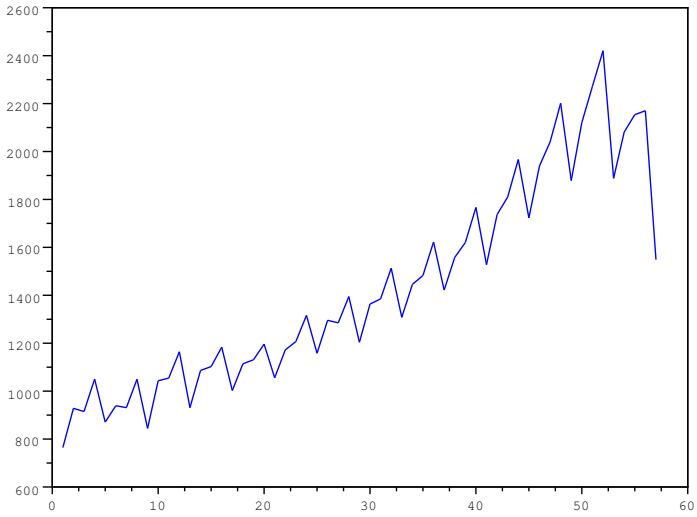

Fig. 3.1. Latvia's seasonally unadjusted GDP series

Note: Series spans from 1995Q1 till 2009Q1. Horizontal axis represents time.

Search for optimal priors

First, the grid search is performed for BARDL(2,1)(D+E). The weight vector $[k\ l]$ is 2-dimensional, one element, k, for the dependent variable and one, l, for a single explanatory variable x, both ranging from .05 to 1 with step size .05. The overall tightness, θ, is set to range from .6 to 2.5 with step .1, and the lag decay, ϕ, from 0 to 1 with step .2. So, the grid size is $20 \times 20 \times 20 \times 6$ containing overall 48000 prior combinations for each one-period ahead forecast with sample size ranging from 17 to 51. The minimum RMSE for the whole sample is attained at the coordinate [19 2 3 1] with the corresponding values $[k\ l\ \theta\ \phi]$=[.95 .1 .8 0] with a boundary value at $\phi = 0$. The boundary for ϕ can not be decreased further since negative values would presume lags of a higher order be more informative which is counterintuitive.

Fig. 3.2(a) and 3.2(b) show the inverse of the RMSE as a function of the prior for the whole sample. Fig. 3.2(a) shows the inverse of the RMSE as a function of the weight vector (the x and y axes represent k and l, respectively) given the rest of parameters, θ and ϕ, at their RMSE-minimizing values. It can be seen that the values of k have the major impact on the RMSE with acceptable range about (.4,1), otherwise the RMSE increases substantially. On the contrary, values of l have less influence on the RMSE given k, nonetheless, a peak is evident at $l = .1$ for all acceptable values of k.

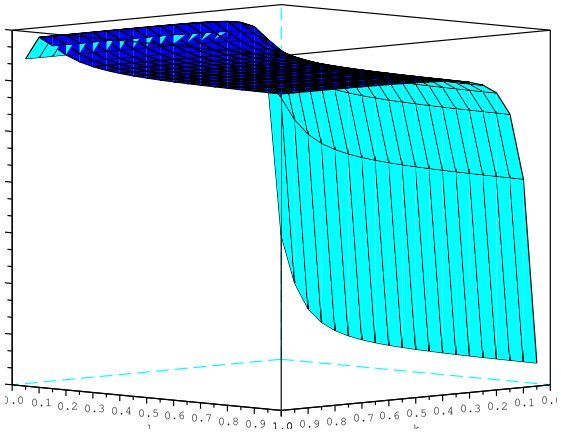

(a) Full sample. Optimal $k = .95$ and optimal $l = .1$.

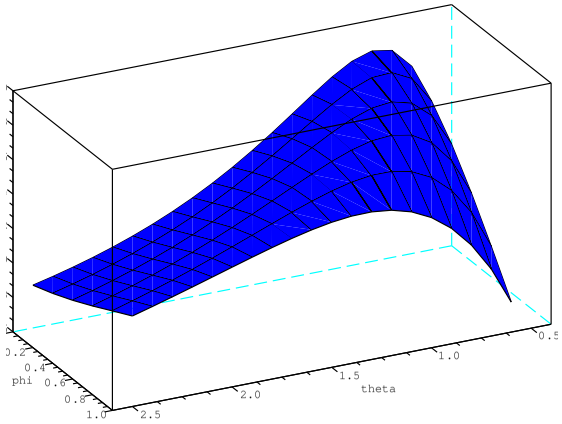

(b) Full sample. Optimal $\theta = .9$ and optimal $\phi = 0$.

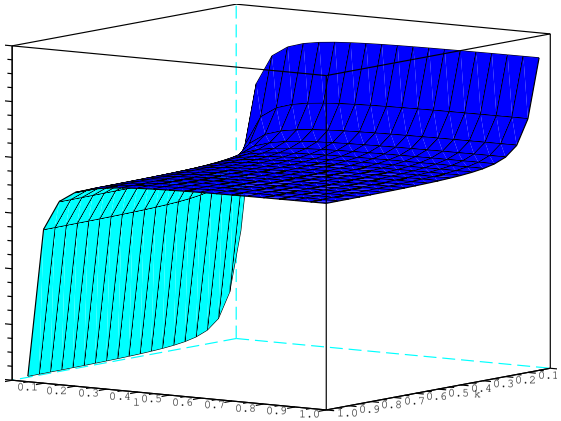

(c) Second half of the sample. Optimal $k = .05$ and optimal $l = 1$.

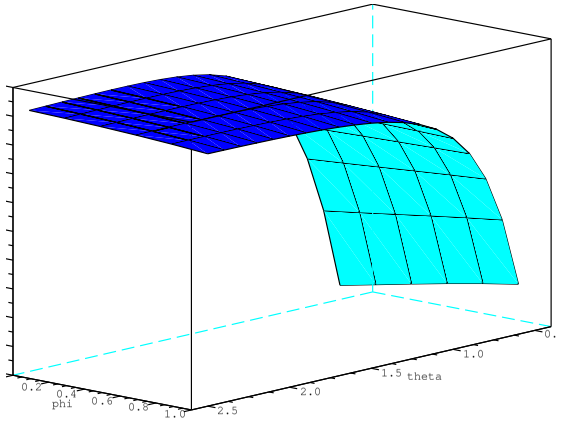

(d) Second half of the sample. Optimal $\theta = 2$ and optimal $\phi = 0$.

Fig. 3.2. Results from grid search for optimal prior for BARDL(2,1)(D+E)

Note: Fig. 3.2(a) and 3.2(b) represent a full sample, whereas Fig. 3.2(c) and 3.2(d) represent the second half of the sample. The figures on the left (3.2(a) and 3.2(c)) show $RMSE^{-1}$ (z axis) as a function of a weight vector (k,l) (x and y axis, respectively) at the RMSE-minimizing θ and ϕ. The figures on the right (3.2(b) and 3.2(d)) show $RMSE^{-1}$ (z axis) as a function of θ and ϕ (x and y axis, respectively) at the RMSE-minimizing weight vector.

Similarly, Fig. 3.2(b) shows the inverse of the RMSE as a function of θ and ϕ (representing x and y axes, respectively) given the RMSE-minimizing weight vector. It can be seen that the values of both θ and ϕ have a nontrivial impact on RMSE at its optimum with the maximizing values .8 and 0, respectively. The maximizing value of $\phi = 0$ might be due to the small number of lags, which is one for each RHS variable in this model.

Now, calculating the minimum RMSE for the second half of the sample, the optimum value is attained at the coordinate [1 20 15 1] with the corresponding values [k l θ ϕ]=[.05 1 2 0] with three boundary values for k, l and ϕ. It can already be seen that the optimal prior weight is different compared to the full sample. Fig. 3.2(c) and 3.2(d) show the inverse of the RMSE as a function of the prior for the second half of the sample. Fig. 3.2(c) looks almost like the inverse of Fig. 3.2(a). Now, the RMSE is increasing with k, with an optimum at the lowest k considered; other values of k would significantly increase the RMSE at all levels of l, the latter being also critical for optimal RMSE with acceptable range about (.3,1), otherwise the forecast error increases substantially. This observation is in line with our hypothesis that, during sharp decline in the economy, explanatory variables containing most recent information are more important than the lagged dependent variable.

Fig. 3.2(d) shows that, for the second half of the sample, the optimal tightness parameter is higher compared to the full sample, with acceptable values in about (1,2.5), otherwise the forecast error increases substantially. This observation is as expected since the model coefficients should be given more flexibility during a rapid change in an economy. For acceptable θ, the values of lag decay parameter, ϕ, is of less importance. The forecasting performance of BARDL(2,1)(D+E) for the first half of the sample is not impressive and thus not presented here.

Having explored BARDL(2,1)(D+E), we now check the results for BARDL(3,2) (D+E) whose forecasting performance for all sample spaces considered, as it can be seen in Table 3.1, is promising. The grid space is formed by k and l being from .05 to 1 with step .05, θ from .1 to 1 with step .1, and ϕ from 0 to 1 with step .1. The coordinate for the least RMSE for full sample is [20 7 2 1] with the prior values [k l θ ϕ]=[1 .35 .2 0], showing some resemblance with the results for BARDL(2,1)(D+E). The inverse RMSE for full sample around the optimal prior values is shown in Fig. 3.3(a) and 3.3(b). The behavior of the inverse RMSE around its optimal value is similar to that of BARDL(2,1)(D+E).

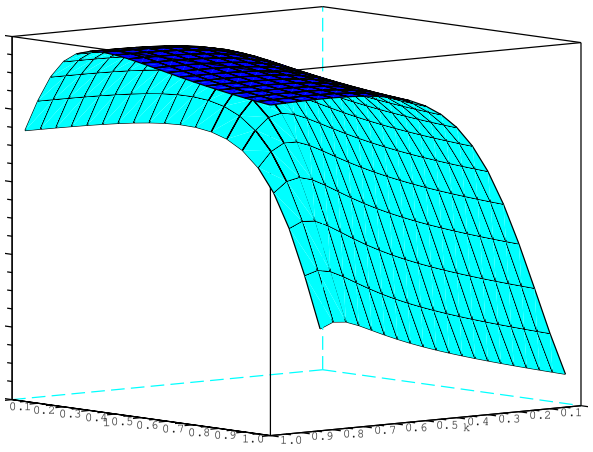

(a) Full sample. Optimal $k = 1$ and optimal $l = .35$.

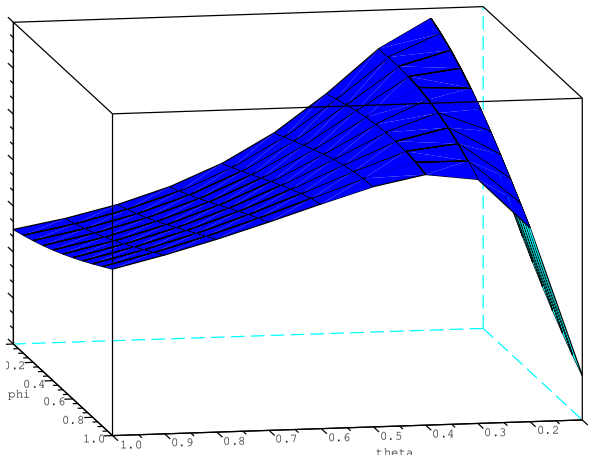

(b) Full sample. Optimal $\theta = .2$ and optimal $\phi = 0$.

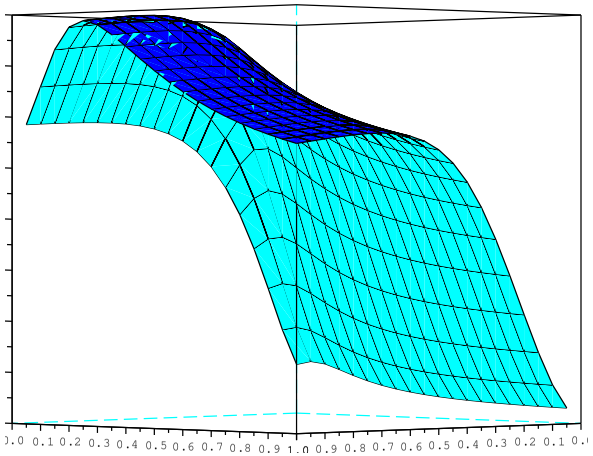

(c) First half of the sample. Optimal $k = .8$ and optimal $l = .25$.

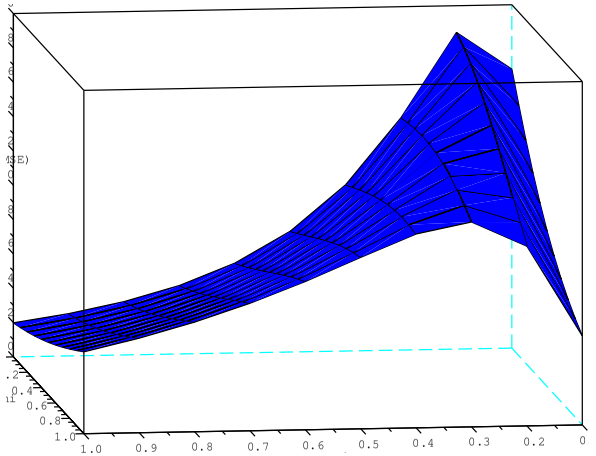

(d) First half of the sample. Optimal $\theta = .2$ and optimal $\phi = 0$.

Fig. 3.3. Results from grid search for optimal prior for BARDL(3,2)(D+E)

Note: Fig. 3.3(a) and 3.3(b) represent a full sample, whereas Fig. 3.3(c) and 3.3(d) represent the first half of the sample. The figures on the left (3.3(a) and 3.3(b)) show $RMSE^{-1}$ (z axis) as a function of a weight vector (k,l) (x and y axis, respectively) at the RMSE-minimizing θ and ϕ. The figures on the right (3.3(c) and 3.3(d)) show $RMSE^{-1}$ (z axis) as a function of θ and ϕ (x and y axis, respectively) at the RMSE-minimizing weight vector.

We can see from Table 3.1 about the model's BARDL(3,2)(D+E) comparatively competitive forecasting performance for the first half of the sample. Fig. 3.3(c) and 3.3(d) show the inverse RMSE around its optimum as a function of prior parameters for the first half of the sample. We see that the results are similar to the results from a full sample with optimal $k = .8$, $l = .25$, $\theta = .2$ and $\phi = 0$. It can also be seen that l has more influence on the RMSE compared to the full sample, with lowest RMSE concentrating on the lowest part of l space.

Regarding the results for the second half of the sample, the coordinate of the optimal value is [20 20 10 1], with all values being at a boundary and suggesting a greater θ (i.e., more flexibility for coefficient values). An extensive search for the optimal θ resulted to its value around 10^5 with RMSE being the same as for FARDL(3,2)(D+E) at least up to and including the 7^{th} digit after a comma, shown in Table 3.1. The latter result might suggest that during a sharp decline in an economy one might wish to set the overall tightness parameter, θ, so loose that one is more comfortable to use frequentist version of ARDL.

3.1.4 Conclusions

1. Bayesian inference requires an analyst to set priors. Setting the right prior is crucial for precise forecasts. This section analyzes how optimal prior changes with business cycle, specifically, when an economy is hit by a recession.

2. The results show that when economy is growing, the optimal overall tightness parameter is less than unity, and the optimal weight vector sets a higher weight on a lagged dependent variable compared to other explanatory variables. However, a swift economic downturn changes the optimal prior considerably in two directions.

3. First, a lagged dependent variable loses its dominance as the key explanatory variable and, instead, more current information contained in leading indicator-type variables is of greater importance to improve forecasts. This changes the structure of the weight prior, setting smaller weight on the lagged dependent variable compared to variables containing more recent information.

4. Second, greater uncertainty brought by a rapid economic downturn requires more space for coefficient variation, which is set by the overall tightness parameter. The results show that, in economic downturn, the optimal overall tightness parameter may increase to such an extent that Bayesian autoregressive distributed lag model (ARDL) becomes equivalent to frequentist ARDL.

5. The results imply that a greater uncertainty in an economy requires more skills from an analyst to set the right prior such that, during great economic uncertainty, one may become more comfortable using frequentist rather than Bayesian inference.

3.2 Factor methods

Factor methods are popular data dimension-reduction techniques. There are several types of factor methods, mainly, statically and dynamically computed factors. Although there are many papers studying the factor methodology, none studies their robustness against rapid change in the dynamics of the target variable. This section tries to fill the gap by studying the robustness of the static versus dynamic factor models under a rapid change in the dynamics of the target variable. The results show that static factors are more robust than their dynamic counterparts.

3.2.1 Introduction

The choice between static and dynamic factors in now-/forecasting GDP is unresolved. Some papers find dynamic factors superior over the static ones (see, for example, den Reijer (2005)). Other papers find little or no advantage of dynamic over static factors. For example, Schumacher (2005) finds that dynamic factors only slightly outperform static factors. D'Agostino and Giannone (2007) find static and dynamic factors perform similarly. Marcellino and Schumacher (2008), among other results, report that information content of now- and forecasts hardly change if factors are estimated by static rather than dynamic principal components analysis. Ajevskis and Davidsons (2008) also find similar performance between static and dynamic factors. Finally, there are papers that argue for static over dynamic factors. For example, Boivin and Ng (2005) state that static factors are easier to construct than dynamic factors, and are favored on practical grounds. This section contributes to the now-/forecasting literature by comparing GDP nowcasting performance of dynamic versus static factor models along a business cycle. For the fulfillment of the task, we had to choose the size of factor-forming set of variables, i.e., we had to decide whether to use a large-scale or a small-scale factor model, and what data to use.

Regarding the choice between large-scale and small-scale factor models, the following empirical evidence is observed. First, several papers on large-scale factor models compare the models only to simple benchmarks, instead of the best-performing models, and find large-scale factor models superior. For example, Siliverstovs and Kholodilin (2010) use a large-scale approximate dynamic factor model from 562 indicators and compare its now-/forecasting performance to, what they call, a naive constant-growth model, and find the factor model being superior. As another example, Ajevskis and Davidsons (2008) use large-scale static and approximate dynamic factor models from 126 indicators, compare them to a benchmark autoregressive model, and find factor models tending to be superior over the benchmark. There is another kind of papers that finds that large-scale factor models can not improve GDP now-/forecasting compared to non-factor models. For example, Banerjee, Marcellino and Masten (2010), inter alia, forecast the indus-

trial production in Germany, and find that large-scale factor models extracted from 90 monthly series can not improve upon the forecasting performance of a simple autoregressive model, and conclude that factors per se may not increase the forecasting precision of models. Likewise, Gupta and Kabundi (2008, 2009) although have a misleading abstract, find that a large-scale factor model performs worse than a vector autoregressive model in forecasting South Africa's GDP. Finally, there are papers that argue for small-scale over the large-scale factor models (see, for example, Schneider and Spitzer (2004), Boivin and Ng (2006), Caggiano, Kapetanios and Labhard (2009)). Given the lack of empirical evidence or rationale for clear advantage of large-scale over small-scale factor models in GDP now-/forecasting, our choice falls to using parsimonious, small-scale factor models.

Considering the choice of data, we choose Latvian data since it possesses a pronounced switch of business cycle phases - there is a period of high GDP growth that is followed by a rapid recession. Thus, we are able to compare nowcasting errors between two cases - when the model span and the nowcasting period are contained within a single business cycle phase versus the case when the model span and the nowcasting period stretch beyond a single business cycle phase. Although our choice falls to the Latvian data, the exercise described in the section might be repeated on any data with a pronounced switch of business cycle phases, including generated data. Considerations of using other data are left for further research.

Note that this section does not discuss the now-/forecasting performance of Markov-switching factor models (see, among others, Kim and Yoo (1995), Chauvet (1998), Kim and Nelson (1998), Chauvet and Hamilton (2005), and Camacho, Perez-Quiros and Poncela (2012)).

The section is organized as follows. Subsection 3.2.2 describes the methodology of factor models and their estimation. Subsection 3.2.3 presents the results for the nowcasting performance of static, dynamic and mixed factor models versus a random walk (RW), autoregressive (AR), and vector autoregressive (VAR) models during a smooth growth phase as well as during a pronounced switch of business cycle phases. Finally, Subsection 3.2.4 concludes.

3.2.2 Static and dynamic factor models

This section discusses the estimation of static and exact dynamic factors, and is mainly in line with Doz and Lenglart (1999) and Dubois and Michaux (2010).

Consider an $(n+1)$-dimensional vector autoregressive model of order r, VAR(r):

$$\begin{bmatrix} y_t \\ x_{1t} \\ \vdots \\ x_{nt} \end{bmatrix} = \begin{bmatrix} a_0 + a_{01}y_{t-1} + \cdots + a_{0r}y_{t-r} + \cdots + a_{011}x_{1,t-1} + \cdots + a_{0nr}x_{n,t-r} + u_{0t} \\ a_1 + a_{11}y_{t-1} + \cdots + a_{1r}y_{t-r} + \cdots + a_{111}x_{1,t-1} + \cdots + a_{1nr}x_{n,t-r} + u_{1t} \\ \vdots \\ a_n + a_{n1}y_{t-1} + \cdots + a_{nr}y_{t-r} + \cdots + a_{n11}x_{1,t-1} + \cdots + a_{nnr}x_{n,t-r} + u_{nt} \end{bmatrix}, \tag{3.13}$$

where y_t is a scalar dependent variable at time $t = 1, \ldots, T$, $x_t = (x_{1t}, \ldots, x_{nt})'$ is an $n \times 1$ vector of endogenous explanatory variables at time t, $u_t = (u_{0t}, \ldots, u_{nt})'$ is an $(n+1) \times 1$ vector of innovation processes at time t with $E(u_t) = 0$, $E(u_t u_t') = \Sigma_u$, $E(u_t u_s') = 0$ for $s \neq t$ and $t = 1, 2, \ldots$. If n is large, model (3.13) incurs in a curse-of-dimensionality problem. A cure for this problem is to use a relatively small number of factors that are weighted averages of the predictors. We will consider two types of factor extractions - static and exact dynamic. Static factors are obtained *à la* Stock and Watson (1998) as follows. It is assumed that x_t can be represented as

$$x_t = \Lambda F_t + e_t, \tag{3.14}$$

where F_t is a $k \times 1$ vector of common factors at time t, Λ is an $n \times k$ matrix of factor loadings, and e_t is an $n \times 1$ vector of white noise processes at time t. It is assumed that

$$E(y_{t+1}|F_t, x_t, y_t, F_{t-1}, x_{t-1}, y_{t-1}, \ldots) = E(y_{t+1}|F_t, y_t, F_{t-1}, y_{t-1}, \ldots). \tag{3.15}$$

The assumption in (3.15) permits the dimension reduction of the matrix of explanatory variables from n to k. F_t is obtained by principal components analysis, i.e., by selecting k eigenvectors ν_j, $j = 1, 2, \ldots, k$ (that are of unit length) of $x'x$, where $x = (x_1, \ldots, x_T)'$, associated with the largest k eigenvalues of $x'x$ and projecting x on the eigenvectors, $F_j = x\nu_j$, $j = 1, 2, \ldots, k$; F_t then is the tth column of $(F_1, \ldots, F_k)'$.

The dynamic factor model is estimated as in Doz and Lenglart (1999), that develops an exact dynamic factor model, where factors are extracted from a relatively small number of variables. The procedure is described as follows. If n is the number of the variables under study, T the number of observations for each variable, x_{it} the value taken by the x_i variable at time t, and if $F_1, \ldots, F_k$, $k < n$ are the unobservable factors, the model has the following form:

$$x_{it} = \lambda_{i1}F_{1t} + \cdots + \lambda_{ik}F_{kt} + u_{it}$$

for $i = 1, \ldots, n$ and for all t. Each common factor F_j contributes to the explanation of the x_i variable with a loading equal to λ_{ij}. The idiosyncratic terms $(u_{it})_{t \in Z}$ are assumed

to be independent of each other and independent of the common factors:

$$E(u_{it}u_{js}) = 0 \; \forall i \neq j, \; \forall (t,s)$$
$$E(u_{it}F_{js}) = 0 \; \forall (i,j), \; \forall (t,s).$$

In the model designed for individual data, the common and idiosyncratic factors are assumed to be white noises, i.e.,

$$E(u_{it}u_{is}) = 0 \; \forall i, \; \forall t \neq s$$
$$E(F_{it}F_{is}) = 0.$$

The model designed for individual data cannot be directly applied to time series, which generally show temporal autocorrelations. For this reason, it is called a static factor model. Using matrix notations

$$x_t = (x_{1t}, \ldots, x_{nt})', \; F_t = (F_{1t}, \ldots, F_{kt})'$$
$$u_t = (u_{1t}, \ldots, u_{nt})', \Lambda = (\lambda_{ij})_{\substack{1 \leq i \leq n \\ 1 \leq j \leq k}},$$

this model can be written as follows:

$$x_t = \Lambda F_t + u_t,$$

where

$$E(F_t) = 0$$
$$E(u_t) = 0$$
$$E(u_t u_t') = D = \text{diag}(d_1, \ldots, d_n)$$
$$E(F_t u_s') = 0, \; \forall (t,s), \; t \neq s$$
$$E(u_t u_s') = 0, \; \forall (t,s), \; t \neq s.$$

It is easy to see that the common factors are only defined up to a linear transformation, that is, it is always possible to premultiply F_t by any invertible matrix, as soon as Λ is postmultiplied by the inverse of the same matrix. Generally, it is assumed that $\text{Var}(F_t) = I_k$, so that F_t and Λ are defined up to a rotation matrix (at the estimation stage, they are fixed by imposing supplementary identifying constraints; see below). If it is imposed that $\text{Var}(F_t) = I_k$, then

$$\text{Var}(x_t) = \Lambda\Lambda' + D,$$

such that

$$\mathrm{Var}(x_{it}) = \sum_{j=1}^{k} \lambda_{ij}^2 + d_i, \quad i = 1, \ldots, n.$$

Each λ_{ij}^2 represents the part of x_i's variance which is explained by F_j; thus, $h_i^2 = \sum_{j=1}^{k} \lambda_{ij}^2$ represents the total contribution of the factors to x_i's variance. On the other hand, $\mathrm{Var}(u_i) = d_i$ is the part of x_i's variance which is not explained by the common factors.

There are two main methods to estimate the static model: principal components analysis (PCA) and the maximum likelihood (ML) under a Gaussian hypothesis. The first one does not need to make preliminary assumption about the number of factors, while this is necessary for the ML estimation. On the other hand, the ML gives efficient estimates of the parameters, which is not the case for PCA. Both methods are implemented as follows. At the first stage, the PCA is performed. Then, the ML estimation is run for the the appropriately chosen number of factors. Since we consider exact factor models, it is assumed that the processes (u_{it}) are uncorrelated with each other at all leads and lags. In this dynamic framework, the likelihood under the Gaussian assumption is not equal to the static model's likelihood. However, Doz and Lenglart (1999) show that, in a stationary framework, the estimators obtained by the maximization of the static model's likelihood are consistent estimators of the parameters. In brief, it is supposed that each of the real processes (F_{it}) and (u_{it}) is weakly stationary and can be autocorrelated, but that the model is estimated by a standard ML procedure as if those processes were Gaussian and were not autocorrelated. The stationarity of the processes (F_{it}) and (u_{it}) implies that the process (x_t) is stationary as well. The parameters of the model can be written in a vector $\mu = (\mathrm{vec}\Lambda', d')'$,where $d = (d_1, \ldots, d_n)'$. The estimator $\hat{\mu}_T$, which is obtained this way is then an M-estimator of μ. Doz and Lenglart (1999) show that this estimator is consistent. Shortly, denote $z_{it} = x_{it} - \bar{x}_i$ and $z_t = (z_{1t}, \ldots, z_{nt})'$ for any t, $S = \frac{1}{T}\sum_t z_t z_t'$ the empirical covariance matrix of the observations and $\Sigma = \Lambda\Lambda' + D$ the theoretical covariance matrix. The quasi-likelihood of the model is computed under the Gaussian assumption as if neither the factors, nor the idiosyncratic components were autocorrelated. Up to a constant term, the quasi-likelihood can be written as

$$\begin{aligned}\mathcal{L}_T(z, \mu) &= \frac{1}{T}\sum_{t=1}^{T} \ln I_t(z, \mu) \\ &= -\frac{1}{2}\ln(\det(\Lambda\Lambda' + D)) - \frac{1}{2}\mathrm{tr}((\Lambda\Lambda' + D)^{-1}S)\end{aligned}$$

Let μ_0 be the true value of the parameter μ. It is assumed that μ belongs to a set of the form $R^{nk} \times [\alpha, +\infty)^n$, $\alpha > 0$, which contains μ_0. Under this assumption, Σ is an invertible matrix, so the quasi-likelihood is well defined. The proof that the M-estimator $\hat{\mu}_T$, that maximizes $\mathcal{L}_T(z, \mu)$, is consistent, relies on several steps. First, Doz and Lenglart (1999) show that, in order to maximize the function on $R^{nk} \times [\alpha, +\infty)^n$, it is sufficient

to maximize the function on a compact subset of $R^{np} \times [\alpha, +\infty)^n$. Then, they show that the function has properties which are sufficient to ensure the consistency.

Given the consistency of the factor loadings, a dynamic factor model with the common factors following an ARMA(p, q) process and the idiosyncratic components following an AR(l) process can be written as

$$\begin{aligned} x_{it} &= m_i + \lambda_{i1}F_{1t} + \ldots + \lambda_{ik}F_{kt} + u_{it} \\ (1 - \phi_{j1}L - \ldots - \phi_{jp}L^p)F_{jt} &= (1 - \theta_{j1}L - \ldots - \theta_{jq}L^q)\epsilon_{jt} \\ (1 - \rho_{i1}L - \ldots - \rho_{il}L^l)u_{it} &= \xi_{it} \end{aligned} \tag{3.16}$$

for $i = 1, \ldots, n$, $j = 1, \ldots, k$ and for all t, where ϵ_{jt} and ξ_{it} are the innovations of F_t and u_{it} at time t, l is the order of the AR process governing u_{it}, and the processes (ϵ_{jt}) and (ξ_{it}) are mutually independent. For identification purposes, the variance of the factor idiosyncratic components, ϵ_{jt}, is set to take the value 0.25.

Model (3.16) can be put into the state-space representation

$$x_t = Z\alpha_t + e_t \tag{3.17}$$

$$\alpha_t = A\alpha_{t-1} + R\eta_t, \tag{3.18}$$

where the processes (e_t) and (η_t) are serially uncorrelated and mutually uncorrelated at all leads and lags, and

$$\begin{aligned} E(e_t) &= 0 \\ \text{Var}(e_t) &= H \\ E(\eta_t) &= 0 \\ \text{Var}(\eta_t) &= Q. \end{aligned}$$

In our case, the state-space form of the model, (3.17) and (3.18), is the following:

$$x_t = \begin{bmatrix} \Lambda & 0_{n\times k(p+q-1)} & I_n & 0_{n\times n(l-1)} \end{bmatrix} \begin{bmatrix} F_t \\ \vdots \\ F_{t-p+1} \\ \epsilon_t \\ \vdots \\ \epsilon_{t-q+1} \\ u_t \\ \vdots \\ u_{t-l+1} \end{bmatrix}$$

$$
\begin{bmatrix} F_t \\ \vdots \\ F_{t-p+1} \\ \epsilon_t \\ \vdots \\ \epsilon_{t-q+1} \\ u_t \\ \vdots \\ u_{t-l+1} \end{bmatrix} = \begin{bmatrix} \phi & \theta & 0_{k\times nl} \\ I_{k(p-1)\times kp} & 0_{k(p-1)\times kq} & 0_{k(p-1)\times nl} \\ 0_{k\times kp} & 0_{k\times kq} & 0_{k\times nl} \\ 0_{k(q-1)\times kp} & I_{k(q-1)\times kq} & 0_{k(q-1)\times nl} \\ 0_{n\times kp} & 0_{n\times kq} & \rho \\ 0_{n(l-1)\times kp} & 0_{n(l-1)\times kq} & I_{n(l-1)\times nl} \end{bmatrix} \begin{bmatrix} F_{t-1} \\ \vdots \\ F_{t-p} \\ \epsilon_{t-1} \\ \vdots \\ \epsilon_{t-q} \\ u_{t-1} \\ \vdots \\ u_{t-l} \end{bmatrix}
$$

$$
+ \left[\begin{array}{c|c} I_k & \\ 0_{k(p-1)\times k} & 0_{k(p+1)\times n} \\ I_k & \\ \hline & 0_{k(q-1)\times n} \\ 0_{[k(q-1)+nl]\times k} & I_n \\ & 0_{n(l-1)\times n} \end{array}\right] \begin{bmatrix} \epsilon_t \\ \xi_t \end{bmatrix},
$$

where

$$
x_t = \begin{bmatrix} x_{1t} \\ \vdots \\ x_{nt} \end{bmatrix} \qquad \Lambda = \begin{bmatrix} \lambda_{11} & \cdots & \lambda_{1k} \\ \vdots & \vdots & \vdots \\ \lambda_{n1} & \cdots & \lambda_{nk} \end{bmatrix} \qquad F_t = \begin{bmatrix} F_{1t} \\ \vdots \\ F_{kt} \end{bmatrix} \qquad \epsilon_t = \begin{bmatrix} \epsilon_{1t} \\ \vdots \\ \epsilon_{kt} \end{bmatrix}
$$

$$
u_t = \begin{bmatrix} u_{1t} \\ \vdots \\ u_{nt} \end{bmatrix} \qquad \xi_t = \begin{bmatrix} \xi_{1t} \\ \vdots \\ \xi_{nt} \end{bmatrix} \qquad \phi = \begin{bmatrix} \phi_1' \\ \vdots \\ \phi_p' \end{bmatrix}' \qquad \phi_i = \text{diag}\left(\begin{bmatrix} \phi_{1i} \\ \vdots \\ \phi_{ki} \end{bmatrix} \right)
$$

$$
\theta = \begin{bmatrix} \theta_1' \\ \vdots \\ \theta_q' \end{bmatrix}' \qquad \theta_j = \text{diag}\left(\begin{bmatrix} -\theta_{1j} \\ \vdots \\ -\theta_{kj} \end{bmatrix} \right) \qquad \rho = \begin{bmatrix} \rho_1' \\ \vdots \\ \rho_l' \end{bmatrix}' \qquad \rho_s = \text{diag}\left(\begin{bmatrix} \rho_{1s} \\ \vdots \\ \rho_{ns} \end{bmatrix} \right)
$$

and is estimated by an ML using the Kalman filter (Kalman, 1960). The initial values for F_t, Λ, and u_t are obtained from performing a static factor analysis, the initial values for ϕ and θ are obtained from running an ARMA(p, q) on F_t, and initial values for ρ and Var(ξ_t) are obtained from running an AR process on u_t.

3.2.3 Results

The dependent variable in the model (3.13) is Latvia's quarterly GDP series from 1995Q1 till 2009Q3. The explanatory variables considered are i) an aggregate output in mining and quarrying, manufacturing, electricity, gas and water supply, and construction industries

(cp), ii) imports, iii) exports, iv) a ratio of exports over imports (nx), and v) money supply M1 (m). All series are quarterly, expressed in logs, and once regularly and once seasonally differenced, except m, that is not seasonally differenced. Table 3.2 lists the data and their sources.

Table 3.2.

The list of data used in the section

Series	Definition	Source
GDP	Gross domestic product	Central Statistical Bureau of Latvia
C	Output in mining and quarrying industry	Central Statistical Bureau of Latvia
D	Output in manufacturing industry	Central Statistical Bureau of Latvia
E	Output in electricity, gas and water supply industry	Central Statistical Bureau of Latvia
F	Output in construction industry	Central Statistical Bureau of Latvia
cp	Sum of C,D,E and F	Derived by the author
exp	Exports	Central Statistical Bureau of Latvia
imp	Imports	Central Statistical Bureau of Latvia
nx	Ratio of exports over imports, exp/imp	Derived by the author
m	Monetary aggregate M1, quarterly average	Bank of Latvia

Note: All national accounts series are chain-priced as of 2000.

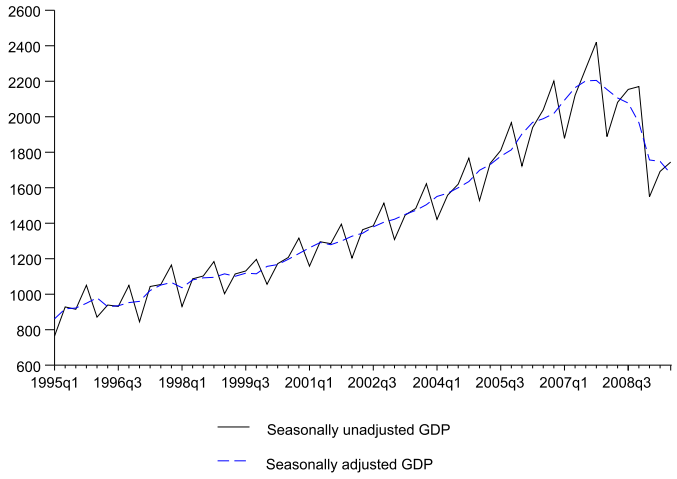

Fig. 3.4. Latvia's quarterly GDP series

Note: The first five observations get lost to make the seasonally unadjusted series stationary. If the rest part is divided in halves, the first half contains a smooth growth (a matter to calculate within-a-business-cycle-phase RMSFEs), whereas the second half contains a pronounced switch of business cycle phases from growth to a deep recession (a matter to calculate between-business-cycle-phases RMSFEs). Source: Central Statistical Bureau of Latvia.

We produce one-period ahead forecasts for GDP, given that all explanatory variables are known for the forecasting horizon (we call this exercise 'nowcasting'). All calculations

are performed in Scilab with the aid of its econometrics toolbox Grocer (see Dubois and Michaux (2010)).

Fig. 3.4 shows seasonally unadjusted as well as seasonally adjusted GDP series. The first five observations get lost to make the seasonally unadjusted series stationary. If the rest part is divided in halves, the first half contains smooth growth (a matter to calculate within-a-business-cycle-phase RMSFEs), whereas the second half contains a pronounced switch of business cycle phases from growth to a deep recession (a matter to calculate between-business-cycle-phases RMSFEs). Table 3.3 shows root mean squared forecast error (RMSFE) for the full sample, the first half of the sample ($\text{RMSFE}^{\text{within}}_{\text{phase}}$) and the second half of the sample ($\text{RMSFE}^{\text{between}}_{\text{phases}}$) from pseudo real-time nowcasts beginning at sample size 19 from a random walk (RW), autoregressive (AR) and vector autoregressive (VAR) models versus static, dynamic and mixed factor-augmented VAR (FAVAR) models, where factors are formed from variables *cp*, *nx* and *m*.

In this table, VAR models are specified by their endogenous variables (first parenthesis) and a lag order (second parenthesis). FAVAR models are specified by their endogenous variables (first parenthesis) and a lag order (second parenthesis).

Static factors are specified by a combination of three symbols 'fsi', where the first symbol 'f' denotes that the variable is a factor, the second symbol 's' means that the factor is obtained in a static manner, and the third symbol 'i' denotes the order of the factor. In this section, we will use only two kinds of static factors: 'fs1' and 'fs2', which are static first and second common factors, accordingly. Dynamic factors are specified by a symbol combination 'fdij(p,q)', where 'f' stands for being a factor, 'd' stands for being a dynamic one, 'ij' stands for being the i-th out of j simultaneously estimated factors, and the numbers '(p,q)' mean that the factor's dynamics in (3.16) are specified by an ARMA(p,q) process. Note that for simplicity, the indiosyncratic component in (3.16) is set to follow an AR(1) for all dynamic factors, regardless of their ARMA specifications. The least RMSFE for each sample space is framed.

Table 3.3 shows that it is better to use FAVAR with two static factors calculated from these three endogenous variables rather than VAR with the same three variables. It is also shown that the least nowcasting errors for a within-a-phase period are obtained by a parsimonious VAR model, whereas for the whole series and for a between-phases period - by a static FAVAR. Notably, none of the many dynamic and mixed factor FAVAR models specified by various ARMA dynamics is superior over the static FAVAR model. Table 3.4 shows the ranking of the models reported in Table 3.3. The ranking for a static two-factor FAVAR (model 9) for the whole series, within-a-phase, and between-phases period is {1,10,1} for variable set {*cp*,*nx*,*m*} out of overall 44 models. The ranking for a dynamic two-factor FAVAR(2,1) (model 33), for the whole series, within-a-phase, and between-phases period is {18,5,19} for variable set {*cp*,*nx*,*m*} out of overall 44 models. Unreported results also show that the ranking of the static FAVAR seems more stable with

Table 3.3.

A comparison of pseudo real-time nowcasting performance

N°	Model	RMSFE	$\text{RMSFE}^{\text{within}}_{\text{phase}}$	$\text{RMSFE}^{\text{between}}_{\text{phases}}$
1	RW	0.0318026	0.0258327	0.0370907
2	AR(1)	0.0289930	0.0174793	0.0375119
3	AR(2)	0.0290639	0.0176315	0.0375493
4	VAR(GDP,cp)(2)	0.0228362	0.0142717	0.0292916
5	VAR(GDP,cp,nx)(2)	0.0220654	0.0167891	0.0265320
6	VAR(GDP,cp,m)(2)	0.0220319	0.0162761	0.0268117
7	VAR(GDP,cp,nx,m)(2)	0.0226704	0.0181780	0.0266128
8	FAVAR(GDP,fs1)(2)	0.0287139	0.0226375	0.0339835
9	FAVAR(GDP,fs1,fs2)(2)	0.0210557	0.0172919	0.0244165
10	FAVAR(GDP,fd11(0,1))(2)	0.0311457	0.0213921	0.0388925
11	FAVAR(GDP,fd11(0,2))(2)	0.0305525	0.0219591	0.0375667
12	FAVAR(GDP,fd11(1,0))(2)	0.0311457	0.0213921	0.0388925
13	FAVAR(GDP,fd11(1,1))(2)	0.0314799	0.0208982	0.0397221
14	FAVAR(GDP,fd11(1,2))(2)	0.0311515	0.0226767	0.0381240
15	FAVAR(GDP,fd11(2,0))(2)	0.0305525	0.0219591	0.0375667
16	FAVAR(GDP,fd11(2,1))(2)	0.0303605	0.0219320	0.0372617
17	FAVAR(GDP,fd11(2,2))(2)	0.0310249	0.0220388	0.0383071
18	FAVAR(GDP,fd12(0,1))(2)	0.0309646	0.0207075	0.0389871
19	FAVAR(GDP,fd12(0,2))(2)	0.0307738	0.0222302	0.0377692
20	FAVAR(GDP,fd12(1,0))(2)	0.0309646	0.0207075	0.0389871
21	FAVAR(GDP,fd12(1,1))(2)	0.0285258	0.0210894	0.0347042
22	FAVAR(GDP,fd12(1,2))(2)	0.0281124	0.0227255	0.0328676
23	FAVAR(GDP,fd12(2,0))(2)	0.0307738	0.0222302	0.0377692
24	FAVAR(GDP,fd12(2,1))(2)	0.0299854	0.0217410	0.0367513
25	FAVAR(GDP,fd12(2,2))(2)	0.0297038	0.0236371	0.0349993
26	FAVAR(GDP,fd12(3,2))(2)	0.0285759	0.0220408	0.0341588
27	FAVAR(GDP,{fd12,fd22}(0,1))(2)	0.0284383	0.0208394	0.0347163
28	FAVAR(GDP,{fd12,fd22}(0,2))(2)	0.0319796	0.0219138	0.0399636
29	FAVAR(GDP,{fd12,fd22}(1,0))(2)	0.0284383	0.0208394	0.0347163
30	FAVAR(GDP,{fd12,fd22}(1,1))(2)	0.0285031	0.0206047	0.0349731
31	FAVAR(GDP,{fd12,fd22}(1,2))(2)	0.0218924	0.0174116	0.0258022
32	FAVAR(GDP,{fd12,fd22}(2,0))(2)	0.0319796	0.0219138	0.0399636
33	FAVAR(GDP,{fd12,fd22}(2,1))(2)	0.0257208	0.0162391	0.0329062
34	FAVAR(GDP,{fd12,fd22}(2,2))(2)	0.0220505	0.0184446	0.0253147
35	FAVAR(GDP,{fd12,fd22}(3,2))(2)	0.0230179	0.0183390	0.0271056
36	FAVAR(GDP,fd12(0,1),fs2)(2)	0.0231555	0.0170880	0.0281907
37	FAVAR(GDP,fd12(0,2),fs2)(2)	0.0226396	0.0161733	0.0278978
38	FAVAR(GDP,fd12(1,0),fs2)(2)	0.0231555	0.0170880	0.0281907
39	FAVAR(GDP,fd12(1,1),fs2)(2)	0.0250382	0.0198881	0.0295278
40	FAVAR(GDP,fd12(1,2),fs2)(2)	0.0214880	0.0176637	0.0249052
41	FAVAR(GDP,fd12(2,0),fs2)(2)	0.0226396	0.0161733	0.0278978
42	FAVAR(GDP,fd12(2,1),fs2)(2)	0.0220661	0.0156102	0.0272847
43	FAVAR(GDP,fd12(2,2),fs2)(2)	0.0217977	0.0179440	0.0252448
44	FAVAR(GDP,fd12(3,2),fs2)(2)	0.0216962	0.0174938	0.0253990

Note: Performance is compared between RW, AR, VAR, static, dynamic and mixed FAVAR models in terms of RMSFE for the full sample, first half of the sample ($\text{RMSFE}^{\text{within}}_{\text{phase}}$) and second half of the sample ($\text{RMSFE}^{\text{between}}_{\text{phases}}$). Factors are formed from *cp*, *nx* and *m*. The least RMSFE in each sample space is framed. Source: author's calculations.

Table 3.4.

Model ranking based on the RMSFEs reported in Table 3.3

N°	Model	Rank	$\text{Rank}^{\text{within}}_{\text{phase}}$	$\text{Rank}^{\text{between}}_{\text{phases}}$
1	RW	42	44	28
2	AR(1)	26	12	30
3	AR(2)	27	14	31
4	VAR(GDP,cp)(2)	13	1	31
5	VAR(GDP,cp,nx)(2)	8	7	7
6	VAR(GDP,cp,m)(2)	6	6	9
7	VAR(GDP,cp,nx,m)(2)	12	17	8
8	FAVAR(GDP,fs1)(2)	25	40	20
9	FAVAR(GDP,fs1,fs2)(2)	1	10	1
10	FAVAR(GDP,fd11(0,1))(2)	38	28	38
11	FAVAR(GDP,fd11(0,2))(2)	31	34	32
12	FAVAR(GDP,fd11(1,0))(2)	38	28	38
13	FAVAR(GDP,fd11(1,1))(2)	41	26	42
14	FAVAR(GDP,fd11(1,2))(2)	40	41	36
15	FAVAR(GDP,fd11(2,0))(2)	31	34	32
16	FAVAR(GDP,fd11(2,1))(2)	30	33	29
17	FAVAR(GDP,fd11(2,2))(2)	37	36	37
18	FAVAR(GDP,fd12(0,1))(2)	35	22	40
19	FAVAR(GDP,fd12(0,2))(2)	33	38	34
20	FAVAR(GDP,fd12(1,0))(2)	35	22	40
21	FAVAR(GDP,fd12(1,1))(2)	23	27	22
22	FAVAR(GDP,fd12(1,2))(2)	19	42	18
23	FAVAR(GDP,fd12(2,0))(2)	33	38	34
24	FAVAR(GDP,fd12(2,1))(2)	29	30	27
25	FAVAR(GDP,fd12(2,2))(2)	28	43	26
26	FAVAR(GDP,fd12(3,2))(2)	24	37	21
27	FAVAR(GDP,{fd12,fd22}(0,1))(2)	20	24	23
28	FAVAR(GDP,{fd12,fd22}(0,2))(2)	43	31	43
29	FAVAR(GDP,{fd12,fd22}(1,0))(2)	20	24	23
30	FAVAR(GDP,{fd12,fd22}(1,1))(2)	22	21	25
31	FAVAR(GDP,{fd12,fd22}(1,2))(2)	5	11	6
32	FAVAR(GDP,{fd12,fd22}(2,0))(2)	43	31	43
33	FAVAR(GDP,{fd12,fd22}(2,1))(2)	18	5	19
34	FAVAR(GDP,{fd12,fd22}(2,2))(2)	7	19	4
35	FAVAR(GDP,{fd12,fd22}(3,2))(2)	14	18	10
36	FAVAR(GDP,fd12(0,1),fs2)(2)	15	8	14
37	FAVAR(GDP,fd12(0,2),fs2)(2)	10	3	12
38	FAVAR(GDP,fd12(1,0),fs2)(2)	15	8	14
39	FAVAR(GDP,fd12(1,1),fs2)(2)	17	20	17
40	FAVAR(GDP,fd12(1,2),fs2)(2)	2	15	2
41	FAVAR(GDP,fd12(2,0),fs2)(2)	10	3	12
42	FAVAR(GDP,fd12(2,1),fs2)(2)	9	2	11
43	FAVAR(GDP,fd12(2,2),fs2)(2)	4	16	3
44	FAVAR(GDP,fd12(3,2),fs2)(2)	3	13	5

Note: The top rank in each sample space is framed.

respect to change of variables than that of the dynamic FAVAR. Indeed, the dynamic factor model turns from the best-performing nowcasting model for the set of variables {*cp*,*imp*,*m*}, to the worst nowcasting model for the set of variables {*cp*,*imp*,*nx*,*m*}, where the only difference between the variable sets is an addition of a single variable to the former set.

Table 3.5 shows a summary of mean and minimum RMSFE of static and dynamic factors between phases. Clearly, static factors have been found to be on average by 20 per cent more precise, in terms of RMSFE, than dynamic factors. Also, the table shows the mean RMSFE for the best performing static and dynamic factor over the five different data sets considered; again, the static factor shows more robustness in terms of forecasting precision than a particular best-performing dynamic factor specification.

Table 3.5.

A summary of factor model comparisons between phases

Factor	Mean RMSFE	Min RMSFE
static	100	84
dynamic	120	82
best static	86	-
best dynamic	109	-

Note: Numbers are normalized such that mean RMSFE of static factor is 100%.

It also has been found that, although some of the mixed FAVAR perform decently, static FAVAR model appears to be the most precise and robust with respect to the change of the factor-forming set of variables for the whole series as well as for the between-phases period. Also, if one considers one-factor models, it is found that one-factor static FAVAR outperforms one-factor dynamic FAVARs except for the within-a-phase period, where the performance is similar.

Plotting stationary GDP, static first common factor, and dynamic first common factor formed from the variable set {*cp*,*nx*,*m*}, where dynamic factors are generated by various ARMA specifications, starting from ARMA(0,1) and ending at ARMA(2,2) (to save space, Fig. 3.5 shows only the results for the ARMA(1,1) case), it has been found that the dynamic common factor, regardless of dynamics specification, hardly detects the recession period and never its depth. On the the other hand, the static first common factor is able to detect the recession and its depth and, thus, static factor methodology is considered to be more robust against rapid change in the dynamics of the target variable.

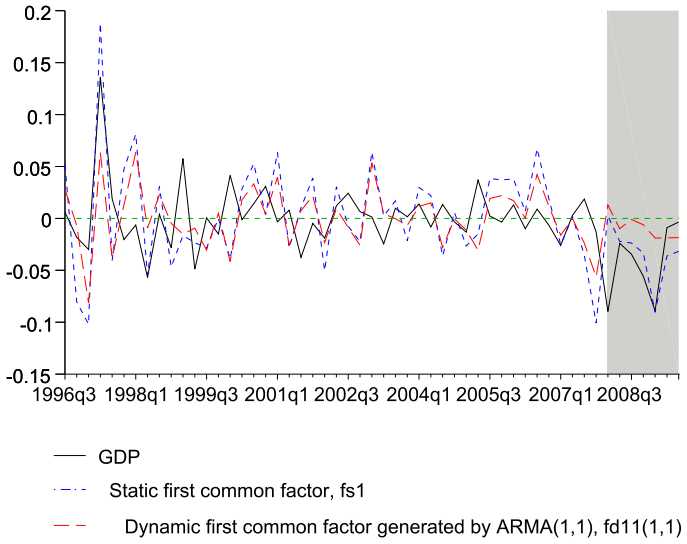

Fig. 3.5. Static versus dynamic factor following ARMA(1,1)

3.2.4 Conclusions

1. The choice between static and dynamic factors in nowcasting and forecasting literature is unresolved.

2. This section builds a small-scale factor model for Latvia's gross domestic product to study the exact dynamic versus static factor model performance during a pronounced switch of business cycle phases due to the latest recession.

3. It is shown that a static factor-augmented vector autoregressive model tends to improve upon the nowcasting performance of the vector autoregressive models during the switch business cycle phases, while exact dynamic factor models tend to fail to detect the timing and depth of the recession regardless of their autoregressive moving average specifications. As regards the period of smooth economic growth, static and dynamic factor models appear to show similar performance with potentially slight superiority of dynamic factor models.

4. The results imply that static-factor models are more robust and therefore should be preferred over the dynamic-factor models for economic forecasting.

3.3 Overview of the forecasting system

As a summary, Fig. 3.6 shows a chart of the main forecasting methods considered in this thesis. The signal-targeting methods in blue are developed, while the all-pass forecasting methods in green are studied for their robustness.

For one-dimensional data, if all frequencies are forecasted then ARIMA is a good benchmark. If business cycle frequencies are estimated then this thesis suggests the ABK filter. If, in addition to the dependent variable, there are some $n < 10$ additional explanatory variables, and if all frequencies are forecasted, then Bayesian methods perform decently, although, they are not robust against rapid changes in economic environment.

In practice, however, there are often many potential explanatory variables available. For example, Bank of Latvia collects more than 200 variables just for short-term forecasting of GDP alone. Therefore, methods capable of using high-dimensional data are demanded. One of the most successful methods dealing with high-dimensional datasets is factor methodology. This thesis suggests that using static factors from principal component analysis yields more robust performance than using exact dynamic factors estimated by Kalman filter.

If, however, business/trend-cycle frequencies are of interest, this thesis proposes using high-dimensional RMDFA which is shown to compete well with the generalized principal components method and have some clear advantages over the latter in forecasting, effect decomposition, and potential performance in presence of many redundant variables.

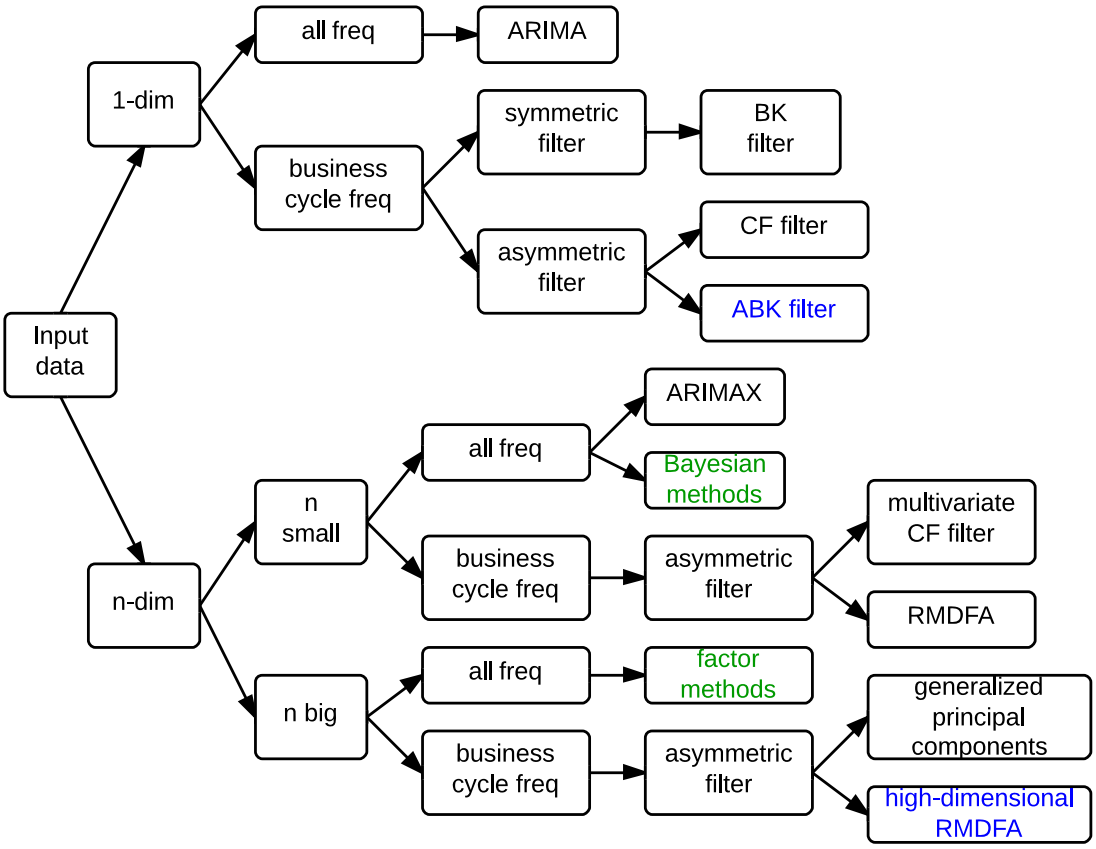

Fig. 3.6. A summary chart of the considered data and their forecasting methods in the thesis

Fig. 3.7 shows a short-term economic forecasting system developed at Bank of Latvia that uses the methods developed in this thesis.

The system involves technical stuff (system maintainer, model developer and filter developer), economists with their domain expert, and decision makers. The system maintainer is responsible for maintaining the forecasting system's technical part, which involves data collection from various in-house and world wide web sources, and storing that information in an orderly manner in the local database. The system maintainer is also responsible to run the software of allpass models (developed by a model developer) and business cycle or trendcycle filters (developed by filter developer), collect their results and produce local reports to economists on a regular basis (particularly, twice a month).

On the economists' side, there is a domain expert who produces his expert judgment on what happens and what will happen in the economy. There is a mutual information flow between the expert judgment and the results produced by the technical stuff - 1) expert judgment is influenced by the local reports produced by the technical stuff, and 2) the expert judgment enters as one of the 'models' in the allpass models block in order to improve i) communication, and possibly also ii) forecasting performance.

The domain expert, together with his fellow economists, partly based on the local reports produced by the technical stuff, then produce a global report to the decision makers.

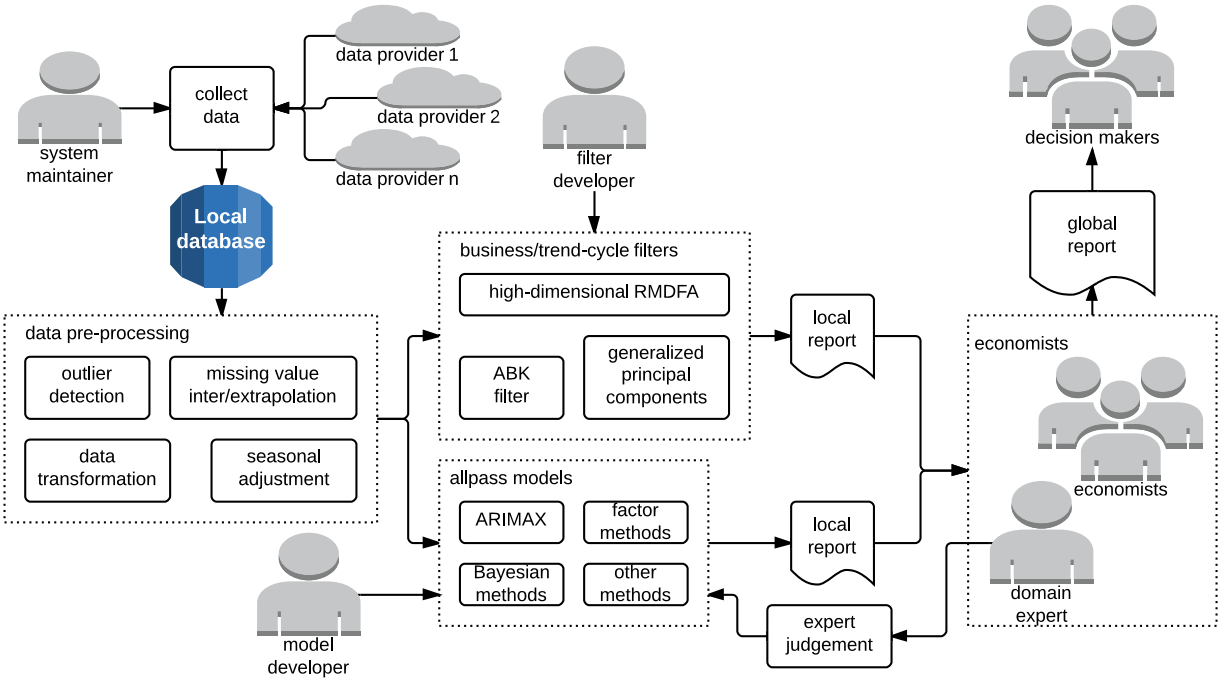

Fig. 3.7. A flowchart of the short-term economic forecasting system at Bank of Latvia using the methods considered in the thesis

MAIN CONCLUSIONS

The main objective of the thesis is to develop robust forecasting methods that are able to work with noisy and high-dimensional macroeconomic data. In order to fulfil the objective of the thesis, the following tasks have been proposed: 1) develop a univariate asymmetric babdpass filter for end-point estimation problems, 2) compare the performance of the developed asymmetric filter to the currently most popular alternative in macroeconomics, 3) develop a method suitable for forecasting and signal extraction using high-dimensional data subject to revisions, 4) assess the properties of the above method and compare with the currently best alternative in macroeconomics and 5) investigate the robustness issues for Bayesian Minnessota prior and factor forecasting models in macroeconomics.

The main objective of the thesis has been achieved; the proposed tasks have been accomplished:

1. An asymmetric filter has been developed for frequency band extraction at the end-points of univariate series.

2. The developed filter's performance has been compared to the currently most popular alternative in macroeconomics - the Christiano-Fitzgerald filter - in monte carlo simulations.

3. A method has been developed for signal extraction and forecasting using high-dimensional and noisy datasets.

4. The properties of the developed high-dimensional filter have been assessed and compared to the factor methodology.

5. Robustness issues of Bayesian and factor methods have been studied when the dynamics of the target variable is subject to a rapid change.

6. The forecasting system module of Bank of Latvia has been developed.

Main conclusions:

1. The developed asymmetric band pass filter outperforms the Christiano-Fitzgerald filter within two years from the end-point.

2. The developed high-dimensional filter is better suited for signal forecasting, effect decomposition and for dealing with irrelevant explanatory variables than the factor methodology.

3. The Bayesian Minnessota prior is not robust against rapid change in the dynamics of the target variable, thus making the forecasts imprecise if the prior is unchanged.

4. Static factor models are more robust than the dynamic factor models against unexpected change in dynamics of the data, and thus are to be preferred.

Main theses of defense:

1. The developed high-dimensional filtration algorithm allows for signal extraction and forecasting using high-dimensional and noisy datasets.

2. The developed univariate asymmetric filter is more precise than the Christiano-Fitzgerald filter for business-cycle frequency extraction at the end-points of univariate macroeconomic series.

3. The traditional methods for forecasting with many time series - Bayesian Minnessota prior and exact dynamic factors - are subject to robustness issues when the dynamics of the target are subject to rapid change.

The approbation of the thesis has been achieved by presenting the results at 11 international scientific conferences and seminars, by publishing 11 articles in international scientific journals or conference proceedings, by implementing the methods at the Central statistical bureau of Latvia for producing the flash release of Latvia's GDP since year 2009. The developed methods are used for forecasting purposes at Bank of Latvia since 2011.

GALVENIE SECINĀJUMI

Promocijas darba galvenais mērķis ir izstrādāt robustas prognozēšanas metodes, kas ir piemērotas strādāt ar trokšņainiem un liela apjoma datiem ar pielietojumu makroekonomikā. Lai sasniegtu promocijas darba galveno mērķi, tika izvirzīti šādi uzdevumi: 1) izstrādāt asimetrisku joslas filtru gala punkta novērtēšanas problēmām viendimensionāliem datiem, 2) salīdzināt izstrādātā asimetriskā filtra sniegumu ar pašreiz populārāko alternatīvu makroekonomikā, 3) izstrādāt prognozēšanas un signāla novērtēšanas metodi, kas būtu piemērota daudzdimensionāliem un trokšņainiem datiem, 4) novērtēt izstrādātās metodes īpašības un salīdzināt ar pašreiz labāko alternatīvu makroekonomikā, 5) novērtēt robustuma problēmas Beijesa un faktoru prognozēšanas modeļiem.

Promocijas darba mērķis ir sasniegts, tika izpildīti izvirzītie promocijas darba uzdevumi:

1. Izstrādāts joslas filtrs gala punkta novērtēšanas problēmām viendimensionāliem datiem.

2. Veiktas monte karlo simulācijas, lai salīdzinātu izstrādātā filtra sniegumu ar pašreiz populārāko joslas filtru makroekonomikā - Kristiano-Ficdžeralda filtru.

3. Izstrādāta metode signāla novērtēšanai un prognozēšanai ar daudzdimensionāliem un trokšņainiem datiem.

4. Novērtētas daudzdimensionālā filtra īpašības un salīdzinātas ar faktoru metodoloģiju.

5. Novērtētas Beijesa un faktoru modeļu robustuma problēmas, kad mērķa rādītāja dinamikā notiek straujas izmaiņas.

6. Attīstīts Latvijas Bankas prognozēšanas sistēmas modulis.

No rezultātiem izrietošie galvenie secinājumi:

1. Izstrādātais viendimensionālais filtrs pārspēj Kristiano-Ficdžeralda filtru divu gadu posmā no galapunkta.

2. Izstrādātā daudzdimensionālā filtrēšanas metode ir piemērotāka signāla prognozēšanai, efektu dekompozīcijai un strādāšanai ar liekiem izskaidrojošiem rādītājiem nekā faktoru metodoloģija.

3. Straujas dinamikas izmaiņas mērķa rādītājā būtiski maina optimālo Beijesa prioru, padarot prognozes neprecīzas, ja priors netiek mainīts.

4. Statisko faktoru modeļi ir robustāki par dinamisko faktoru modeļiem pret negaidītām izmaiņām datu dinamikā.

Aizstāvēšanai izvirzāmās tēzes:

1. Izstrādātā daudzdimensionālā filtrēšanas metode ļauj novērtēt un prognozēt signālu, izmantojot daudzdimensionālus un trokšņainus datus.

2. Izstrādātais viendimensionālais filtrs ir precīzāks par Kristiano-Ficdžeralda filtru biznesa cikla frekvenču novērtēšanai viendimensionālu makroekonomisko laikrindu galapunktos.

3. Tradicionālās laikrindu prognozēšanas metodes ar daudziem rādītājiem - Beijesa Minesotas priors un precīzie dinamiskie faktori - ir pakļautas robustuma problēmām, kad atkarīgā rādītāja dinamika strauji mainās.

Darba aprobācija tika veikta, prezentējot darba rezultātus 11 starptautiskajās zinātniskajās konferencēs un semināros, publicējot 11 zinātniskos rakstus starptautiskajos zinātniskajos izdevumos, pielietojot metodes LR Centrālajā statistikas pārvaldē Latvijas IKP ātrā novērtējuma ražošanā kopš 2009. gada. Izstrādātās metodes tiek lietotas prognozēšanā Latvijas Bankā kopš 2011. gada.

BIBLIOGRAPHY

[1] Ajevskis, V. and Davidsons, G. (2008), "Dynamic Factor Models in Forecasting Latvia's Gross Domestic Product", Working Paper, 2/2008, Latvijas Banka.

[2] Altissimo, F., R. Cristadoro, M. Forni, M. Lippi, G. Veronese (2010), "New Eurocoin: Tracking Economic Growth in Real Time", *The Review of Economics and Statistics*, MIT Press, vol. 92(4), 1024-1034.

[3] Altissimo, F., A. Bassanetti, R. Cristadoro, M. Forni, M. Hallin, M. Lippi, L. Reichlin and G. Veronese (2001), "A real time coincident indicator for the euro area business cycle", CEPR Discussion Paper No. 3108.

[4] Amisano, G. and J. Geweke (2013), "Prediction using several macroeconomic models", ECB working paper series No 1537, European Central bank.

[5] Asai, M., M. Caporin and M. McAleer (2012), "Forecasting Value-at-Risk Using Block Structure Multivariate Stochastic Volatility Models," KIER Working Papers 812, Kyoto University, Institute of Economic Research.

[6] Athanasopoulos, G., R.J. Hyndman, H. Song, and D.C. Wu (2011), "The tourism forecasting competition", *International Journal of Forecasting*, Elsevier, vol. 27(3), 822-844.

[7] Banerjee, A., M. Marcellino and I. Masten (2010), "Forecasting with factor-augmented error correction models", Discussion Papers 09-06R, Department of Economics, University of Birmingham.

[8] Baxter, M. and R. G. King (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.

[9] 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.

[10] Boivin, J. and S. Ng (2005), "Understanding and comparing factor-based forecasts", *International Journal of Central Banking*, 1(3).

[11] Boivin, J. and S. Ng (2006), "Are more data always better for factor analysis?", *Journal of Econometrics* 132, 169-194.

[12] Bollerslev, T. (1986), "Generalized autoregressive conditional heteroskedasticity", *Journal of Econometrics*, vol. 31(3), 307-327.

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

[14] Box, G.E.P. and Jenkins, G.M. (1970), *Time Series Analysis: Forecasting and Control*, San Francisco: Holden-Day.

[15] Box, G., G. M. Jenkins, and G. Reinsel (1994), *Time Series Analysis: Forecasting and Control*, 3rd Edition, Prentice Hall.

[16] Brockwell, P. and Davis, R. (1987), *Time Series: Theory and Methods*, New York: Springer-Verlag.

[17] Burns, A. F. and W. C. Mitchell (1946), *Measuring business cycles*, New York: NBER.

[18] Buss, G. (2009), "Comparing Forecasts of Latvia's GDP Using Simple Seasonal ARIMA Models and Direct Versus Indirect Approach", MPRA Paper 16832, University Library of Munich, Germany. (Indexed in: RePEc, SciVerse, Scirus, Econlit)

[19] Buss, G. (2010), "A Note on Now-/Forecasting with Dynamic Versus Static Factor Models along a Business Cycle", 10th International Vilnius Conference on Probability Theory and Mathematical Statistics: Abstracts of Communications, Vilnius, Lithuania, 28 June - 2 July, 2010, p 119.

[20] Buss, G. (2010), "Asymmetric Baxter-King Filter", Scientific Journal of RTU, 5th series, Computer Science, 42. vol, pp 95-99. (Indexed in: EBSCO, RePEc, SciVerse, Scirus, Econlit)

[21] Buss, G. (2010), "Comparing forecasts of Latvia's GDP using simple seasonal ARIMA models and direct versus indirect approach: an overview", The results of statistical scientific research 2010, Research papers, Ed. O. Krastins, I. Vanags, Riga: Central Statistical Bureau of Latvia, pp 50-56.

[22] Buss, G. (2010), "Economic Forecasts with Bayesian Autoregressive Distributed Lag Model: Choosing Optimal Prior in Economic Downturn", Aplimat: Journal of Applied Mathematics, vol 3, pp. 191-200. (Indexed in: RePEc, SciVerse, Scirus, Econlit)

[23] Buss, G. (2010), "Forecasts with Single-Equation Markov-Switching Model: an Application to the Gross Domestic Product of Latvia", Journal of Applied Economic Sciences, Vol 5, Issue 2, pp 49-59. (Indexed in: Scopus, RePEc, SciVerse, Scirus, Econlit)

[24] Buss, G. (2010), "Forecasts with Single-Equation Markov-Switching Model: an Application to the Gross Domestic Product of Latvia", Acta Societatis Mathematicae Latviensis: Abstract of the 8th Latvian Mathematical Conference, Valmiera, Latvia, 9-10 April, 2010, p 17.

[25] Buss, G. (2011), "An Application of Direct Filter Approach: New Economic Indicators for Latvia", Scientific Journal of RTU, 5th series, Computer Science, 48 vol, pp 75-81. (Indexed in: EBSCO)

[26] Buss, G. (2011), "A Band Pass Filter for Real-Time Signal Extraction", Abstracts of 16th International Conference on Mathematical Modelling and Analysis, Sigulda, Latvia, 25-28 May, 2011, p 22.

[27] Buss, G. (2011), "Preliminary Results on Asymmetric Baxter-King Filter", Aplimat 2011: 10th International conference on applied mathematics: Proceedings, Bratislava, Slovakia, 1-4 February, 2011, pp 1499-1508.

[28] Buss G. (2012). "A new real-time indicator for the Euro Area GDP", Working Papers 2012/02, Latvijas Banka. (Indexed in: RePEc, SciVerse, Scirus, Econlit, Google Scholar, Microsoft Academic Search)

[29] Buss, G. (2012), "Introduction to regularized DFA", Scientific Journal of RTU, series 5, vol. 48, pp 48-56. (Indexed in: EBSCO)

[30] Caggiano, G., Kapetanios, G. and Labhard, V. (2009), "Are more data always better for factor analysis? Results for the Euro Area, the six largest Euro Area countries and the UK", Working Paper Series No 1051/May 2009, European Central Bank.

[31] Camacho, M., G. Perez-Quiros and P. Poncela (2012), "Markov-switching dynamic factor models in real time", CEPR Discussion Papers 8866, C.E.P.R. Discussion Papers.

[32] Campbell, J. Y. and P. Perron (1991), "Pitfalls and opportunities: what macroeconomists should know about unit roots", NBER Technical Working Papers 0100, National Bureau of Economic Research, Inc.

[33] Caporello, G., A. Maravall and F. J. Sanchez (2001), "Program TSW reference manual", Banco de Espana Working Papers 0112, Banco de Espana.

[34] Chauvet, M. (1998), "An econometric characterization of business cycle dynamics with factor structure and regime switches", *International Economic Review*, 39(4), 969-96.

[35] Chauvet, M. and J. Hamilton (2005), "Dating business cycle turning points", NBER Working Papers 11422.

[36] Christiano, L. J. and T. J. Fitzgerald (2003), "The band pass filter", *International Economic Review*, vol. 44(2), 435-465.

[37] D'Agostino, A. and D. Giannone (2007), "Comparing alternative predictors based on large-panel factor models", CEPR Discussion Papers 6564, C.E.P.R. Discussion Papers.

[38] Doan, T., R. B. Litterman and C. A. Sims (1984), "Forecasting and Conditional Projection Using Realistic Prior Distributions", *Econometric Reviews*, 3(1), 1-100.

[39] Doz, C. and F. Lenglart (1999), "Analyse factorielle dynamique: test du nombre de facteurs, estimation et application a l'enquete de conjoncture dans l'industrie", *Annales d'Economie et de Statistique*, No. 54, 91-127.

[40] Dubois, E. and Michaux, E. (2010), "Grocer 1.41: an econometric toolbox for Scilab", available at http://dubois.ensae.net/grocer.html.

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

[42] Forni, M., M. Hallin, M. Lippi, and L. Reichlin (2000), "The generalized dynamic factor model: identification and estimation", *The Review of Economics and Statistics*, vol. 82(4), 540-554.

[43] Forni, M., Hallin, M., Lippi, M. and Reichlin, L. (2005), "The Generalized Dynamic Factor Model, One-sided Estimation and Forecasting", *Journal of the American Statistical Association*, vol. 100, 830-840.

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

[45] Guay, A. and P. St-Amant (2005), "Do the Hodrick-Prescott and Baxter-King filters provide a good approximation of business cycles?", *Annales d'Economie et de Statistique*, issue 77.

[46] Gupta, R. and A. Kabundi (2008), "Forecasting macroeconomic variables using large datasets: dynamic factor models versus large-scale BVARs", Working Papers 200816, University of Pretoria, Department of Economics.

[47] Gupta, R. and A. Kabundi (2009), "A large factor model for forecasting macroeconomic variables in South Africa", Working Papers 137, Economic Research Southern Africa, University of Cape Town.

[48] Hamilton, J.D. (1985), "Uncovering Financial Market Expectations of Inflation", *Journal of Political Economy* 93, 1224-41.

[49] Hamilton, J. D. (1989), "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle", *Econometrica*, vol. 57(2), 357-384.

[50] Hamilton, J.D. (1994): *Time Series Analysis*, Princeton, New Yersey: Princeton University Press.

[51] Hodges, J. S. and D. J. Sargent (2001), "Counting degrees of freedom in hierarchical and other richly-parameterized models", *Biometrika*, vol. 88(2), 367-379.

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

[53] Hoerl, A. E. and R. W. Kennard (1970), "Ridge regression: applications to nonorthogonal problems", *Technometrics*, vol. 12(1), 69-82.

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

[55] Kalman, R.E. (1960), "A New Approach to Linear Filtering and Prediction problems", *Journal of Basic Engineering, Transactions of the ASME* Series D, 82, 35-45.

[56] Kim, M. J. and J. S. Yoo (1995),"New index of coincident indicators: A multivariate Markov switching factor model approach", *Journal of Monetary Economics*, 36(3), 607-630.

[57] Kim, C. J. and C. R. Nelson (1998), "Business cycle turning points, a new coincident index, and tests of duration dependence based on a dynamic factor model with regime switching", *Review of Economics and Statistics*, 80(2), 188-201.

[58] King, R. G. and S. T. Rebelo (1993), "Low frequency filtering and real business cycles", *Journal of Economic Dynamics and Control*, vol. 17(1-2), 207-231.

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

[60] LeSage, J. P. (1999), *Applied Econometrics using MATLAB*.

[61] LeSage, J. P., A. Krivelyova (1999), "A Spatial Prior for Bayesian Vector Autoregressive Models", *Journal of Regional Science*, 39, 297-317.

[62] LeSage, J. P., M. Magura (1991), "Using Interindustry Input-Output Relations as a Bayesian prior in Employment Forecasting Models", *International Journal of Forecasting*, 7, 231-238.

[63] LeSage, J. P., Z. Pan (1995), "Using Spatial Contiguity as Bayesian Prior Information in Regional Forecasting Models", *International Regional Science Review*, 18, 33-53.

[64] Litterman, R. B. (1979), "Techniques of Forecasting Using Vector Autoregressions", Working Paper 115, Federal Reserve Bank of Minneapolis.

[65] Litterman, R. B. (1986), "Forecasting with Bayesian Vector Autoregressions - Five Years of Experience", *Journal of Business & Economic Statistics*, 4, 25-38.

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

[67] Marcellino, M. and C. Schumacher (2008), "Factor-MIDAS for now- and forecasting with ragged-edge data: A model comparison for German GDP", CEPR Discussion Papers 6708.

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

[69] Moody, J. E. (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.

[70] Priestley, M. B. (1981), *Spectral Analysis and Time Series*, Academic Press, Inc.

[71] Raviv, E., K.E. Bouwman and D. van Dijk (2013), "Forecasting Day-Ahead Electricity Prices: Utilizing Hourly Prices," Tinbergen Institute Discussion Papers 13-068/III, Tinbergen Institute.

[72] Reijer, A. H. J. den (2005), "Forecasting Dutch GDP using large scale factor models", DNB Working Papers 028, Netherlands Central Bank, Research Department.

[73] Schneider, M. and M. Spitzer (2004), "Forecasting Austrian GDP using the generalized dynamic factor model", Working Papers 89, Oesterreichische Nationalbank (Austrian Central Bank).

[74] Schumacher, C. (2005), "Forecasting German GDP using alternative factor models based on large datasets", Discussion Paper Series 1: Economic Studies 2005, 24, Deutsche Bundesbank, Research Centre.

[75] Schumacher, C. (2009), "Factor forecasting using international targeted predictors: the case of German GDP", Discussion Paper, Series 1: Economic Studies, No 10/2009, Deutsche Bundesbank.

[76] Siliverstovs, B., and K. A. Kholodilin (2010), "Assessing the real-time informational content of macroeconomic data releases for now-/forecasting GDP: evidence for Switzerland", KOF Working papers 10-251, KOF Swiss Economic Institute, ETH Zurich.

[77] Stock, J.H., and Watson, M.W. (1991), "A Probability Model of the Coincident Economic Indicators", in Lahiri, K. and Moore, G.H. (eds.) *Leading Economic Indicators: New Approaches and Forecasting Records*, Cambridge, England: Cambridge University Press.

[78] Stock, J.H., and Watson, M.W. (1998), "Diffusion Indexes", NBER Working Paper 6702, National Bureau of Economic Research, Inc.

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

[80] Stock, J.H., and Watson, M.W. (2003), "Combination Forecasts of Output Growth", Working Paper.

[81] Stock, J.H., and Watson, M.W. (2004), "Forecasting with many predictors", Prepared for the *Handbook of Economic Forecasting*.

[82] Theil, H. and A. S. Goldberger (1961), "On Pure and Mixed Statistical Estimation in Economics", *International Economic Review*, 2, 65-78.

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

[84] Tikhonov, A. N. and V. Y. Arsenin (1977), *Solutions of Ill-Posed Problems*, Washington: V. H. Winston & Sons.

[85] 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.

[86] Watson, M. W. (1986), "Univariate detrending methods with stochastic trends", *Journal of Monetary Economy*, Elsevier, vol. 18, 49-75.

[87] Wiener, N. (1949), *Extrapolation, interpolation, and smoothing of stationary time series*, MIT Press, ISBN 0-262-73005-7.

Empirical Study Based on TRAMO/SEATS and Census X-12-ARIMA", KOF-Working Paper 96, ETH-Zurich.

of Turning-Points", Lecture Notes in Economic and Mathematical Systems, 547, Springer-Verlag Berlin Heidelberg.

[88] 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.

[89] Wildi, M. (2009), "Real-Time US-Recession Indicator (USRI): A Classical Cycle Perspective with Bounceback", IDP-Working Paper, IDP-WP-09Jun-06.

[90] Wildi, M. (2010), "Real-Time Signal extraction: a shift of perspective", *Estudios de Economia Aplicada*, Vol 28(3), 497-518.

[91] 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.

[92] Wildi, M. (2012), "Elements of forecasting and signal extraction", IDP-Working Paper, available at http://blog.zhaw.ch/idp/sefblog/uploads/elements_1904.pdf.