Forecasting with Large Unbalanced Datasets: The Mixed-Frequency Three-Pass Regression Filter

Christian Hepenstrick and Massimiliano Marcellino

Legal Issues

DISCLAIMER

The views expressed in this paper are those of the author(s) and do not necessarily represent those of the Swiss National Bank. Working Papers describe research in progress. Their aim is to elicit comments and to further debate.

COPYRIGHT©

The Swiss National Bank (SNB) respects all third-party rights, in particular rights relating to works protected by copyright (information or data, wordings and depictions, to the extent that these are of an individual character).

SNB publications containing a reference to a copyright (© Swiss National Bank/SNB, Zurich/year, or similar) may, under copyright law, only be used (reproduced, used via the internet, etc.) for non-commercial purposes and provided that the source is mentioned. Their use for commercial purposes is only permitted with the prior express consent of the SNB.

General information and data published without reference to a copyright may be used without mentioning the source. To the extent that the information and data clearly derive from outside sources, the users of such information and data are obliged to respect any existing copyrights and to obtain the right of use from the relevant outside source themselves.

LIMITATION OF LIABILITY

The SNB accepts no responsibility for any information it provides. Under no circumstances will it accept any liability for losses or damage which may result from the use of such information. This limitation of liability applies, in particular, to the topicality, accu‑racy, validity and availability of the information.

Forecasting with Large Unbalanced Datasets:
The Mixed-Frequency Three-Pass Regression Filter[*]

Christian Hepenstrick[†]   Massimiliano Marcellino[‡]

January 2016

Abstract

In this paper, we propose a modification of the three-pass regression filter (3PRF) to make it applicable to large mixed-frequency datasets with ragged edges in a forecasting context. The resulting method, labeled MF-3PRF, is very simple but compares well to alternative mixed frequency factor estimation procedures in terms of theoretical properties, finite sample performance in Monte Carlo experiments, and empirical applications to GDP growth nowcasting and forecasting for the USA and a variety of other countries.

Keywords: Dynamic Factor Models, Mixed Frequency, GDP Nowcasting, Forecasting, Partial Least Squares

JEL-Codes: E37, C32, C53

[†]Swiss National Bank
[‡]Bocconi University, IGIER and CEPR

1 Introduction

The recent crisis has emphasized the importance for policy-makers and economic agents of a real-time assessment of the current state of the economy and its expected developments when a large but incomplete information set is available. The main obstacle is the delay with which key macroeconomic indicators, such as GDP and its components, but also fiscal variables, regional/sectoral indicators and disaggregated data, are released. For example, even in the USA, where the economic data production process is quite efficient, GDP data are only available on a quarterly basis and the advance estimate is only published with an approximately 4 week delay. Moreover, preliminary data are often subsequently revised, particularly around turning points in the business cycle. Furthermore, a large and growing number of more timely leading and coincident indicators is available, at a monthly or even higher frequency, based in particular on financial, survey and internet data, albeit occasionally subject to short samples, missing observations and other data irregularities. This has stimulated a vast amount of econometric research on how to exploit this timely and higher frequency, but irregular, information to generate estimates for the key low-frequency economic indicators. Banbura, Giannone, Reichlin (2011), Banbura, Giannone, Modugno, Reichlin (2012) and Foroni and Marcellino (2013, 2014) provide overviews of the alternative econometric techniques, most of which rely on the Kalman filter.

In this paper, we are particularly interested in the use of a large set of monthly indicators for backcasting, nowcasting and short-term forecasting a quarterly target variable, GDP growth in our empirical applications. Various approaches are available, based on variable selection methods (e.g., Bulligan, Marcellino, Venditti (2015)), pooling of single-indicator models (e.g., Kuzin, Marcellino, Schumacher (2013)), Bayesian techniques (e.g., Carriero, Clark and Marcellino (2015)) and mixed-frequency factor models (e.g., Giannone, Reichlin and Small (2008)). We focus on factor models, as they performed reasonably well in several empirical analyses and are perhaps the most common when using a large, irregular, mixed-frequency dataset.

Recent applications using mixed-frequency factor models are, for example, Banbura and Runstler (2011), who further extend the model of Giannone, Reichlin and Small (2008), Banbura and Modugno (2014), who discuss the maximum likelihood estimation of factor models on datasets with arbitrary pattern of missing data, and Marcellino, Porqueddu and Venditti (2015), who allow for stochastic volatility in a Bayesian estimation framework.[1] These papers are based on the Kalman filter, which is most efficient but requires the specification of a parametric model in high frequency and can be computationally very demanding for large datasets and high frequency mismatch. Therefore, Marcellino and Schumacher (2010) propose a simpler alternative, which combines the Stock and Watson (2002) EM algorithm for factor estimation from irregular datasets with the Mixed-Data Sampling

---

[1]Barhoumi et al. (2008) compare small bridge equations and forecast equations in which the bridging between monthly and quarterly data is achieved through a regression on factors extracted from large monthly datasets. Angelini et al. (2011) provide an out-of-sample evaluation of the method presented by Banbura and Runstler (2011) and compare the forecasting performance of this approach to that obtained by pooling the forecasts from selected bridge equations.

(MIDAS) regression technique (Ghysels, Santa-Clara and Valkanov (2004)) or its unrestricted counterpart (U-MIDAS, Foroni, Marcellino and Schumacher (2015)). In this paper, we propose an even simpler approach, which is based primarily on OLS regressions.

Our starting point is the three-pass regression filter (3PRF), developed by Kelly and Pruitt (2013, 2014) Henceforth, we refer to Kelly and Pruitt (2014) as KP. The 3PRF, which is an extension of Partial Least Squares (PLS), permits to obtain targeted factors for forecasting a specific variable of interest in a simple and intuitive manner. Moreover, it has a number of (asymptotic) optimality properties, performs well in finite samples compared to more complex alternatives, and produces good nowcasts and short-term forecasts for a variety of macroeconomic and financial variables, see KP.

As a first step, in Section 2, after an overview of the 3PRF, we discuss its extension to handle mixed-frequency datasets, possibly with ragged edges created by publication delays that cause missing values for some of the variables at the end of the sample (Wallis (1986)). The resulting method, called mixed-frequency three pass regression filter (MF-3PRF), inherits the asymptotic optimality properties of 3PRF.

In Section 3, we assess the finite-sample backcasting, nowcasting and short-term forecasting performance of the MF-3PRF by means of a set of Monte Carlo experiments. We also compare its performance to the mixed-frequency U-MIDAS factor model of Marcellino and Schumacher (2010), Factor-U-MIDAS, and 3PRF using data aggregated to the quarterly frequency.

In Section 4, we apply the MF-3PRF and Factor-U-MIDAS for backcasting, nowcasting and short-term forecasting quarterly GDP growth in the USA, based on a large set of approximately 100 monthly indicators. We also introduce and assess a number of extensions of the MF-3PRF, finding that, overall, the basic version works reasonably well.

Both the evaluation using simulated data in Section 3 and the empirical application in Section 4 are supportive of the new MF-3PRF. To further assess its performance, in Section 5, we repeat the backcasting, nowcasting and short-term forecasting of quarterly GDP growth for a large number of other countries, in each case based on either a large monthly national dataset or a global dataset that pools together all of the national datasets and includes over 800 variables. Specifically, we consider first the euro area as a whole and its four largest member states: France, Germany, Italy and Spain. Next, we consider two other large developed economies, the UK and Japan. Finally, two large emerging countries, Brazil and India, and two smaller ones, Korea and Taiwan, are considered. For the majority of countries, the MF-3PRF approach again outperforms both the standard benchmarks and the competing factor approaches. This finding, combined with its simplicity and substantial computational speed, makes the MF-3PRF a very useful tool to handle large mixed-frequency and irregular datasets.

2 The Mixed-Frequency 3PRF

In this section we first review the 3PRF, closely following the notation of KP, to whom we refer for additional details. Next, we extend the 3PRF to mixed-frequency data, possibly with ragged edge.

2.1 The 3PRF

Let us consider the following model:

$$y_{t+1} = \beta_0 + \beta' \mathbf{F}_t + \eta_{t+1}, \tag{1}$$

$$\mathbf{z}_t = \lambda_0 + \boldsymbol{\Lambda} \mathbf{F}_t + \omega_t, \tag{2}$$

$$\mathbf{x}_t = \phi_0 + \boldsymbol{\Phi} \mathbf{F}_t + \varepsilon_t, \tag{3}$$

where y is the target variable of interest; $\mathbf{F}_t = (\mathbf{f}_t', \mathbf{g}_t')'$ are the $K = K_f + K_g$ common driving forces of all variables, the unobservable factors; $\beta = (\beta_f', \mathbf{0}')'$, such that y only depends on $\mathbf{f}$; $\mathbf{z}$ is a small set of L proxies that are driven by the same underlying forces as y, such that $\boldsymbol{\Lambda} = (\boldsymbol{\Lambda}_f, \mathbf{0})$ with $\boldsymbol{\Lambda}_f$ nonsingular; $\mathbf{x}_t$ is a large set of N weakly stationary variables, driven by both $\mathbf{f}$ and $\mathbf{g}$; and $t = 1, ..., T$. To achieve identification, when N and T diverge, the covariance of the loadings is assumed to be the identity matrix, and the factors are orthogonal to one another.[2] For the sake of space, we refer to KP for precise conditions on the factors, loadings, allowed temporal and cross-sectional dependence of the residuals, and existence of proper central limit theorems.

With respect to the factor model analyzed by, e.g., Stock and Watson (2002), here the large dataset $\mathbf{x}_t$ is possibly driven by more factors than the target variable y. Asymptotically and with a strong factor structure, this does not matter for forecasting, as if we include more factors than those strictly needed in (1), then their estimated loadings will converge to zero. However, in finite samples, or if the $\mathbf{f}_t$ are weak while $\mathbf{g}_t$ are strong factors, estimating and using only the required factors $\mathbf{f}_t$ in (1) would be very convenient. This is a well known problem, see, e.g., Boivin and Ng (2006), who suggest some form of variable pre-selection prior to factor extraction.

KP provide a general, elegant and simple solution to the problem of estimating in the model (1)-(3) $\mathbf{f}_t$ only. Their approach can be represented by the three following steps (that give the name to the procedure):

- Pass 1: run a (time-series) regression of each element of $\mathbf{x}$, $x_{\mathbf{i}}$, on $\mathbf{z}$:

$$x_{\mathbf{i},t} = \phi_{0,i} + \mathbf{z}_t' \phi_i + \epsilon_{i,t},$$

where $i = 1, ...N$, and retain the OLS estimates $\widehat{\phi}_i$.

---

[2]Specifically, defining $\mathbf{J}_T = \mathbf{I}_T - \frac{1}{T}\iota_T\iota_T'$ where $\mathbf{I}_T$ is a T-dimensional identity matrix and ι_T a T-vector of ones (and similarly $\mathbf{J}_N$) and assuming that $N^{-1}\boldsymbol{\Phi}'\mathbf{J}_N\boldsymbol{\Phi} \underset{N \longrightarrow \infty}{\longrightarrow} \mathbf{P}$, $N^{-1}\boldsymbol{\Phi}'\mathbf{J}_N\phi_0 \underset{N \longrightarrow \infty}{\longrightarrow} \mathbf{P}_1$, $T^{-1}\mathbf{F}'\mathbf{J}_T\mathbf{F} \underset{T \longrightarrow \infty}{\longrightarrow} \Delta_F$, for identification, we require, as KP, that $\mathbf{P} = \mathbf{I}$, $\mathbf{P}_1 = \mathbf{0}$, and that Δ_F is diagonal, positive definite, and that each diagonal element is unique.

- Pass 2: run a (cross-sectional) regression of $\mathbf{x}_t$ on $\widehat{\phi}_i$:

$$x_{\mathbf{i},t} = \phi_{0,t} + \widehat{\phi}'_i \mathbf{F}_t + \varepsilon_{i,t},$$

where $t = 1, ..., T$, and retain the OLS estimates $\widehat{\mathbf{F}}_t$.

- Pass 3: run a (time-series) regression of y_{t+1} on $\widehat{\mathbf{F}}_t$:

$$y_{t+1} = \beta_0 + \beta' \widehat{\mathbf{F}}_t + \eta_{t+1},$$

retain the OLS estimates $\widehat{\beta}_0$ and $\widehat{\beta}'$, and use them in combination with $\widehat{\mathbf{F}}_t$ to construct the forecast $\widehat{y}_{t+1} = \widehat{\beta}_0 + \widehat{\beta}' \widehat{\mathbf{F}}_t$.

KP show that the 3PRF factor estimator $\widehat{\mathbf{F}}_t$ is consistent for the space spanned by the true factors. Moreover, they demonstrate that the 3PRF based forecast $\widehat{y}_{t+1} = \widehat{\beta}_0 + \widehat{\beta}' \widehat{\mathbf{F}}_t$ converges to the unfeasible forecast $\beta_0 + \beta' \mathbf{F}_t$ when N and T diverge. In addition,

$$\frac{\sqrt{T}(\widehat{y}_{t+1} - \beta_0 + \beta' \mathbf{F}_t)}{Q_t} \xrightarrow{d} N(0, 1),$$

where Q_t is defined in KP.

For the case in which there is just one $\mathbf{f}_t$ factor, KP suggest directly using the target variable y as proxy $\mathbf{z}$. They refer to this case as target-proxy 3PRF. In the case of more factors, they propose to either use proxies suggested by theory, or a simple automated procedure, which can be implemented in the following steps, indicating a proxy by r_j with $j = 1, ..., L$.

- Pass 1: set $r_1 = y$, and obtain the 3PRF forecast $\widehat{y}_t^1$ and the associated residuals $e_t^1 = y_t - \widehat{y}_t^1$.

- Pass 2: set $r_2 = e^1$, and obtain the 3PRF forecast $\widehat{y}_t^2$ using r_1 and r_2 as proxies. Obtain the associated residuals $e_t^2 = y_t - \widehat{y}_t^2$.

- ...

- Pass L: set $r_L = e^{L-1}$, and obtain the 3PRF forecast $\widehat{y}_t^L$ using $r_1, r_2, ... r_L$ as proxies.

Finally, KP study the relationship between 3PRF and PLS, see Wold (1975) and Helland (1990) for theoretical results on PLS and, e.g., Kapetanios, Marcellino and Papailias (2015) for an application to GDP forecasting. Specifically, KP show that PLS is a special case of 3PRF obtained when the predictors are standardized, the first two regression passes of 3PRF are run without an intercept, and the proxies are automatically selected.

2.2 The MF-3PRF

We now consider the case in which the target variable y (or the proxies $\mathbf{z}$) are sampled at a lower frequency than the indicators $\mathbf{x}$. Next, we allow some of the components of $\mathbf{x}$ to also be available only at low frequency. Finally, we study the case of ragged edges in $\mathbf{x}$.

2.2.1 Low frequency target or proxy variables and high frequency indicators

This is an empirically common situation. It arises, for example, when the target variable is GDP growth or GDP deflator inflation, which are available on a quarterly basis, while the indicators are monthly, e.g., industrial production and its components, labor market variables, financial indicators and survey variables.

To cope with the frequency mismatch, we propose modifying the steps of 3PRF as follows.

- Pass 1: as in 3PRF. Run at low frequency by temporally aggregating the high-frequency indicators.

- Pass 2: as in 3PRF, run at high frequency to obtain high-frequency predictive factor(s).

- Pass 3: use the U-MIDAS approach to construct a nowcasting or forecasting model that links the low-frequency target variable to the high-frequency factors.

We label the resulting procedure the mixed-frequency three-pass regression filter, MF-3PRF.

A few comments on each step are in order. First, as the regression in Pass 1 of 3PRF is static, running it in MF-3PRF with, say, quarterly rather than monthly indicators leads to consistent parameter estimators, which have the same properties as in 3PRF. As an alternative, one could interpolate the low-frequency target variable and run Pass 1 in high frequency. Empirically, we find no major differences between the two approaches; see Section 4.4.

Second, in step 2, MF-3PRF uses monthly rather than quarterly indicators, but the properties of the resulting factor estimators are the same as in 3PRF, as only the cross-sectional dimension is exploited in this step.

Third, in step 3, we suggest the use of the U-MIDAS method of Foroni, Marcellino and Schumacher (2015), as the frequency mismatch in macroeconomic applications is typically small (often monthly/quarterly). This requires splitting the estimated monthly factors $\widehat{\mathbf{F}}_t$ into three quarterly factors ($\widehat{\mathbf{F}}_t^1$, $\widehat{\mathbf{F}}_t^2$, and $\widehat{\mathbf{F}}_t^3$), where the first (second/third) new quarterly series contains the values of $\widehat{\mathbf{F}}_t$ in the first (second/third) month of each quarter. Next, $\widehat{\mathbf{F}}_t^1$, $\widehat{\mathbf{F}}_t^2$, and $\widehat{\mathbf{F}}_t^3$ are used as explanatory variables for y_t in the third step of 3PRF, thus balancing the frequency of the left- and right-hand side variables while maintaining the linearity of the equation and still using all available information.[3]

---

[3] In the case of a larger frequency mismatch, e.g., for quarterly and daily data, U-MIDAS can be replaced with MIDAS; see, e.g., Ghysels, Santa-Clara and Valkanov (2004) and Ghysels, Sinko and Valkanov (2007), which however requires the use of nonlinear least squares for parameter estimation. Moreover, empirically, the lags of $\widehat{\mathbf{F}}_t^1$, $\widehat{\mathbf{F}}_t^2$, and $\widehat{\mathbf{F}}_t^3$ could also matter for forecasting the target variable, and the proper lag length can be selected by means of information criteria.

The third step of MF-3PRF, the forecasting step, can be implemented based on the direct method, as in KP. As an alternative, a VAR model for $\mathbf{F}_t$ can be added to the system in (1)-(3):

$$\mathbf{F}_t = \sum_{k=1}^{K} \mathbf{\Theta}_k \mathbf{F}_{t-k} + \mathbf{v}_t. \tag{4}$$

It can be used, in combination with the estimated factors $\widehat{\mathbf{F}}$, to generate forecasts of $\widehat{\mathbf{F}}_{t+h}$, to be inserted in an iterated procedure to predict y_{t+h}. See, e.g., Marcellino, Stock and Watson (2006) for a comparison of the properties of direct and iterated forecasts.

Note that the MF-3PRF inherits all of the theoretical properties of 3PRF, under the same assumptions, as it is based on the same three steps and an efficient use of all available information. Its finite sample properties will be evaluated in Section 3 by means of Monte Carlo experiments and compared with those of 3PRF run on aggregated quarterly data to assess the additional value of the high-frequency indicators.

Finally, Kapetanios, Marcellino and Papailias (2015) apply U-MIDAS to obtain a mixed-frequency version of PLS (MF-PLS). As we noted in the previous subsection that PLS is a special case of 3PRF, MF-PLS can also be considered as a special case of MF-3PRF.

2.2.2 Low frequency indicators

Not only the target or proxy variables in (1)-(2) but also some of the indicators in $\mathbf{x}_t$ in (3) can be available in low frequency only. This situation can be addressed using the Kalman filter, as in the case of ragged edges described in the next subsection, which requires the specification of a model for the unobservable factors (e.g., as in (4)). As an alternative, the EM algorithm of Stock and Watson (2002) can be applied, together with standard principal component analysis (PCA) and without specifying a model for the factors. In both cases, the systematically missing observations in the low-frequency variables are replaced with their best estimates, and then the analysis proceeds as in the previous subsection.

Let us consider in some detail the working of the EM algorithm; further details are provided in, e.g., Marcellino and Schumacher (2010) and Angelini, Henry and Marcellino (2006). Let us assume for simplicity that the $\mathbf{x}$ variables are standardized. Focusing on a variable i, let the vector $\mathbf{x}_i^{\text{obs}}$ contain the available observations and formulate the relationship between observed and not fully observed data by

$$\mathbf{x}_i^{\text{obs}} = \mathbf{A}_i \mathbf{x}_i, \tag{5}$$

where $\mathbf{A}_i$ is in general a $T \times Tk$ matrix that addresses the mixed frequencies (or missing values) with full row rank T. In the event that no observations are missing, $\mathbf{A}_i$ is the identity matrix. In the event that an observation every k is missing (point-in-time or stock sampling), $\mathbf{A}_i$ is the identity matrix with the rows corresponding to missing observations deleted. In the event of average or flow

sampling, again with frequency k and assuming, for example, that $k = 3$, the $\mathbf{A}_i$ matrix takes the form

$$\mathbf{A}_i = \begin{matrix} 1 & 1 & 1 & 0 & 0 & 0 & ... & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & ... & 0 & 0 & 0 \\ ... & ... & ... & ... & ... & ... & ... & ... & ... & ... \\ 0 & 0 & 0 & 0 & 0 & 0 & ... & 1 & 1 & 1 \end{matrix} .$$

The EM algorithm proceeds as follows:

1. Provide an initial (naive) guess of observations $\widehat{\mathbf{x}}_i^{(0)}$, $\forall i$. These guesses, together with the fully observable monthly time series, yield a balanced dataset $\widehat{\mathbf{x}}^{(0)}$. Standard PCA provides initial monthly factor estimates $\widehat{\mathbf{F}}^{(0)}$ and loadings $\widehat{\boldsymbol{\Phi}}^{(0)}$.

2. **E-step:** An updated estimate of the missing observations for variable i is provided by the expectation of $\mathbf{x}_i$ conditional on observations $\mathbf{x}_i^{\text{obs}}$, factors $\widehat{\mathbf{F}}^{(j-1)}$ and loadings $\widehat{\boldsymbol{\Phi}}_i^{(j-1)}$ from the previous iteration:

$$\widehat{\mathbf{x}}_i^{(j)} = \widehat{\mathbf{F}}^{(j-1)}\widehat{\boldsymbol{\Phi}}_i^{(j-1)} + \mathbf{A}_i'(\mathbf{A}_i\mathbf{A}_i')^{-1}\left(\mathbf{x}_i^{\text{obs}} - \mathbf{A}_i\widehat{\mathbf{F}}^{(j-1)}\widehat{\boldsymbol{\Phi}}_i^{(j-1)}\right). \tag{6}$$

The update consists of two components: the common component from the previous iteration $\widehat{\mathbf{F}}^{(j-1)}\widehat{\boldsymbol{\Phi}}_i^{(j-1)}$, plus the low-frequency idiosyncratic component $\mathbf{X}_i^{\text{obs}} - \mathbf{A}_i\widehat{\mathbf{F}}^{(j-1)}\widehat{\boldsymbol{\Phi}}_i^{(j-1)}$, distributed by the projection coefficient $\mathbf{A}_i'(\mathbf{A}_i\mathbf{A}_i')^{-1}$ on the high-frequency periods. Repeat this E-step for each i to obtain a balanced dataset.

3. **M-step:** Re-estimate the factors and loadings, $\widehat{\mathbf{F}}^{(j)}$ and $\widehat{\boldsymbol{\Phi}}^{(j)}$ by PCA of $\widehat{\mathbf{x}}^{(j)}$, and return to step 2 until convergence.

After convergence, the EM algorithm provides high-frequency (monthly) factor estimates and estimates of the missing values of the time series $\widehat{\mathbf{x}}_i$, for each i. We can then form a balanced panel $\widehat{\mathbf{x}}$ and apply to it the MF-3PRF described in the previous subsection.

2.2.3 Ragged edges

Mixed-frequency sampling generates systematic patterns of missing observations in $\mathbf{x}$, but other types of missing observations are also frequent in empirical analysis, due to the different start and/or end dates of some indicators. The case of missing observations at the end of the sample, due to different release timing, was labeled the "ragged edge" by Wallis (1986). We focus on this situation, as it is more relevant for forecasting, but missing observations at the start of the sample can be handled in a similar manner (as can scattered missing observations).

In addition to the EM algorithm that can also be applied in this context with a proper choice of the $\mathbf{A}_i$ matrix, we consider three approaches suggested in the factor MIDAS case by Marcellino and Schumacher (2010), to whom we refer for additional details and the computational aspects. First,

the use of the Kalman filter, where the VAR for the factors in (4) contains the state equations to be used in combination with the observation equations in (3). Note that this approach can also be used to cope with mixed frequencies in **x**; see, e.g., Mariano and Murasawa (2003, 2010).

The second method simply requires fitting time series models to the variables with ragged edges (AR(2), for example) and then replacing the missing observations at the end of each time series with their forecast values.

The third method is even simpler, using vertical re-alignment of the data, as suggested, e.g., by Altissimo et al. (2010) for estimating the New Eurocoin indicator.

The use of the Kalman filter to address ragged edges is optimal, in the sense of producing the best linear estimates of the missing observations conditional on the correct specification of the state space form. However, it requires the specification and estimation of a model for the unobservable factors. The second and third methods are much easier to implement but can be suboptimal. We will experiment empirically with all three procedures.

3   Monte Carlo Evaluation

In this section, we assess the finite sample performance of the MF-3PRF, including relative to that of 3PRF and PCA. In the first subsection, we describe the data generation process (DGP) used in the experiments and the structure of the forecasting exercise. In the second subsection, we present the results.

3.1   DGP and structure of the exercise

We use the same DGP as KP but assume that the data are generated at the monthly frequency while y_t is only observable at quarterly frequency.

The factors are generated as

$$f_t = \rho_f f_{t-1} + u_{f,t}, \quad u_{f,t} \sim i.i.d.N(0,1), \tag{7}$$

$$\mathbf{g}_t = \rho_g \mathbf{g}_{t-1} + \mathbf{u}_{g,t}, \quad \mathbf{u}_{g,t} \sim i.i.d.N(0,\mathbf{\Sigma}_g),$$

where $k_f = 1$, $k_g = 4$, such that $k = 5$. The matrix $\mathbf{\Sigma}_g$ is diagonal with elements 1.25, 1.75, 2.25, 2.75 (and hence the irrelevant factors are dominant with respect to f_t). We fix $\rho_f = 0.9$ and let ρ_g take values of 0.3 or 0.95.

The (unobservable) monthly target variable is

$$\tilde{y}_t = f_t + \sigma_y \eta_t, \quad \eta_t \sim i.i.d.N(0,1), \tag{8}$$

where σ_y is set at a value that guarantees that the R^2 of the infeasible best forecast of $\tilde{y}_t$ is 0.50. In accordance with the empirical application, we model the observed quarterly variable as the growth

rate (GDP growth in the application) of the quarterly average of the index I_t, which is created by the unobservable monthly variable. In formulae,

$$y_t = \frac{(I_{t-2} + I_{t-1} + I_t)/3}{(I_{t-5} + I_{t-4} + I_3)/3} - 1, \tag{9}$$

where

$$I_{t-\tau} = I_{t-5} \prod_{j=0}^{4-\tau} (1 + \tilde{y}_{t-4+j}).$$

Note that (9) calculates the exact quarterly growth rate, whereas empirical applications often approximate it by $y_t \simeq (\tilde{y}_{t-4} + 2\tilde{y}_{t-3} + 3\tilde{y}_{t-2} + 2\tilde{y}_{t-1} + \tilde{y}_t)/3$, as suggested by Mariano and Murasawa (2003). [4]

Regarding the **x** variables we have

$$\mathbf{x}_t = \mathbf{\Phi}\mathbf{F}_t + \varepsilon_t,$$

where $\mathbf{F}_t = \left[f_t, \mathbf{g}_t'\right]'$, the elements of $\mathbf{\Phi}$ are drawn from standard normals, and the idiosyncratic errors can be both temporally and cross-sectionally correlated. In particular

$$\varepsilon_{i,t} = a\varepsilon_{i,t-1} + \widetilde{\varepsilon}_{i,t},$$
$$\widetilde{\varepsilon}_{i,t} = (1 + d^2)v_{i,t} + dv_{i-1,t} + dv_{i+1,t},$$

where each v_i is standard normal, $a = 0, 0.9$, and $d = 0, 1$. We set $T = 100, 200, 400$, $N = 100$ and draw 500 replications from each DGP, but to save space we only report results for some parameter combinations, making the full set of results available upon request.

We assume that the quarterly value of y_t is released at the end of the following quarter, and we are interested in one-step-ahead forecasting, nowcasting, and backcasting. We define forecasting as the situation in which we have some information on a quarter $t-1$ and want to predict the target in the next quarter y_t. An example would be that we are somewhere in Q4 2013 and attempt to predict GDP growth in the next quarter, i.e., in Q1 2014. If we are at the very beginning of Q4 2013, i.e., in October 2013, we do not yet have any data on Q4. In the Monte Carlo exercise we represent this as a situation in which indicators are available up to the third month of the quarter $t-2$, and we predict the target y_t. At the other extreme, we are at the end of December 2013 and predict GDP growth in Q1 2014. In this situation, most indicators would be available up to and including

---

[4]This expression is obtained by first approximating the arithmetic means of the monthly indices with geometric means $y_t \simeq (I_{t-2}I_{t-1}I_t)^{1/3}/(I_{t-4}I_3)^{1/3} - 1$, where the initial index level, I_{t-5}, is set to 1 without loss of generality because it merely scales the later levels. A second step log-approximates the quarterly growth rate, $y_t \simeq \log(1 + y_t) = \frac{1}{3}(\log(I_{t-2}) + \log(I_{t-1}) + \log(I_t) - \log(I_{t-4}) + \log(I_{t-3}))$, and a third step finally uses $\log(I_{t-\tau}) = \sum_{j=0}^{4-\tau} \log(1 + \tilde{y}_{t-4+j}) \simeq \sum_{j=0}^{4-\tau} \tilde{y}_{t-4+j}$. Substituting this approximation of the monthly index levels into the log-approximation of the quarterly growth rates yields the linear expression suggested by Mariano and Murasawa (2003).

November. Thus, this forecast is represented in the Monte Carlo exercise as the situation in which data are available up to the 2^{nd} month of the the quarter $t-1$.

We define nowcasting as the situation in which one wishes to predict the current quarter. Taking up the previous example, the nowcasts for Q1 2014 would be generated at some point during Q1 2014. If the nowcast is performed in January, most indicators would be available up to December. In the Monte Carlo simulation this corresponds to the exercise whereby we forecast the target in quarter t and have data up to the third month of the quarter $t-1$. If the nowcast is performed in March, however, we simulate an economy in which data are available up to and including the second month of the quarter t.

Backcasting, finally, is the case in which we assume that we are in April 2014. We therefore observe, for most indicators, data up to March 2014; i.e., the indicator dataset for Q1 2014 is complete. Using this information, we want to predict what GDP growth had likely been in Q1 2014. In the corresponding Monte Carlo simulation, we observe data up to and including the third month of quarter t and forecast the target of quarter t, y_t.

In summary, in the Monte Carlo experiments, we will consider 7 different information sets or points in time, during each of which we will forecast the target in quarter t. In the first simulation, we observe data only up to the third month of the quarter $t-2$. In the following simulations, we then gradually extend the dataset in six steps until we are in the backcasting situation where we observe the indicators up to and including the third month of the quarter t.

We will evaluate the performance of a number of different models. We first compare a simple version of the MF-3PRF with 3PRF based on aggregated quarterly data for the indicators $\mathbf{x}$. In the MF case, we begin by extracting the monthly factor f_t using the first two passes of the MF-3PRF. In the third pass we run a U-MIDAS regression using the factor estimates for the latest five months. Next, we compare MF-3PRF with mixed-frequency PCA. For the PCA we extract the first five principal components from the monthly $\mathbf{x}$ variables and then translate them into a GDP forecast using U-MIDAS. We also report the results for the infeasible best forecast based on the true relevant factor.

3.2   Results

The selected results of the Monte Carlo exercise are reported in the three panels of Figure 1. Panel A compares the RMSFE of the quarterly 3PRF and the MF-3PRF. The lines with filled markers refer to a sample size of 200, while the empty markers refer to the larger sample ($T = 400$). Unsurprisingly, in all cases, the RMSFEs decline monotonically with the enlargement of the information set. For the quarterly 3PRF, the information set expands only once the quarter is full, hence the steps. Interestingly, for the smaller sample, the quarterly 3PRF performs slightly better whenever a quarter's data are fully observed. This is because, in the quarterly specification, the U-MIDAS regression estimates only one parameter, whereas the monthly specification requires five parameters

to be estimated. Once the sample size is increased to $T = 400$, parameter estimation becomes more precise and the quarterly specification is dominated by the mixed-frequency specification. Similar results are obtained for other DGP parameter values.

Panels B and C of Figure 1 report the RMSFE for the MF-3PRF, PCA, and infeasible best for two different parameter specifications. We find that the ranking of MF-3PRF and PCA is not clear cut. With a sample size of 200, MF-3PRF performs better than PCA when there is substantial persistence in the idiosyncratic errors ($a = 0.9$) as can be seen in Panel B. PCA performs better if the relevant factor is strongly autocorrelated ($\rho_g = 0.9$) but there is no temporal or cross-sectional correlation in the idiosyncratic errors ($a = d = 0$) as shown in Panel C.

It is important to see that the relative performance can change with the sample size. Consider, for example, Panel C: if the training sample becomes smaller, the performance of PCA deteriorates strongly. This is because with PCA, 25 parameters need to be estimated (five lags for each of the five factors), while the MF-3PRF requires the estimation of only five parameters, as there is only one 3PRF factor. This disadvantage of PCA of course becomes less relevant with larger samples, as can be seen from Panel B, where the PCA performance improves more than the MF-3PRF's with the doubling of the sample size. The infeasible best forecast of course always dominates; however, the extent of the outperformance is specification dependent.

Overall, the Monte Carlo evaluation confirms the good finite-sample properties of MF-3PRF and the usefulness of adding higher-frequency information for nowcasting and forecasting. However, as in the case of KP, the ranking of the PCA and 3PRF factor estimation methods is not clear-cut but depends on the parameterization of the DGP and on the sample size. Hence, we now turn to an assessment of their relative merits in empirical applications.

4   Nowcasting US GDP Growth

To illustrate empirically the workings of the MF-3PRF, we focus on nowcasting and short-term forecasting of GDP growth, a key issue for policy makers that has also attracted substantial interest in the academic literature. In this section, we focus on the USA; in the next section, we will consider a variety of other countries. After discussing the data and in-sample and out-of-sample results, we also consider the performance of a set of extensions and modifications of the basic MF-3PRF procedure.

4.1   Data

We use a dataset collected on 28 January 2014, with values dating back to January 1984. Some series are also available before 1984, but we prefer to begin our analysis around the beginning of the great moderation (see, for example, McConnell and Perez-Quiros, 2000). The dataset consists of 98 monthly variables spanning hard indicators, such as industrial production and retail sales, soft indicators, such as the ISM survey and its subcomponents, financial indictors, such as the VIX and

various interest rates, and price indicators, such as CPI and PPI indices. All variables are seasonally adjusted and, if nonstationary, transformed by calculating the growth rate. The target is quarterly real GDP growth expressed at an annual rate.

4.2 In-sample results

We begin the empirical evaluation of the MF-3PRF model with a specification that is as simple and intuitive as possible: the first pass is estimated at a quarterly frequency with the target variable as the only proxy. Specifically, the 98 monthly predictor variables are aggregated to quarterly frequency. The nonstationary variables are then transformed into growth rates. We regress each resulting predictor variable on real GDP growth and retain the slope estimates. In the second pass, we extract the 3PRF factor by regressing the cross-section of predictor variables on the slope estimates from the first pass.

Figure 2, panel A, plots the resulting factor over time. The dynamics align well with the common view of the history of the US business cycle. The factor declines sharply at the beginning of periods that the NBER classifies as recessions (shaded areas), remains at low levels during the recessions, and recovers with the beginning of the expansion. The initially very slow recovery after the recession of the early 2000s is highly visible as is the very sluggish recovery following the Great Recession.

This good in-sample performance is of course not unique to the 3PRF approach. Panel B of Figure 2 plots the first three common components of the dataset against the 3PRF factor (all factors are normalized). Clearly, the first PCA factor captures the growth dynamics in the US as well. Interestingly, however, the levels of the first common component and the 3PRF factor diverge somewhat after the Great Recession. A similar effect is also apparent when comparing the third common component to the 3PRF factor.

Below, we will use the factors to forecast GDP growth. It is therefore natural to ask how well they align in-sample with real GDP growth. For this purpose, we regress quarterly GDP growth on the static MF-3PRF factor. The first column of Table 1 presents the R^2 and adjusted R^2 of three versions of this regression. The first version ("quarter average") regresses quarterly GDP growth on the quarterly average of the factor. The second version ("monthly: $t, t-1, t-2$") regresses quarterly GDP growth on the three monthly values of the factor during the quarter of interest, according to the U-MIDAS approach. The third version ("growth rate of index") first constructs a monthly index of the factor, then calculates the quarterly averages of this index, and finally calculates the quarterly growth rates onto which quarterly GDP growth is then regressed. All three versions yield an R^2 on the order of 60%. Repeating the same exercise with the first principal component instead of the MF-3PRF factor yields a slightly lower R^2 (second column of Table 1; "PCA1"). Additionally including the second component still yields R^2 that are slightly below the fit of the MF-3PRF factor, as can be seen from the third column ("PCA2"). Even when including the third component ("PCA3"), the fit remains below the MF-3PRF R^2 for two of the three regressions. Note that the

PCA3 regressions use three times as many covariates as the MF-3PRF regression. In fact, when considering the adjusted R^2, the fit of the MF-3PRF regressions is always better than the principal-component-based regressions (lower panel of Table 1). This result is not particularly surprising, as the 3PRF approach specifically gears the factors toward being able to explain variation in GDP growth. Nevertheless, the results suggest that the 3PRF approach may yield quite a parsimonious model, which may be helpful in terms of out-of-sample properties.

4.3 Forecasting exercise

In the next step, we turn to the out-of-sample analysis. Beginning in Q1 2000, in each quarter, we replicate the ragged edge as observed on 28 January 2014. We estimate the model on this pseudo real-time dataset and make a first backcast of GDP growth (Q1 2000, $h = 1$), a nowcast (Q2 2000, $h = 2$), and three short-term forecasts (Q3 2000, Q4 2000, Q1 2001, $h = 3, 4, 5$, respectively), applying the direct forecasting approach to the U-MIDAS equations of the third pass of MF-3PRF. The dataset is then successively enlarged and the pseudo out-of-sample exercise is repeated in every quarter. Note that we do not consider vintage data but use the final dataset.[5]

In this exercise, the ragged edge problem becomes central. On 28 January 2014, for example, the Philly Fed Manufacturing Survey for January was already available. Industrial production and retail sales are only available up to December 2013, while data on international trade, construction spending, and consumer credit are only observed up to November 2013. To begin, we use the Kalman filter as outlined in Section 2.2.3 to address the ragged edge. Figure 3 presents the estimated pseudo real-time factors. The black vertical line denotes January 2000, when the out-of-sample exercise begins. For comparability, all factors are normalized by their mean and standard deviation prior to 2000. Several observations are worth mentioning.

In the in-sample period, the estimate of the factor is little revised: the real-time estimates in black generally overlap with one another and also agree with the latest estimate in red that uses all available data (up to 28 January 2014). The slight differences that sometimes emerge are due to the fact that the expanding dataset leads to changes in the slope estimates in the first pass and to changes in the estimated law of motion for the factor. However, regarding the model's general interpretation of economic history, in the sense of identifying whether growth has been accelerating or decelerating, these revisions are irrelevant.

In the out-of-sample evaluation period, the revisions become more important. In Figure 3, the solid black lines represent the real-time estimates of the factor for the period for which there is some data available to inform the factor estimate. For the period thereafter, the model generates a forecast that is entirely driven by the factor's law of motion. This forecast is represented by the grey dotted lines. Clearly, the end-point of the factor estimate is an important determinant of the forecasted factor dynamics and ultimately of the forecasting performance. In the period after the

---

[5]Bernanke and Boivin (2003) and Schumacher and Breitung (2008) demonstrate that the conclusions regarding forecasting performance do not change significantly if vintage data instead of final data are used.

14

Great Recession, for example, the real-time factor estimates are, in several instances, considerably below the level that the model identifies based on the full dataset. This happens if the data with a short publication lag, i.e. the data that are available at the ragged edge, send a worse signal than the data that become available subsequently. Note that the model used in Figure 3 has two mechanisms to discount the early signals if warranted: first, the regression in Pass 1 ensures that the data are only used up to the degree that is actually relevant for the target, and second, the Kalman filter discounts noisy series.

As outlined in Section 2.2.3, one can alternatively use vertical re-alignement of the data (Shift) or univariate AR(2) forecasts (AR) to fill the ragged edge. Table 2 presents the RMSFEs and MAFEs of the three approaches to addressing the ragged edge. The table also presents the performance of a number of benchmark models. The three simple benchmarks are the forecast derived from an AR(p), where p is chosen according to AIC (referenced as AR), from a random walk (referenced as RW), and the average growth rate in the past (referenced as CONST). We also consider the forecasting performance of principal-component-based dynamic factor models, where the factors are first estimated at a monthly frequency and the GDP forecast is then obtained using U-MIDAS. We used one to three factors (referenced as PCA1, PCA2, PCA3), as in the in-sample analysis.

When $h = 1$ or 2, the Kalman-3PRF provides the lowest values for both evaluation criteria, and the gains with respect to the other methods are often statistically significant, according to a Diebold-Mariano (1995) test. When $h = 3$, AR-3PRF performs best, though the differences with respect to Kalman-3PRF are small and not significant. When $h = 4$ or 5, Shift-3PRF is best. Hence, for any horizon, 3PRF outperforms both the PCA-based dynamic factor models and the simple univariate benchmarks. The gains in terms of RMSFE with respect to PCA-based models are approximately 10% for $h > 1$ and even larger with respect to the simple benchmarks.

Figure 4 plots actual GDP growth in black against the real-time forecasts produced by the Kalman-3PRF model and two of the benchmarks. The MF-3PRF model captures the general dynamics quite well for the backcast ($h = 1$) and the nowcast ($h = 2$). Moreover, for $h = 3$ and $h = 4$, the model captures the dynamics outside of the crisis. However, once the longest forecasting horizon is reached ($h = 5$), the forecast is essentially constant except for a dip following the crisis. This is reminiscent of the general finding in the literature that time-series methods typically provide good back- and nowcasts, but they become less reliable once the forecasting horizon is extended. Regarding the relative performance of the three models, we observe that for $h = 1, 2$ the factor model captures the GDP dynamics better than does the AR model - particularly so during the crisis. We also see that during the initial recovery phase, the PCA1 model tends to underpredict GDP growth, while MF-3PRF is better able to capture the recovery.

Figure 5 provides an impression of the relative performance of the models over time. It shows the 4-year rolling RMSEs of the three models. Three observations are worth emphasizing here. First, prior to the crisis, all three models performed quite similarly. Second, the AR benchmark performs

poorly during the crisis because this model does not benefit from the information contained in the high-frequency indicators and therefore cannot detect the rapid deceleration in economic activity in real-time. Third, the PCA1 model clearly underperforms relative to the MF-3PRF after the crisis for $h = 2, 3$.

4.4   Robustness analysis

In this subsection, we consider various alternative formulations and extensions of the 3PRF method to assess whether even better results can be obtained than with the benchmark MF-3PRF specification used above.

A first modification to our model concerns Pass 1. In our baseline specification, we simply used contemporaneous GDP growth as the proxy. One may ask if proxies that are specific to the forecasting horizon improve performance. In particular, one could use lagged GDP as proxy for the backcast, contemporaneous GDP as proxy in the nowcast, a one quarter lead when performing the one-step-ahead forecast ($h = 3$), and so forth. With this approach, the relative importance of the indicators for the factor extraction can vary depending on the forecasting horizon, e.g., indicators with lagging (leading) properties will be strongly (lightly) weighted in the backcast but will have little (large) importance for the factor extraction for a forecast. Panel A in Figure 6 presents the RMSFE and MAFE associated with a model in which the weights used for the factor extraction are horizon specific and the resulting factors are then translated into a GDP forecast using a U-MIDAS regression. Figure 6 also presents the forecast performance of the Kalman-3PRF benchmark specification discussed above, of the PCA1 model, and of the simple AR benchmark. For horizons $h = 1, 2, 3, 4$ the performance changes little, but we observe a small improvement for $h = 5$.

We also experimented with interpolating the GDP level using a spline to obtain monthly GDP growth rates. Using these, we can then run Pass 1 at the monthly frequency. The resulting out-of-sample performance is very similar to that of the reference specification.

A second modification concerns Pass 3. In our benchmark specification, we first extract the factors and then translate the monthly factors into GDP forecasts by estimating horizon-specific U-MIDAS equations. Alternatively, the iterated or indirect approach can be used. One may first use the law of motion estimated in the Kalman filter to forecast the factor. In a second step, the contemporaneous correlation of the factors and GDP growth is used to translate the factor forecast into a GDP forecast. We consider two versions of this indirect forecasting approach. In the first version, we estimate a contemporaneous U-MIDAS regression and use this equation to generate the GDP forecast. In the second version, we aggregate the monthly factor to the quarterly frequency and only then estimate a simple regression in which the quarterly factor explains GDP growth. As can be seen from Panel B in Figure 6, the performance of these indirect approaches is virtually indistinguishable from our Kalman-3PRF benchmark specification.

A third modification uses the L-autoproxy procedure with $L = 2, 3, 4, 5$, as suggested by KP and

reviewed in Section 2.1, to extract additional forecasting factors. Panel C in Figure 6 presents the results: generally, $L = 1$, as used in the Kalman-3PRF benchmark, performs best, and additional factors tend to actually decrease forecasting performance.

Finally, we also experimented with the inclusion of exogenously defined recession dummies in the third pass. This improved forecasting performance somewhat. However, the NBER recession dates contain in-sample information that was as such not available at the time of the recessions. To adress this issue, future research may attempt to model two regimes in Pass 3 based on the factor extracted in Pass 2, or, more generally allow for parameter time variation in Pass 1 and Pass 3.

4.5 Temporal evolution of the nowcasts

We conclude our empirical analysis for the USA by considering the temporal evolution of the nowcasts over the quarter. To fix ideas, let us consider GDP growth in Q1 2014, the official advance estimate for which was released at the end of April 2014. In the out-of-sample analysis discussed above, we used the dataset as available at the end of January 2014. At this point in time, barely any data for Q1 2014 were available (except for financial indicators and some results of the regional business surveys), and hence the nowcast ($h = 2$) was informed primarily by data from the previous quarter. The out-of-sample analysis replicated this ragged edge in each evaluation period and compared the resulting forecasts to actual GDP growth. In the following, we move through Q1 in weekly steps and evaluate how the pseudo out-of-sample forecasting performance changes as more and more data arrive.

Figure 7 presents the results: on the x-axis are the days on which the dataset is collected; on the y-axis, we plot the RMSFEs that result from various models. We consider the Kalman-3PRF benchmark specification and the restricted indirect 3PRF model discussed in the previous section. Additionally, we consider the three PCA models with one to three factors and a simple AR model. Several observations can be made. First, the performance of the AR model does not change during most weeks because no additional data arrive - the only exception is the decrease in the forecasting error from the first to the second week because during the second week the advance GDP estimate for Q4 arrives, which of course improves the forecasting performance of the model. Second, the results from Table 2 are confirmed: the dots for the first week correspond to the values presented in Table 2. Third, the ranking of the models remains rather stable over time, and the Kalman-3PRF reference specification continues to outperform once additional data arrive. However, the improvement for the PCA-based models (PCA1-3) is more pronounced such that the RMSFE converge toward the end of the exercise. In fact, in the very last week, the backcast produced by the PCA3 model slightly outperforms that of the 3PRF models.

Overall, this analysis further highlights the usefulness of higher-frequency information, and supports the use of the MF-3PRF approach for nowcasting.

5 Nowcasting GDP growth in other countries

In this section, we evaluate whether the good performance of the MF-3PRF we have detected for the USA is also common in other countries. Specifically, we consider first the euro area as a whole and its four largest member states: France, Germany, Italy and Spain. Next, we consider two other large developed economies, UK and Japan. Finally, two large emerging countries, Brazil and India, and two smaller ones, Korea and Taiwan are analized. The choice of these countries is dictated primarily by the availability of a large dataset for a sufficiently long period.

The structure of the forecasting exercise is akin to that for the USA such that, in a pseudo real-time context, we consider five different forecast horizons for each country and a variety of competing models, listed in Table 3. In addition, we also consider a very large dataset that includes the indicators of all countries under evaluation including the USA. Additional details regarding the data are reported in the next subsection, while the empirical results are discussed in the second subsection. The final subsection evaluates the usefulness of variable pre-selection in the context of a very large dataset.

5.1 Data

Table 4 reports the number of indicators in each country-specific dataset, the beginning of the estimation period and the beginning of the pseudo out-of-sample evaluation period. The beginnings of the estimation and the evaluation periods are primarily functions of data availability. The extraction date for all countries is 28 January 2014. The global dataset, in which all country-specific data are pooled, features 832 indicators.[6] When estimating the models using the global dataset, estimation and evaluation periods are left unchanged for each country to ensure comparability of the results.[7]

5.2 Results

The results are summarized in Table 5 and Figure 8, while Table 3 lists all the models under evaluation. Let us begin the discussion with the euro area. A first, striking finding is that the use of the pooled information set is generally better than the use of euro area data only, likely due to the substantial remaining heterogeneity among euro area countries (see, e.g., Marcellino, Stock and Watson (2003) for an early evaluation of this issue). A second important result is that the baseline MF-3PRF applied to the pooled dataset produces the lowest RMSFE for $h = 1, 2, 3$ followed by the pooled MF-IR3PRF. For $h = 4$, AR is best. For $h = 5$, the MF-3PRF, PCA1, and AR have essentially equal performance. Among the PCA methods, the use of a single principal component is best, with its relative performance improving with the length of the forecast horizon.

---

[6] The global dataset also contains some variables from other countries, such as China, which is why the number of indicators is larger than the sum of the number of indicators in the individual country sets.

[7] A detailed list of all the variables is available upon request.

Regarding the four euro area countries, for both Italy and Spain, the MF-3PRF using the pooled dataset is generally better (in the case of Spain, local-information-based MF-3PRF or MF-IR3PRF are even better for h=1). For France, the pooled MF-3PRF is again best for $h = 1, 2$, while PCA with two factors is slightly better for $h = 3, 4, 5$. In the case of Germany, the baseline MF-3PRF with German data only is best for $h = 1$, the PCA with two/three factors performs best for $h =2/3$ and the simple AR-benchmark dominates for $h = 4, 5$.

Moving to the UK and Japan, for the former, the PCA approach with 3 factors or AR are generally best. While for Japan, the indicator-based 3PRF and PCA models only outperform the AR benchmark for $h = 1$, where MF-IR3PRF dominates.

Regarding the four emerging markets, we find that for India, PCA1 or PCA2 are best for $h = 1, 2$, while MF-IR3PRF is best for $h = 3, 4, 5$. For Brazil, Korea and Taiwan, MF-3PRF methods generally denominate.

Finally, Figure 9 in the Appendix reports, for each country, the MF-3PRF factors based on the country-specific and pooled datasets. The two types of factors are generally similar, with the global-data-based factor typically being more negative during the Great Recession.

Overall, we find some variety in the relative performance of the 3PRF and PCA, as in the Monte Carlo experiments. However, for many countries, and in particular for those in the euro area but also for Japan, Brazil, Korea and Taiwan, the MF-3PRF approach appears to work quite well for nowcasting and short-term forecasting GDP growth.

5.3 Indicator preselection

Boivin and Ng (2006) document that, in a principal component context, the forecasting performance often improves when the indicators are preselected or weighted. The gains come from the fact that such weighting schemes can avoid situations in which an irrelevant factor comes to dominate the relevant factor. The global dataset used above is a case in point: assume, for example, that we wish to forecast GDP growth for a very closed economy. If we extract the first few principal components from the global dataset, we will likely obtain factors that are related to the global business cycle but less related to the forces driving the cycle of our very closed economy. A preselection may help to avoid such an outcome in a principal components context. In the 3PRF case, in contrast, preselection is likely to be less necessary, as the algorithm automatically weights the indicators according to their relevance for the proxy variable(s). Nevertheless, Pass 1 would offer a very natural and simple opportunity to perform variable preselection, based on the outcome of an F-statistic for the significance of the regressor(s).

In the following, we report the out-of-sample performance of the MF-3PRF model when the indicators are preselected from the global dataset. Specifically, we drop all indicators from the dataset that have a p-value for the F-test that is higher than 1% in the first-step regression.[8] Table

---

[8]We ran similar evaluations for p-Values of 5%, 10%, 20%, and 50%. Qualitatively, the results do not change and

6 presents the results. In column 2, we report the number of selected indicators per country ranging from 57 out of a total of 832 series in the case of India to 473 in the case of the euro area. On average, somewhat less than 40% of the series are selected. Columns 3 to 7 present the relative RMSFEs - the model with preselection relative to the model using all indicators. Asterisks denote significant differences according to the Diebold-Mariano (1995) test. While there are some instances in which the models using the preselected indicator sets perform significantly better (e.g., the back- and nowcast for the US and Spain), there are other instances in which the model using the global dataset performs significantly better. In most cases, however, there is no significant difference with respect to the model using the global dataset.

This suggests that the 3PRF algorithm already performs relatively well in weighting the infor- mation and that additional weighting or preselecting is unnecessary – an encouraging result for the applied forecaster, implying that she need not be particularly concerned about which indicators to include and that when in doubt she should include an indicator in the dataset.

6   Conclusions

In this paper, we extended the three-pass regression filter of KP to allow for the use of large mixed- frequency datasets, possibly with ragged edges, in a forecasting context. The resulting method, labeled MF-3PRF, represents a simple but robust and efficient method for forecasting, nowcasting and backcasting.

The MF-3PRF inherits the good theoretical asymptotic properties of 3PRF. It also performs well in finite samples, according to Monte Carlo experiments, compared both to the single-frequency 3PRF and to the mixed-frequency PCA, though the ranking with the latter is not clear cut.

In the empirical application, the forecaster faces a number of modeling choices, which we sum- marize in Table 7. We found that, in general, MF-3PRF is a very useful approach for forecasting, nowcasting and backcasting GDP growth in the USA and a variety of other countries. Based on our forecasting exercises, our baseline recommendation is the Kalman specification with a U-MIDAS equation in Pass 3 as a powerful, robust and easy to implement short-term forecasting model. In the case of European countries, we suggest that a dataset pooling indicators from all major European countries should be used.

Finally, the method could be extended in a variety of ways, e.g. to cope with integrated variables and parameter time variation, and applied in other contexts, such as nowcasting fiscal variables. We leave these interesting extensions for further research.

---

are therefore not explicitly reported, but are available upon request.

References

[1] 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, 92, 1024-1034.

[2] Angelini, E., J. Henry, M. Marcellino (2006), " Interpolation and backdating with a large information set", Journal of Economic Dynamics & Control 30, 2693-2724.

[3] Angelini, E., G. Camba-Mendez, D. Giannone, G. Runstler, and L. Reichlin (2011), "Short term forecasts of Euro Area GDP growth", Econometrics Journal, 14, C25–C44..

[4] Banbura M, Giannone D, Reichlin L. (2011), " Nowcasting". In *Oxford Handbook on Economic Forecasting*, Clements MP, Hendry DF (eds). Oxford University Press: Oxford.

[5] Banbura M, Giannone D, Modugno, M., Reichlin L. (2012), "Now-casting and the real time data flow", CEPR WP 9112.

[6] Banbura, M. and M. Modugno (2014), "Maximum likelihood estimation of factor models on data sets with arbitrary pattern of missing data", Journal of Applied Econometrics, 29, 133–160,.

[7] Banbura, M., and G. Runstler (2011), "A Look into the factor model black box: Publication lags and the role of hard and soft data in forecasting GDP", International Journal of Forecasting, 27(2), 333-346.

[8] Barhoumi, K., S. Benk, R. Cristadoro, A. Den Reijer, A. Jakaitiene, P. Jelonek, A. Rua, G. Runstler, K. Ruth, and C. Van Nieuwenhuyze (2008): "Short-term forecasting of GDP using large monthly datasets. A pseudo real-time forecast evaluation exercise" , Occasional Paper Series 84, European Central Bank.

[9] Bencivelli, L., Marcellino, M. and Moretti, G.L., (2012), "Selecting predictors by Bayesian model averaging in bridge models", Banca d'Italia WP.

[10] Bernanke, B.S. and Boivin, J. (2003), "Monetary policy in a data-rich environment", Journal of Monetary Economics, 50, 525–546.

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

[12] Bulligan, G., Marcellino, M. and Venditti, F. (2015), "Forecasting economic activity with targeted predictors", International Journal of Forecasting, 31, 188-206.

[13] Carriero, A., Clark, T. and Marcellino, M. (2015), "Real-Time Nowcasting with a Bayesian Mixed Frequency Model with Stochastic Volatility", Journal of the Royal Statistical Society, Series A, forthcoming..

[14] Diebold, F.X. and Mariano, R. (1995), "Comparing Predictive Accuracy", Journal of Business and Economic Statistics 13, 253-263.

[15] Doz, C., D. Giannone, and L. Reichlin (2011), " A Two-step estimator for large approximate dynamic factor models based on Kalman filtering", Journal of Econometrics, 164(1), 188-205.

[16] Forni, M., M. Hallin, M. Lippi, L. Reichlin (2005), "The Generalized Dynamic Factor Model: One-sided estimation and forecasting", Journal of the American Statistical Association 100, 830-840.

[17] Foroni, C., and M. Marcellino (2013), "A Survey of Econometric Methods for Mixed-Frequency Data", Norges Bank, WP 2013/06.

[18] Foroni, C., and M. Marcellino (2014), "A comparison of mixed frequency approaches for nowcasting Euro area macroeconomic aggregates", International Journal of Forecasting, 30, 554-568.

[19] Foroni, C., M. Marcellino, and C. Schumacher (2015), "U-MIDAS: MIDAS regressions with unrestricted lag polynomial", Journal of the Royal Statistical Society, Series A, 178, 57-82.

[20] Ghysels, E., P. Santa-Clara, and R. Valkanov (2004), "The MIDAS touch: MIxed DAta Sampling regression models", mimeo, Chapel Hill, N.C.

[21] Ghysels, E., A. Sinko, and R. Valkanov (2007), " MIDAS regressions: Further results and new directions", Econometric Reviews, 26(1), 53-90.

[22] Giannone, D., L. Reichlin, and D. Small (2008), "Nowcasting GDP and inflation: The real-time informational content of macroeconomic data releases", Journal of Monetary Economics, 55, 665-676.

[23] Helland, I.S. (1990), "Partial Least Squares Regression and Statistical Models", Journal of Statistics, 17, 91-114.

[24] Kapetanios, G., Marcellino, M. and Papailias, F. (2015), "Variable Selection for Large Unbalanced Datasets Using Non-Standard Optimisation of Information Criteria and Variable Reduction Methods ", Computational Statistics and Data Analysis, forthcoming.

[25] Kelly, B. and S. Pruitt (2013), "Market Expectations in the Cross Section of Present Values", Journal of Finance, 68, 1721-1756.

[26] Kelly, B. and S. Pruitt (2014), "The Three-Pass Regression Filter: A New Approach to Forecasting with Many Predictors", Journal of Econometrics, forthcoming

[27] Kuzin, V., M. Marcellino and C. Schumacher (2013), "Pooling versus model selection for nowcasting GDP with many predictors: Empirical evidence for six industrialized countries", Journal of Applied Econometrics, 28(3), 392-411.

[28] Marcellino, M., J. Stock and M.W. Watson, (2003), "Macroeconomic forecasting in the Euro area: country specific versus Euro wide information", European Economic Review, 47, 1-18.

[29] Marcellino, M., J. Stock and M.W. Watson, (2006), "A Comparison of Direct and Iterated AR Methods for Forecasting Macroeconomic Series h-Steps Ahead", Journal of Econometrics, 135, 499-526.

[30] Marcellino, M., and C. Schumacher (2010), "Factor-MIDAS for now- and forecasting with ragged-edge data: A model comparison for German GDP", Oxford Bulletin of Economics and Statistics, 72, 518-550.

[31] Marcellino, M., Porqueddu, M. and Venditti, F. (2015), "Short-term GDP forecasting with a mixed frequency dynamic factor model with stochastic volatility", Journal of Business and Economic Statistics, forthcoming.

[32] Mariano, R., and Y. Murasawa (2003), "A new coincident index of business cycles based on monthly and quarterly series", Journal of Applied Econometrics, 18(4), 427-443.

[33] Mariano, R., and Y. Murasawa (2010), "A coincident index, common factors, and monthly real GDP", Oxford Bulletin of Economics and Statistics, 72(1), 27-46.

[34] Schumacher, C., and Breitung, J. (2008), " Real-time forecasting of German GDP based on a large factor model with monthly and quarterly data", International Journal of Forecasting, 24, 368–398.

[35] Stock, J., and M.Watson (2002), "Forecasting using principal components from a large number of predictors", Journal of the American Statistical Association, 97(460), 1167-1179.

[36] Wallis, K. (1986), "Forecasting with an econometric model: The 'ragged-edge' problem", Journal of Forecasting, 5, 1-13.

[37] Wold, H. (1975), "Estimation of principal components and related models by iterative least squares," in Multivariate Analysis, ed. by P. Krishnaiaah, pp. 391-420. New York: Academic Press.

Table 1: In-sample fit of the static MF-3PRF and PCA factors

		MF-3PRF	PCA1	PCA2	PCA3
R^2	quarter average	**61.1**	59.8	60.5	60.8
	monthly: t, $t-1$, $t-2$	61.2	60.0	61.0	**62.5**
	growth rate of index	**56.5**	54.2	55.4	56.5
adjusted R^2	quarter average	**60.8**	59.5	59.8	59.8
	monthly: t, $t-1$, $t-2$	**60.2**	59.0	58.9	59.4
	growth rate of index	**56.1**	53.4	53.8	54.1

Note: bold numbers indicate the best-fitting model.

Table 2: Root mean squared and mean absolute forecasting errors

Root mean squared forecasting errors

		$h=1$	$h=2$	$h=3$	$h=4$	$h=5$
3PRF reference specifications	Kalman	1.88	2.15	2.60	2.80	2.91
	Shift	2.37*	2.35*	2.67	2.73	2.76
	AR	2.26*	2.39	2.57	2.86	2.83*
simple benchmarks	AR	2.57*	2.71	2.93*	2.92	3.00
	RW	3.00*	3.11*	3.63*	3.60*	3.93*
	CONST	2.90*	2.91*	2.91*	2.92	2.92
principal component benchmarks	PCA1	1.96	2.50*	2.97*	3.04*	3.04*
	PCA2	1.92	2.52*	2.88	2.94	3.05
	PCA3	2.05	2.51*	2.90*	2.92	2.95

Mean absolute forecasting errors

		$h=1$	$h=2$	$h=3$	$h=4$	$h=5$
3PRF reference specifications	Kalman	1.43	1.73	1.88	2.07	2.02
	Shift	1.82*	1.75	2.05	1.95*	1.95
	AR	1.74*	1.83	1.82	2.04	1.95
simple benchmarks	AR	1.84*	1.88	2.03	2.00	2.03
	RW	2.52*	2.24*	2.73*	2.79*	2.89*
	CONST	1.99*	2.00	2.00	2.00	2.01
principal component benchmarks	PCA1	1.48	2.01*	2.23*	2.21	2.15*
	PCA2	1.46	2.03*	2.25*	2.20	2.12
	PCA3	1.61*	2.03*	2.29*	2.26*	2.20

Note: * denotes a p-value smaller than 0.1 for the null hypothesis of equal forecasting performance based on Mariano-Diebold regressions comparing the 3PRF Kalman reference specification to the other forecasts. Kalman, Shift, and AR are three methods for coping with the ragged edge, as presented in Section 2.2.3.

Table 3: Overview of models considered in evaluation across countries (results are presented in Figure 8)

	Legend	Data	Ragged Edge	Factor extraction	GDP forecast
KRS3PRF		Country-specific	Kalman filter	3PRF	U-MIDAS, 5 lags
KRS3PRFpool		Combination of all country-specific datasets			
IR3PRF		Country-specific			Indirect forecast*
IR3PRFpool		Combination of all country-specific datasets			
PCA1		Country-specific		PCA	U-MIDAS, 5 lags
PCA2		Country-specific			
PCA3		Country-specific			
AR		GDP	-	-	AR(p), p chosen with AIC
RW		GDP	-	-	Forecast is latest observed value

* First, factor is forecasted using its law of motion. Second, this forecast is translated into a GDP forecast based on the contemporaneous correlation of GDP and the factor.

Table 4: Size of data sets, estimation and evaluation ranges by country

	Number of indicators	Start of estimation	Start of evaluation
USA	98	January 1984	Q1 2000
Euro area	77	February 1995	Q1 2003
Germany	27	January 1990	Q1 2000
France	24	January 1984	Q1 2000
Italy	31	February 1991	Q1 2003
Spain	33	February 1995	Q1 2007
United Kingdom	29	January 1984	Q1 2000
Japan	51	February 1994	Q4 2001
Brazil	66	January 1990	Q1 2007
India	101	February 1998	Q1 2008
South Korea	29	January 1984	Q2 2002
Taiwan	25	January 1984	Q1 2000
Global dataset	832	country-specific	country-specific

Note: in the case in which the global dataset is used to forecast GDP for a specific country, the beginning of the estimation and evaluation periods are the same as in the case of the country-specific dataset. The extraction date for all countries is 28 January 2014.

Table 5: Best models by country and forecasting horizon

	h = 1	h = 2	h = 3	h = 4	h = 5
US	**KRS3PRF**	**KRS3PRF**	**KRS3PRFpool**	**KRS3PRFpool**	**KRS3PRF**
EA	**IR3PRFpool**	**IR3PRFpool**	**KRS3PRFpool**	AR	*PCA1*
GER	**IR3PRF**	*PCA2*	*PCA1*	AR	AR
FR	**KRS3PRFpool**	**KRS3PRFpool**	*PCA2*	*PCA2*	*PCA2*
IT	**IR3PRFpool**	**KRS3PRFpool**	**KRS3PRFpool**	**KRS3PRFpool**	*PCA3*
ES	**IR3PRF**	**IR3PRF**	**IR3PRF**	**IR3PRF**	**IR3PRF**
UK	**KRS3PRFpool**	**KRS3PRFpool**	**KRS3PRFpool**	**KRS3PRFpool**	AR
JP	**IR3PRF**	*PCA2*	**KRS3PRFpool**	AR	**KRS3PRFpool**
BR	**KRS3PRF**	**KRS3PRFpool**	**KRS3PRF**	*PCA1*	**KRS3PRFpool**
IND	*PCA2*	*PCA1*	**KRS3PRFpool**	**IR3PRF**	**IR3PRF**
KO	**IR3PRF**	**KRS3PRF**	**KRS3PRFpool**	**IR3PRFpool**	**IR3PRFpool**
TW	*PCA2*	**IR3PRF**	**IR3PRF**	**KRS3PRFpool**	*PCA1*

This table presents, for each country evaluated and each forecasting horizon ($h = 1, 2, 3, 4, 5$), the model with the lowest RMSFE in the out-of-sample evaluation. Details on the models compared can be found in Table 3. Bold text denotes 3PRF models, italics refer to PCA models, and standard fonts to the simple AR benchmark.

**Table 6: RMSFE with indicator preselection relative to no selection**

	# of selected indicators	relative RMSFEs				
		$h=1$	$h=2$	$h=3$	$h=4$	$h=5$
USA	328	0.95*	0.94*	1.02	0.99	1.00
Euro area	473	1.01	1.00	1.05*	1.05*	1.05*
Germany	392	1.03	1.08	0.97	0.99	1.02
France	441	1.02	1.00	1.00	1.03*	1.03*
Italy	426	0.98	1.02	1.01	1.01*	1.01
Spain	349	0.89*	0.90*	0.96*	0.96*	0.94*
United Kingdom	341	0.94*	0.96	1.01	1.06	1.00
Japan	272	1.08*	1.00	0.99	1.04	0.98
Brazil	237	1.21*	1.14	1.06*	1.12*	1.03
India	57	1.13	1.11	1.27	1.04	1.01
South Korea	257	1.12*	1.09*	1.05	0.99	0.99
Taiwan	246	0.98	1.04	1.01	1.03	0.97

This table presents in the second column, the number of indicators from the global database that have a p-value smaller than 1% for the F-test for the Pass 1 regression. Columns 3 to 7 present the RMSFE of the MF-3PRF model using this preselected indicator dataset relative to the model using the global dataset containing all 832 indicators. * denotes a p-value smaller than 0.1 for the null hypothesis of equal forecasting performance based on Mariano-Diebold regression comparing the two models.

Table 7: Overview of modeling choices

Modeling dimension	Choices	Short explanation
Ragged Edge	Kalman filter	Use Kalman filter to extract information from all available data
	Realignment	Realign every series such that
	AR(2)	Fill missing observations with AR(2)
Factor extraction	3PRF	Three-pass regression filter
	PCA	Principal component
GDP forecast	U-MIDAS	Regress GDP growth on factor and its lags
	Indirect, restricted	1. Forecast factor using Kalman filter
		2. Calculate quarterly factor
		3. Use contemporaneous correlation of (forecasted) quarterly factor and GDP to get GDP forecast
	Indirect, unrestricted	1. Forecast factor using Kalman filter
		2. Use contemporaneous correlation of monthly factor and its lags and GDP to forecast GDP
Proxy	L-autoproxy	$L = 1$: Use GDP as proxy. $L = 2$: Use residuals from $L = 1$ as additional proxy. $L = 3$: Use residual from $L = 2$ as additional proxy. ...
Data	Country-specific dataset	Use data specific to the country
	Global dataset	Pool all country datasets to one large dataset
Selection	No selection	Use all series
	F-test in Pass 1	Run F-tests for the Pass 1 regressions and use only series with p-values below a threshold

Figure 1: A Monte Carlo experiments

Panel A: $\rho_g = 0.3$, $a = 0.9$, $d = 1$

Panel B: $\rho_g = 0.3$, $a = 0.9$, $d = 1$

Panel C: $\rho_g = 0.95$, $a = 0$, $d = 0$
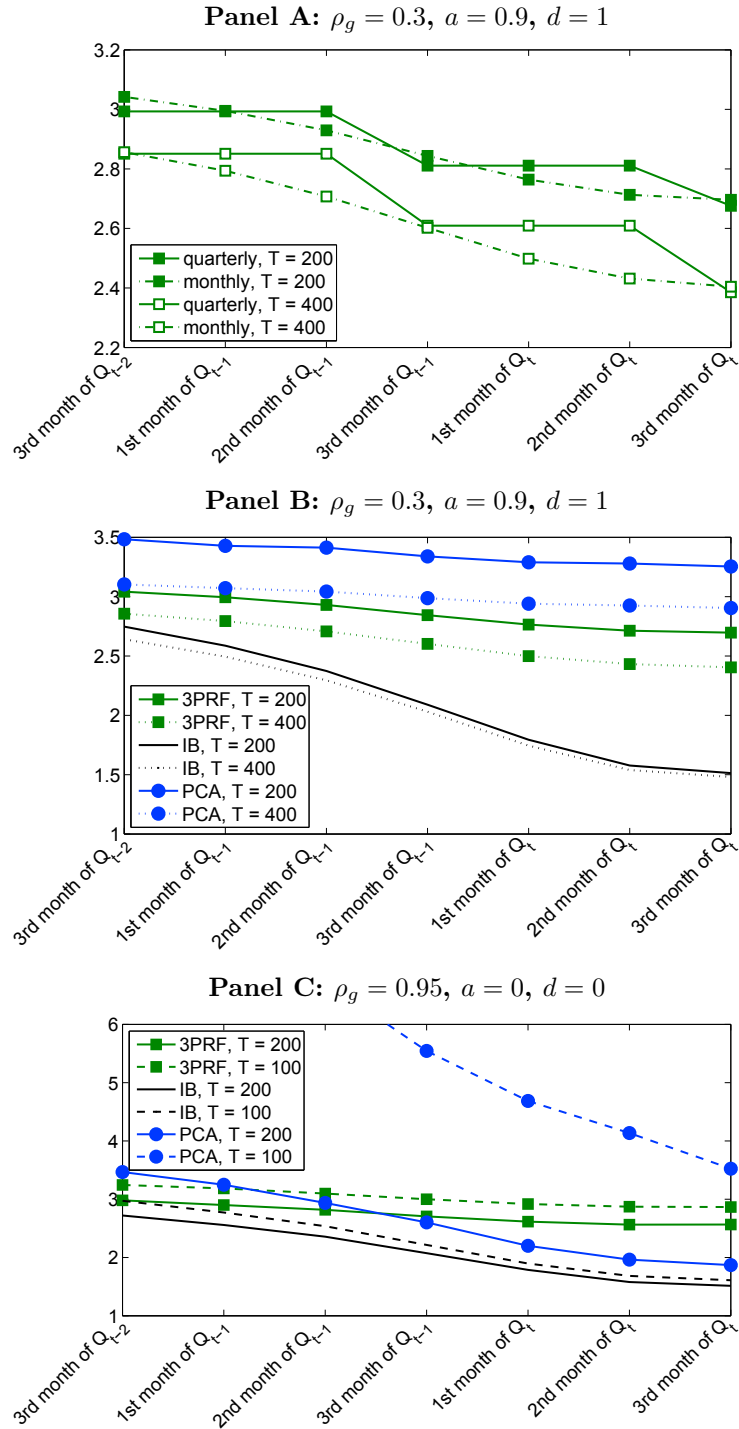
The figures plot on the x-axis the month up to (and including) which data are available and the RMSFE on the y-axis. Panel A presents the performance of the 3PRF and the MF-3PRF models. Panels B and C present the forecasting performance of the MF-3PRF and PCA models, together with the infeasible best forecast (IB) for two different DGPs and different sizes of the training samples.

Figure 2: The static 3PRF factor

Panel A
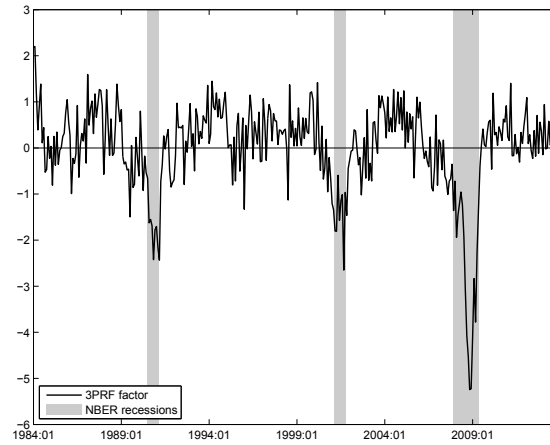
Panel B

Panel A presents the 3PRF factor for the US that is obtained if only cross-sectional data are used to extract the factor. Panel B compares, for the US, the 3PRF factor to the first three principal components. To extract the 3PRF and the PCA factors, only cross-sectional information is used.

Figure 3: Pseudo real-time evolution of the 3PRF factor
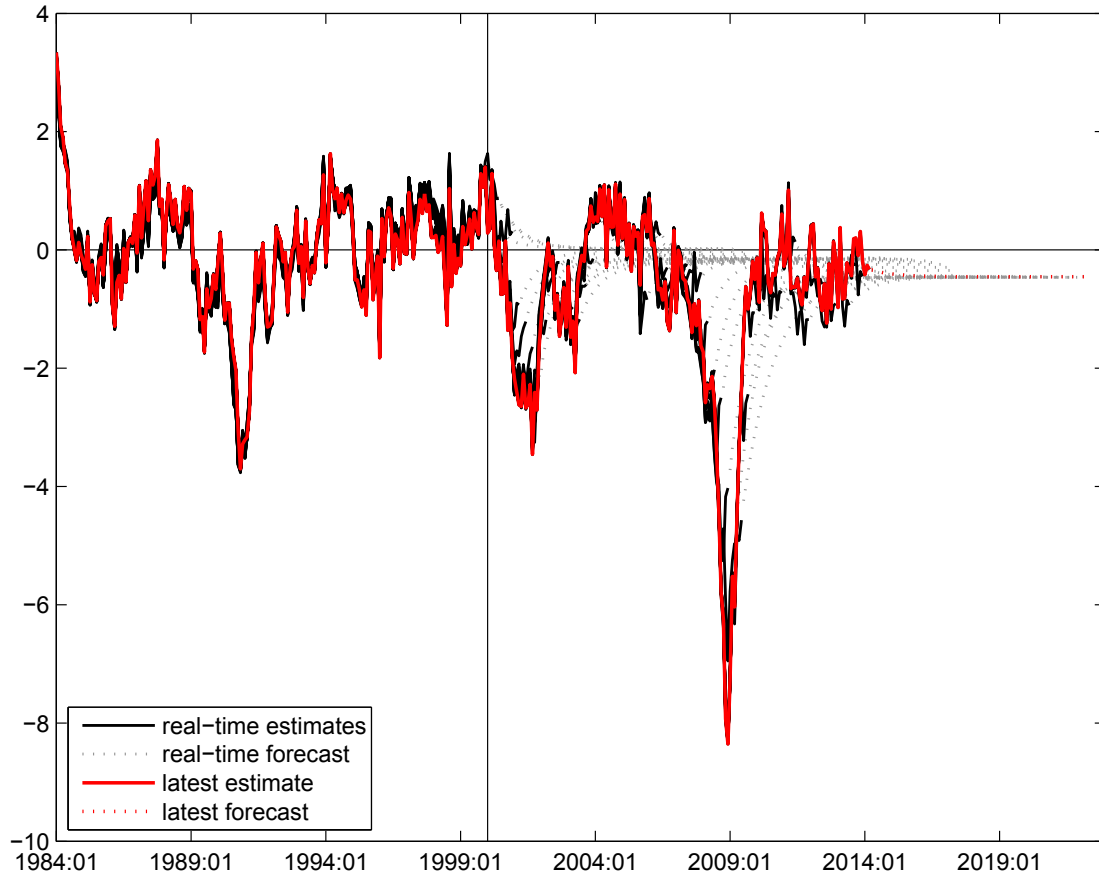
This figure presents the evolution of the pseudo real-time estimates of the 3PRF factor for the USA. The vertical black line denotes January 2000, when the out-of-sample exercise begins.

Figure 4: Recursive GDP forecasts for different forecasting horizons
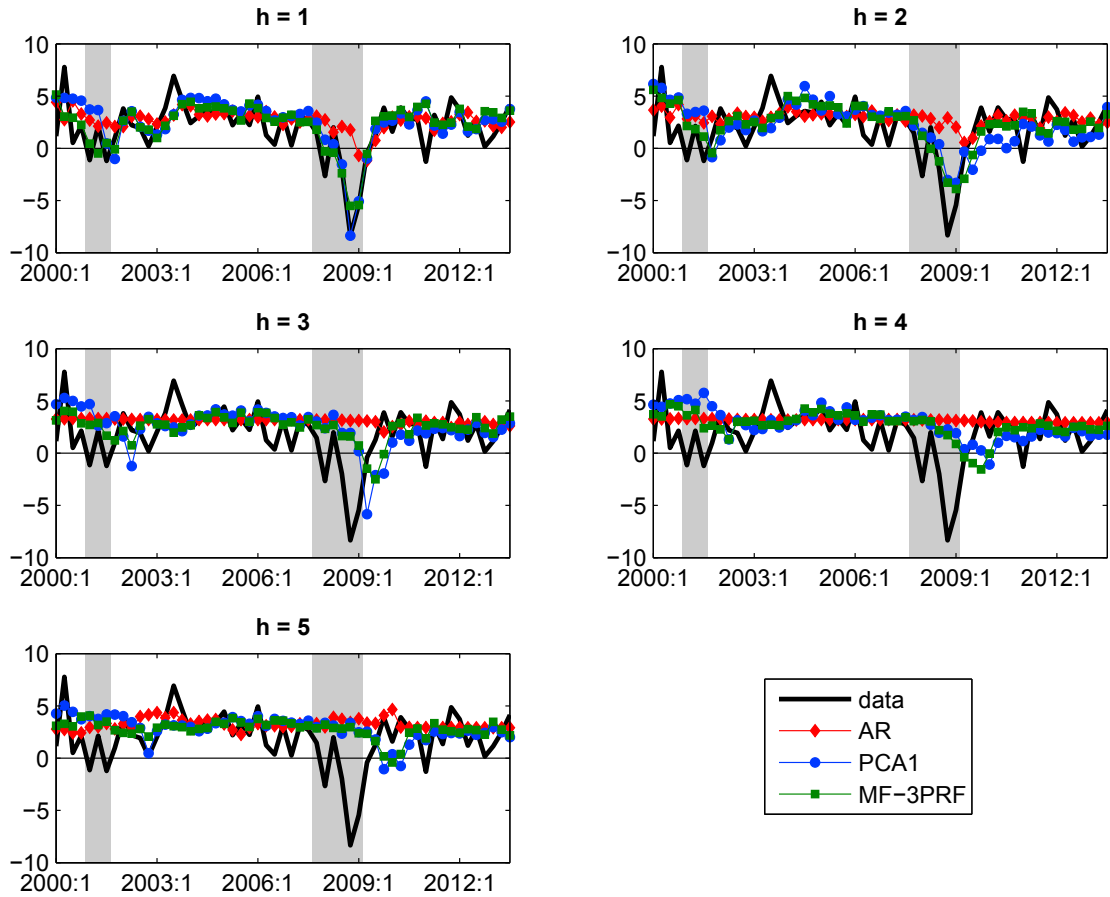
This figure plots the pseudo out-of-sample real-time forecasts for US GDP growth generated by three different models against actual GDP growth in black. The green line with square markers represents the forecasts generated by MF-3PRF, the blue line with the circle markers are the forecasts of the principal-components-based PCA benchmark with one factor, and the red line with the diamond markers represents the forecasts generated by the AR benchmark. The panels represent the different forecasting horizons: $h = 1$ is the backcast, e.g., a forecast made during Q1 2014 for GDP growth in Q4 2013. $h = 2$ is the nowcast, e.g., a forecast for growth in Q1 2014 made in Q1 2014, and $h = 3, 4, 5$ the short-term forecasts, e.g., forecasts for Q2, Q3, and Q4 2014 made in Q1 2014.

**Figure 5: 4-year rolling RMSEs for different forecasting horizons**
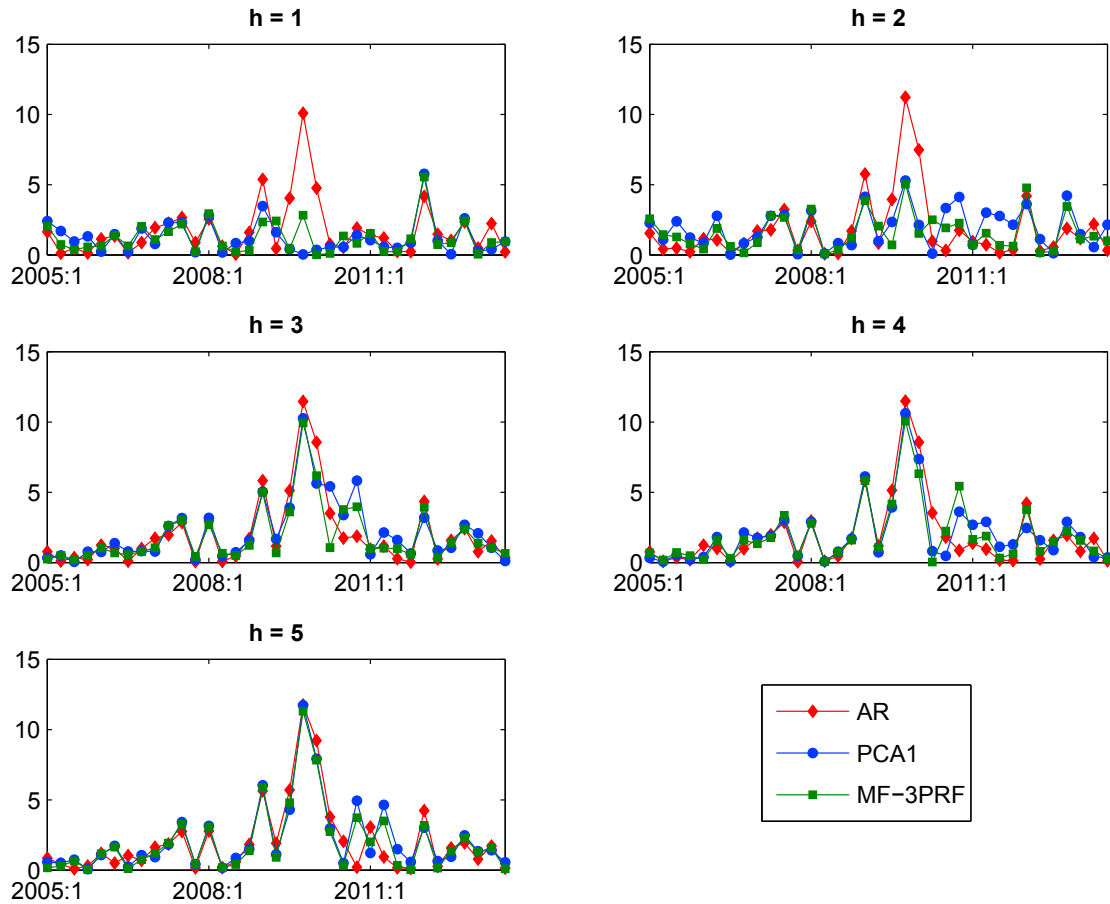
This figure plots the 4-year rolling RMSEs for the pseudo out-of-sample real-time forecasts for US GDP growth generated by three different models. The green line with square markers represents the forecasts generated by MF-3PRF, the blue line with the circle markers are the forecasts of the principal-components-based PCA benchmark with one factor, and the red line with the diamond markers represents the forecasts generated by the AR benchmark. The panel represent the different forecasting horizons: $h = 1$ is the backcast, e.g., a forecast made during Q1 2014 for GDP growth in Q4 2013, $h = 2$ is the nowcast, e.g., a forecast for growth in Q1 2014 made in Q1 2014, and $h = 3, 4, 5$ the short-term forecasts, e.g., forecasts for Q2, Q3, and Q4 2014 made in Q1 2014.
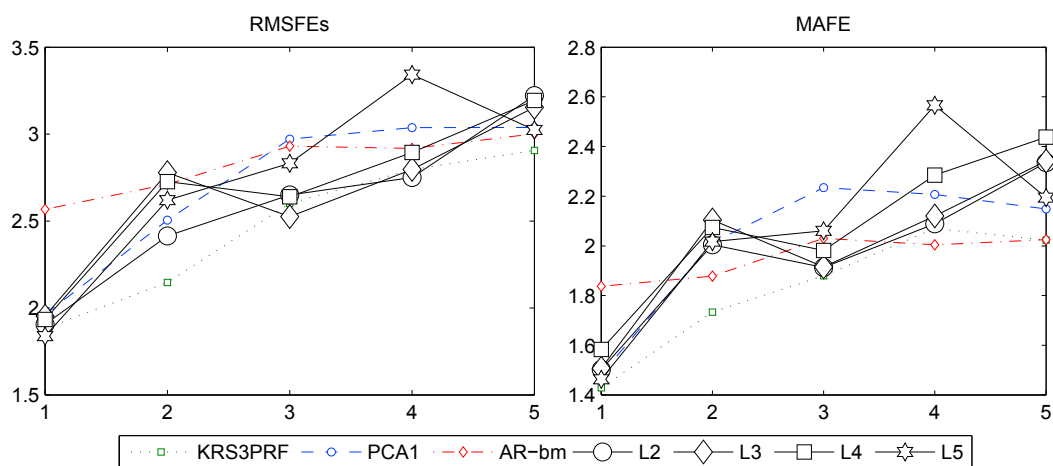
Figure 6: Robustness analysis

Panel A: Proxies specific to the forecasting horizon
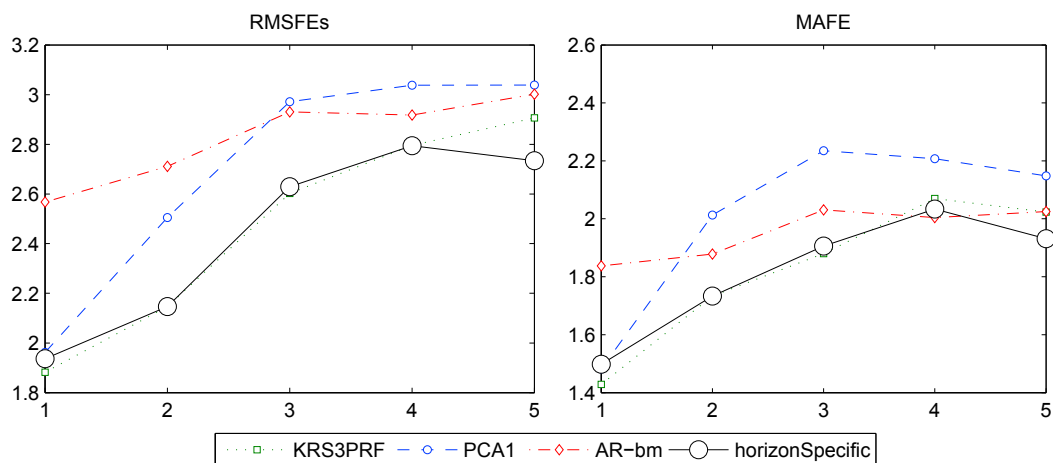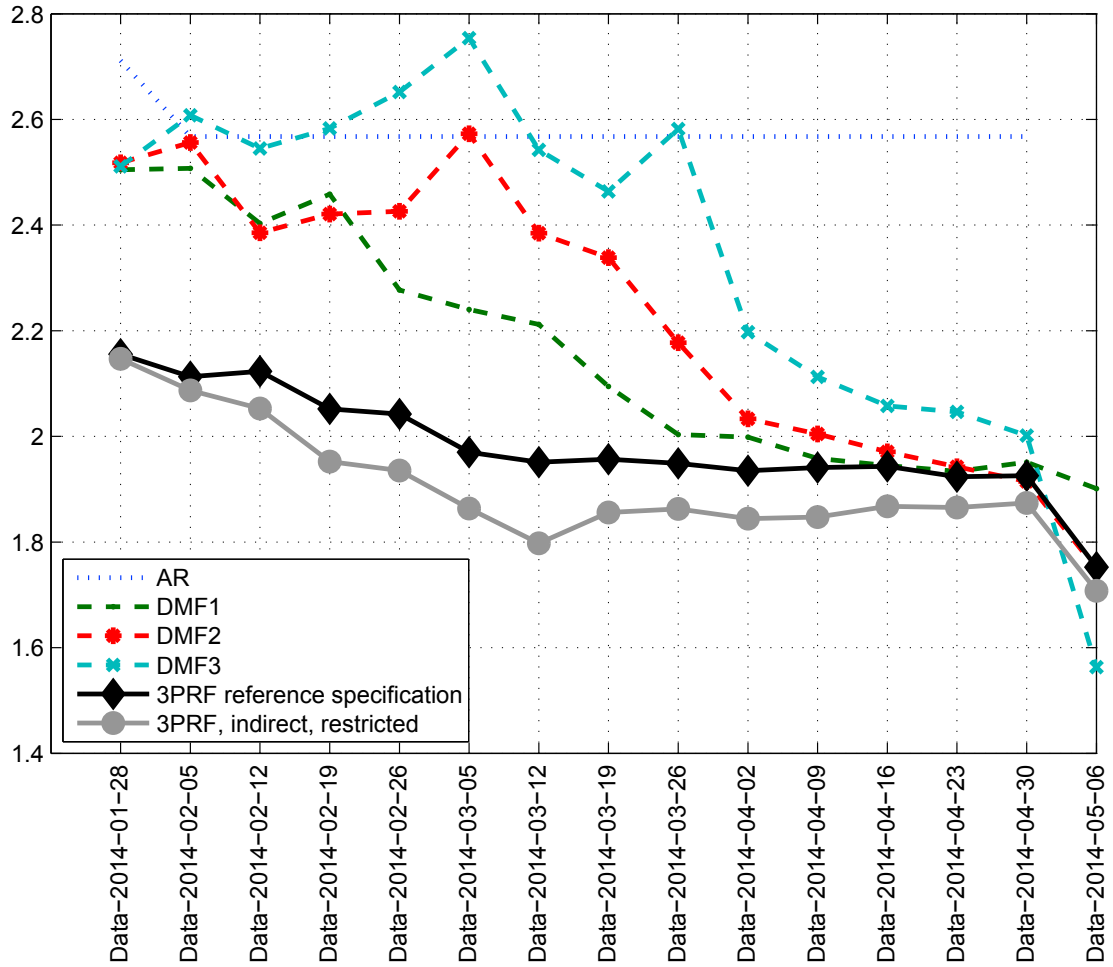

Panel B: Indirect forecasts
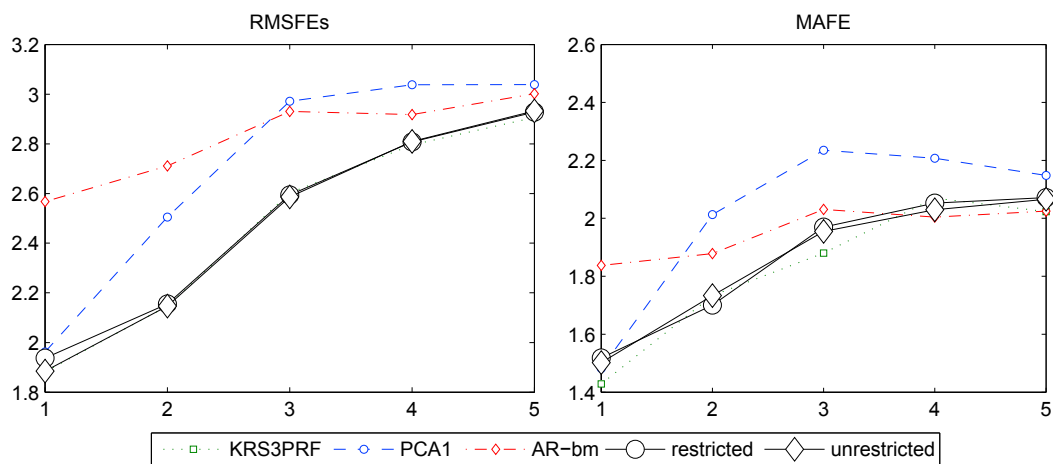
Panel C: L-autoproxy with $L = 1, ..., 5$

This figure presents the forecasting performance of various alternative specifications of the MF-3PRF model.

Figure 7: Temporal evolution of the nowcast

This figure presents the temporal evolution of the nowcasting performance for US GDP of various models as additional data arrive. The dataset is extended in weekly steps beginning with the data structure available on 28 January 2014 and ending with the structure available on 6 May 2014.

Figure 8: RMSFEs for different countries

This figure presents the RMSFEs of the pseudo out-of-sample analysis for different countries. Details on the models compared can be found in Table 3. The data are described in Table 4.

APPENDIX: additional material

Figure 9: 3PRF factor based on pooled dataset vs country-specific dataset
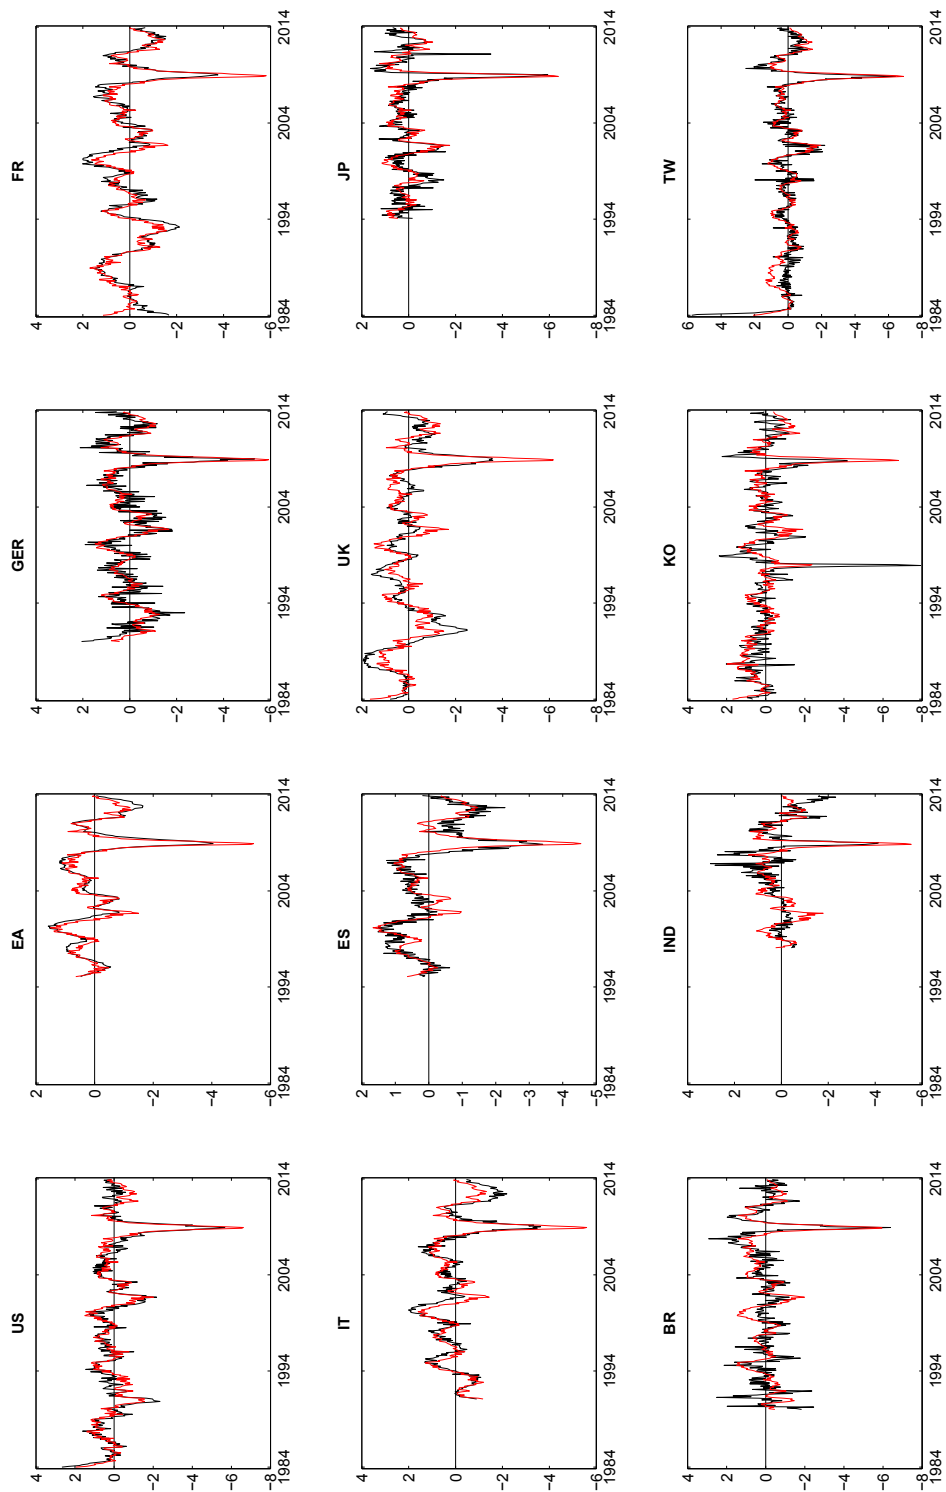
This figure presents the Kalman-3PRF factors that are obtained when using the country-specific database (black) and the global database (red).
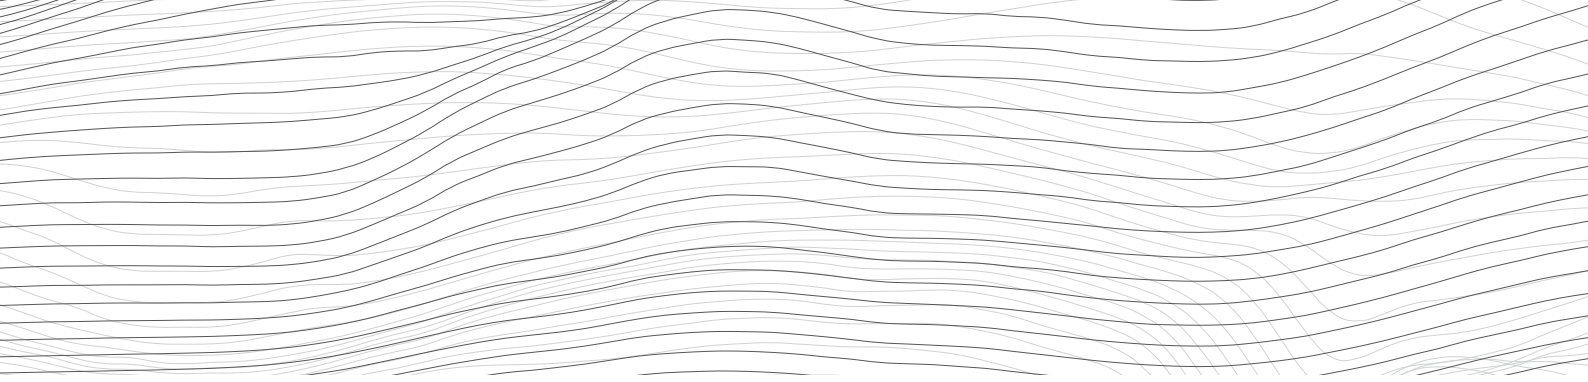
Recent SNB Working Papers

2016-4   Christian Hepenstrick and Massimiliano Marcellino: Forecasting with Large Unbalanced Datasets: The Mixed-Frequency Three-Pass Regression Filter.

2016-3   Alain Galli: How reliable are cointegration-based estimates for wealth effects on consumption? Evidence from Switzerland.

2016-2   Pinar Yeşin: Exchange Rate Predictability and State-of-the-Art Models.

2016-1   Sandra Hanslin and Rolf Scheufele: Foreign PMIs: A reliable indicator for exports?

2015-13   Thomas Nellen: Collateralised liquidity, two-part tariff and settlement coordination.

2015-12   Jacob Gyntelberg, Mico Loretan and Tientip Subhanij: Private information, capital flows, and exchange rates.

2015-11   Philip Sauré: Time-intensive R&D and unbalanced trade.

2015-10   Nikola Mirkov and Andreas Steinhauer: Ben Bernanke vs. Janet Yellen: Exploring the (a)symmetry of individual and aggregate inflation expectations.

2015-9   Aleksander Berentsen, Sébastien Kraenzlin and Benjamin Müller: Exit Strategies and Trade Dynamics in Repo Markets.

2015-8   Thomas Nitschka: Is there a too-big-to-fail discount in excess returns on German banks' stocks?

2015-7   Alin Marius Andries, Andreas M. Fischer and Pinar Yeşin: The impact of international swap lines on stock returns of banks in emerging markets.

2015-6   Jens H.E. Christensen and Signe Krogstrup: Transmission of Quantitative Easing: The Role of Central Bank Reserves.

2015-5   Petra Gerlach-Kristen and Seán Lyons: Mortgage arrears in Europe: The impact of monetary and macroprudential policies.

2015-4   Reto Foellmi, Sandra Hanslin and Andreas Kohler: A dynamic North-South model of demand-induced product cycles.

2015-3   Katarina Juselius and Katrin Assenmacher: Real exchange rate persistence: The case of the Swiss franc-US dollar rate.

2015-2   Lucas Marc Fuhrer, Basil Guggenheim and Silvio Schumacher: Re-use of collateral in the repo market.

2015-1   Pinar Yeşin: Capital flow waves to and from Switzerland before and after the financial crisis.

2014-13   Thomas Nitschka: Have investors been looking for exposure to specific countries since the global financial crisis? – Insights from the Swiss franc bond market.

2014-12   Christian Grisse and Thomas Nitschka: Exchange rate returns and external adjustment: evidence from Switzerland.

2014-11   Rina Rosenblatt-Wisch and Rolf Scheufele: Quantification and characteristics of household inflation expectations in Switzerland.

SCHWEIZERISCHE NATIONALBANK
BANQUE NATIONALE SUISSE
BANCA NAZIONALE SVIZZERA
BANCA NAZIUNALA SVIZRA
SWISS NATIONAL BANK   ⊹