

2013 | 06

Working Paper

Norges Bank Research

A survey of econometric methods for mixed-frequency data

Claudia Foroni and Massimiliano Marcellino

Working papers fra Norges Bank, fra 1992/1 til 2009/2 kan bestilles over e-post:

servicesenter@norges-bank.no

Fra 1999 og senere er publikasjonene tilgjengelige på www.norges-bank.no

Working papers inneholder forskningsarbeider og utredninger som vanligvis ikke har fått sin endelige form. Hensikten er blant annet at forfatteren kan motta kommentarer fra kolleger og andre interesserte. Synspunkter og konklusjoner i arbeidene står for forfatterens regning.

Working papers from Norges Bank, from 1992/1 to 2009/2 can be ordered by e-mail:

servicesenter@norges-bank.no

Working papers from 1999 onwards are available on www.norges-bank.no

Norges Bank's working papers present research projects and reports (not usually in their final form) and are intended inter alia to enable the author to benefit from the comments of colleagues and other interested parties. Views and conclusions expressed in working papers are the responsibility of the authors alone.

A Survey of Econometric Methods for Mixed-Frequency Data*

Claudia Foroni
Norges Bank

Massimiliano Marcellino
European University Institute,
Bocconi University and CEPR

This version: February 6, 2013

Abstract

The development of models for variables sampled at different frequencies has attracted substantial interest in the recent econometric literature. In this paper we provide an overview of the most common techniques, including bridge equations, MIXed DATA Sampling (MIDAS) models, mixed frequency VARs, and mixed frequency factor models. We also consider alternative techniques for handling the ragged edge of the data, due to asynchronous publication. Finally, we survey the main empirical applications based on alternative mixed frequency models.

J.E.L. Classification: E37, C53

Keywords: mixed-frequency data, mixed-frequency VAR, MIDAS, nowcasting, forecasting

*We would like to thank Tommaso Di Fonzo, Eric Ghysels, Helmut Lutkepohl for useful comments on a previous version. The views expressed herein are solely those of the authors and do not necessarily reflect the views of the Norges Bank. The usual disclaimers apply.

1 Introduction

In recent times, econometric models that take into account the information in unbalanced datasets have attracted substantial attention. Policy-makers, in particular, need to assess in real-time the current state of the economy and its expected developments, when only incomplete information is available.

In real-time, the unbalancedness of datasets arises mainly due to two features: the different sampling frequency with which the indicators are available and the so-called "ragged-edge" problem, namely, publication delays cause missing values for some of the variables at the end of the sample, see Wallis (1986). As an example, one of the key indicators of macroeconomic activity, the Gross Domestic Product (GDP), is released quarterly and with a considerable publication lag, while a range of leading and coincident indicators is available more timely and at a monthly or even higher frequency.

In this paper we review the methods proposed so far in the literature to deal with mixed-frequency data and missing values due to publication lags, see Banbura, Giannone, Reichlin (2011) and Banbura, Giannone, Modugno, Reichlin (2012) for complementary overviews with a stronger focus on Kalman filter based factor modeling techniques.

The simplest approach is to aggregate the data to obtain a balanced dataset at the same frequency and to work with a "frozen" final vintage dataset, which eliminates the ragged edge problem. However, in the literature there are also a few methods to avoid pre-filtering associated with temporally aggregated or interpolated data, and to exploit the information contained in the large number of series available in real-time at different frequencies. In what follows, we depict the main features of the bridge models, often employed in central banks and other policy making institutions, especially for nowcasting and short-term forecasting, see e.g. Baffigi, Golinelli and Parigi (2004), Diron (2008) and Bencivelli, Marcellino and Moretti (2012). We then move to one of the main strands of the literature, mixed-data sampling (MIDAS) models, parsimonious specifications based on distributed lag polynomials, which flexibly deal with data sampled at different frequencies and provide a direct forecast of the low-frequency variable (see e.g. Ghysels et al. (2004), Clements and Galvao (2008)). Finally, we consider the state-space approaches, presenting mixed-frequency VAR (MF-VAR) and factor models. Both are system approaches that jointly describe the dynamics of the variable to be explained and of the indicators, where the use of the Kalman filter provides not only predictions of the future observations but also estimates of the current latent state (see Mariano and Murasawa (2003, 2010)). A natural extension in the literature is the combination of the factors with the MIDAS models, and it is based on the use of factors as explanatory variables to exploit the information in large mixed-frequency datasets. The resulting model is labelled Factor-MIDAS by Marcellino and Schumacher (2010).

For each of the alternative approaches to mixed frequency modelling listed above, we first describe their key theoretical features, and then summarize empirical applications.

The paper is organized as follows. In Section 2, we survey the different approaches

to model mixed frequency variables. In Section 3, we discuss the additional estimation issues arising with a ragged-edge structure of the dataset. In Section 4 we compare the main features of the different approaches. In Section 5 we present a summary of the most significant empirical applications in this literature. Finally, in Section 6 we summarize and conclude.

2 Models for mixed-frequency data

Typical regression models relate variables sampled at the same frequency. To ensure the same frequency, researchers working with time series data either aggregate the higher-frequency observations to the lowest available frequency or interpolate the lower-frequency data to the highest available frequency, see Section 2.1. The most common solution in empirical applications is the former, temporal aggregation. The higher-frequency data are aggregated to the lowest-frequency by averaging or by taking a representative value (for example, the last month of a quarter). In pre-filtering the data so that left- and right-hand variables are available at the same frequency, a lot of potentially useful information might be destroyed, and mis-specification inserted in the model. Hence, direct modelling of mixed frequency data can be useful¹.

One of the early approaches to deal with mixed-frequency data focuses on forecasting and relies on bridge equations, see e.g. Baffigi, Golinelli, Parigi (2004), i.e. equations that link the low-frequency variables and time-aggregated indicators. Forecasts of the high-frequency indicators are provided by specific high-frequency time series models, then the forecast values are aggregated and plugged into the bridge equations to obtain the forecast of the low-frequency variable. Details are provided in Section 2.2.

In Section 2.3 we propose a more detailed overview of one of the most recent and competitive univariate approaches, the mixed-data sampling method originally proposed by Ghysels, Santa-Clara and Valkanov (2004). Mixed-data sampling (MIDAS) models handle series sampled at different frequencies, where distributed lag polynomials are used to ensure parsimonious specifications. Whereas early MIDAS studies focused on financial applications, see e.g. Ghysels, Santa-Clara and Valkanov (2006), recently this method has been employed to forecast macroeconomic time series, where typically quarterly GDP growth is forecasted by monthly macroeconomic and financial indicators, see e.g. Clements and Galvao (2008, 2009).

Another common approach in the literature is the state-space representation of the model, where to handle data with different frequencies, the low-frequency variable is

¹Wolhrabe (2009) presents another review of mixed-frequency models. However, his review focuses more on the earliest attempts to tackle the mixed-frequency issues. In particular he reviews in detail the aggregation and interpolation of data, and the bridge and linkage models. In this paper, instead, we focus on the most recent developments of the literature, and especially we go through the MIDAS approach and its recent extensions, the MF-VAR in a classical and Bayesian framework, and different factor models which take into account the mixed-frequency and ragged-edge nature of the dataset.

considered as a high-frequency one with missing observations. The Kalman filter and smoother is then applied to estimate the missing observations and to generate forecasts. Moreover, the dynamics of the low and high-frequency series are jointly analyzed. One of the most compelling approaches at the moment is the one proposed by Zadrozny (1988) for directly estimating a VARMA model with series sampled at different frequencies. In the same fashion, Mariano and Murasawa (2010) set what they call mixed-frequency VAR (MF-VAR from now on), i.e. they introduce a VAR model for partially latent time series and cast it in state-space form, see Section 2.4 for more details. Among the state-space approaches we can also list mixed-frequency factor models employed, for example, to extract an unobserved state of the economy and create a new coincident indicator or forecast and nowcast GDP, see e.g. Mariano and Murasawa (2003, 2010) in Section 2.5.1 for small scale applications and Giannone, Reichlin and Small (2008) and Banbura and Rünstler (2011) for large scale models in Section 2.5.2 and 2.5.3. A similar approach is also followed by Frale et al. (2010, 2011): differently from the other studies, dynamic factor models are applied to a set of small datasets where variables are grouped according to economic theory and institutional considerations, rather than to the entire information set. The separate small factor models are then linked together within a state-space framework. Finally, in Section 2.5.4 we review the literature that proposes to merge the two recent strands in the mixed sampling econometrics: factor models and MIDAS approach. Marcellino and Schumacher (2010) introduce Factor-MIDAS, an approach for now- and forecasting low-frequency variables exploiting information in large sets of higher-frequency indicators.

2.1 Aggregation and interpolation

In most of the empirical applications, the common solution in the presence of a mixed sample frequency is to pre-filter the data so that the left- and right-hand side variables are sampled at the same frequency. In the process, a lot of potentially useful information can be destroyed and mis-specification included in the model.

The standard aggregation methods depend on the stock/flow nature of the variables and, typically, it is the average of the high-frequency variables over one low-frequency period for stocks, and the sum for flows.

Taking the latest available value of the higher frequency variable is another option for both stock and flow variables. The underlying assumption is that the information of the previous high-frequency periods is reflected in the latest value, representative of the whole low-frequency period.

The second option to match frequencies is the interpolation of the low frequency variables, which is rarely used. There are several different interpolating methods, see e.g. Lanning (1986), Marcellino (1998) and Angelini et al. (2006). A common approach is a two-step procedure: first missing data are interpolated, then model parameters are estimated using the new augmented series, possibly taking into account the measurement

error induced by disaggregation. Both steps can be conveniently and jointly run in a Kalman filter set-up, starting with a state-space representation of the model, see e.g. Harvey (1989) and Sections 2.4 and 2.5 below.

2.2 Bridge equations

One of the early econometric approaches in the presence of mixed-frequency data relies on the use of bridge equations, see e.g. Baffigi, Golinelli, Parigi (2004) and Diron (2008). Bridge equations are linear regressions that link ("bridge") high frequency variables, such as industrial production or retail sales, to low frequency ones, e.g. the quarterly real GDP growth, providing some estimates of current and short-term developments in advance of the release. The "Bridge model" technique allows computing early estimates of the low-frequency variables by using high frequency indicators. They are not standard macroeconomic models, since the inclusion of specific indicators is not based on causal relations, but on the statistical fact that they contain timely updated information. In principle, bridge models require that the whole set of regressors should be known over the projection period, allowing for an estimate only of the current period. In practice, anyway, this is not the case, even though the forecasting horizon of the bridge models is quite short, one or two quarters ahead at most.

Taking forecasting GDP as an example, since the monthly indicators are usually only partially available over the projection period, the predictions of quarterly GDP growth are obtained in two steps. First, monthly indicators are forecasted over the remainder of the quarter, usually on the basis of univariate time series models (in some cases VAR have been implemented in order to obtain better forecasts of the monthly indicators), and then aggregated to obtain their quarterly correspondent values. Second, the aggregated values are used as regressors in the bridge equation which allows to obtain forecasts of GDP growth.

Therefore, the bridge model to be estimated is:

$$y_{t_q} = \alpha + \sum_{i=1}^j \beta_i(L) x_{it_q} + u_{t_q} \quad (1)$$

where $\beta_i(L)$ is a lag polynomial of length k , and x_{it_q} are the selected monthly indicators aggregated at quarterly frequency.

The selection of the monthly indicators included in the bridge model is usually based on a general-to-specific methodology and relies on different in-sample or out-of-sample criteria, like information criteria or RMSE performance. Bencivelli, Marcellino and Moretti (2012) propose an alternative procedure based on Bayesian Model Averaging (BMA) that performs quite well empirically.

In order to forecast the missing observations of the monthly indicators which are then aggregated to obtain a quarterly value of x_{it_q} , it is common practice to use autoregressive

models, where the lag length is based on information criteria.

2.3 Mixed-Data Sampling

Distributed lag (DL) models have been typically employed in the literature to describe the distribution over time of the lagged effects of a change in the explanatory variable. In general, a stylized distributed lag model is given by

$$y_{t_q} = \alpha + B(L) x_{t_q} + \varepsilon_{t_q} \quad (2)$$

where $B(L)$ is some finite or infinite lag polynomial operator.

This kind of models underlies the construction of the bridge equations, once all the high frequency values are aggregated to the corresponding low-frequency values.

In order to take into account mixed-frequency data, Ghysels et al. (2004) introduce the Mixed-Data Sampling (MIDAS) approach, which is closely related to the distributed lag model, but in this case the dependent variable y_{t_q} , sampled at a lower-frequency, is regressed on a distributed lag of x_{t_m} , which is sampled at a higher-frequency.

In what follows, we first present the basic features of the model as presented by Ghysels et al. (2004), the corresponding unrestricted version as in Foroni, Marcellino and Schumacher (2012), and then the extensions that have been introduced in the literature.

In terms of notation, $t_q = 1, \dots, T_q$ indexes the basic time unit (e.g. quarters), and m is the number of times the higher sampling frequency appears in the same basic time unit. For example, for quarterly GDP growth and monthly indicators as explanatory variables, $m = 3$. w is the number of monthly values of the indicators that are earlier available than the lower-frequency variable to be estimated. The lower-frequency variable can be expressed at the high frequency by setting $y_{t_m} = y_{t_q}, \forall t_m = mt_q$, where t_m is the time index at the high frequency.

2.3.1 The basic MIDAS model

MIDAS regressions are essentially tightly parameterized, reduced form regressions that involve processes sampled at different frequencies. The response to the higher-frequency explanatory variable is modelled using highly parsimonious distributed lag polynomials, to prevent the proliferation of parameters that might otherwise result, as well as the issues related to lag-order selection.

The basic MIDAS model for a single explanatory variable, and h_q -step-ahead forecasting, with $h_q = h_m/m$, is given by:

$$y_{t_q+mh_q} = y_{t_m+h_m} = \beta_0 + \beta_1 b(L_m; \theta) x_{t_m+w}^{(m)} + \varepsilon_{t_m+h_m} \quad (3)$$

where $b(L^{1/m}; \theta) = \sum_{k=0}^K c(k; \theta) L_m^k$, and $L_m^x x_{t_m}^{(m)} = x_{t_m-x}^{(m)} \cdot x_{t_m+w}^{(m)}$ is skip-sampled from the

high frequency indicator x_{t_m} .

The parameterization of the lagged coefficients of $c(k; \theta)$ in a parsimonious way is one of the key MIDAS features. One of the most used parameterizations is the one known as “Exponential Almon Lag”, since it is closely related to the smooth polynomial Almon lag functions that are used to reduce multicollinearity in the Distributed Lag literature. It is often expressed as

$$c(k; \theta) = \frac{\exp(\theta_1 k + \dots + \theta_Q k^Q)}{\sum_{k=1}^K \exp(\theta_1 k + \dots + \theta_Q k^Q)} \quad (4)$$

This function is known to be quite flexible and can take various shapes with only a few parameters. These include decreasing, increasing or hump-shaped patterns. Ghysels, Santa-Clara and Valkanov (2005) use the functional form with two parameters, which allows a great flexibility and determines how many lags are included in the regression.

Notice that the standard practice in bridge equations of calculating a quarterly series from the monthly indicators corresponds to imposing restrictions on this parameterization function. To be concrete, in the case of the quarterly-monthly example, taking the last month in the quarter to produce a quarterly series amounts to setting $c(2; \theta) = c(3; \theta) = c(5; \theta) = c(6; \theta) = \dots = c(11; \theta) = c(12; \theta) = 0$.

Another possible parameterization, also with only two parameters, is the so-called “Beta Lag”, because it is based on the Beta function:

$$c(k; \theta_1, \theta_2) = \frac{f\left(\frac{k}{K}, \theta_1; \theta_2\right)}{\sum_{k=1}^K f\left(\frac{k}{K}, \theta_1; \theta_2\right)} \quad (5)$$

where $c(x, a, b) = \frac{x^{a-1}(1-x)^{b-1}\Gamma(a+b)}{\Gamma(a)\Gamma(b)}$ and $\Gamma(a) = \int_0^\infty e^{-x} x^{a-1} dx$.

Ghysels, Rubia and Valkanov (2009) propose also three other different parameterizations of the lag coefficients: a linear scheme, with $c(k; \theta) = \frac{1}{K}$, where there are no parameters to estimate in the lagged weight function; an hyperbolic scheme, with $c(k; \theta) = \frac{g\left(\frac{k}{K}, \theta\right)}{\sum_{k=1}^K g\left(\frac{k}{K}, \theta\right)}$, $g(k, \theta) = \frac{\Gamma(k+\theta)}{\Gamma(k+1)\Gamma(\theta)}$ where the gamma function has only one parameter to estimate, but it’s not as flexible as the Beta specification; a geometric scheme, with $c(k; \theta) = \frac{\theta^k}{\sum_{k=1}^\infty \theta^k}$, $|\theta| \leq 1$ and $c(k; \theta)$ are normalized so that they sum up to one.

The parameterizations described above are all quite flexible. For different values of the parameters, they can take various shapes: weights attached to the different lags can decline slowly or fast, or even have a hump shape. Therefore, estimating the parameters from the data automatically determines the shape of the weights and, accordingly, the number of lags to be included in the regression.

The MIDAS model can be estimated using nonlinear least squares (NLS) in a regression of y_t onto $x_{t-h}^{(m)}$. Ghysels, Santa-Clara and Valkanov (2004) show that MIDAS regressions

always lead to more efficient estimation than the typical approach of aggregating all series to the least frequent sampling. Moreover, they also show that discretization biases are the same for MIDAS and distributed lag models and vanish when regressors are sampled more frequently.

The forecast is given by

$$\widehat{y}_{T_m^y+h_m|T_m^x} = \widehat{\beta}_0 + \widehat{\beta}_1 b(L_m; \widehat{\theta}) x_{T_m^x}^{(m)}. \quad (6)$$

Note that MIDAS is h -dependent, and thus needs to be re-estimated for each forecast horizon.

2.3.2 The AR-MIDAS model

Since autoregressive models often provide competitive forecasts to those obtained with models that include explanatory variables, the introduction of an autoregressive term in the MIDAS model is a desirable extension, although not straightforward. Ghysels, Santa-Clara and Valkanov (2004) show that the introduction of lagged dependent variables creates efficiency losses. Moreover, it would result in the creation of seasonal patterns in the explanatory variables.

Consider adding a lower-frequency lag of y_{t_m} , y_{t_m-3} , to the basic model with $m = 3$ (x is monthly and y is quarterly):

$$y_{t_m} = \beta_0 + \lambda y_{t_m-3} + \beta_1 b(L_m; \theta) x_{t_m+w-3}^{(3)} + \varepsilon_{t_m}. \quad (7)$$

As highlighted in Clements and Galvao (2009), this strategy is in general not appropriate. The reason becomes clear when we write the model as:

$$y_{t_m} = \beta_0 (1 - \lambda)^{-1} + \beta_1 (1 - \lambda L_m^3)^{-1} B(L_m; \theta) x_{t_m+w-3}^{(3)} + (1 - \lambda L_m^3)^{-1} \varepsilon_{t_m}. \quad (8)$$

The polynomial on $x_{t-1}^{(3)}$ is a product of a polynomial in $L^{1/3}$ and a polynomial in L . This product generates a seasonal response of y to $x^{(3)}$, irrespective of whether $x^{(3)}$ displays a seasonal pattern.

To avoid this inconvenience, the authors suggest the introduction of the AR dynamics as a common factor:

$$y_{t_m} = \beta_0 + \lambda y_{t_m-3} + \beta_1 b(L_m; \theta) (1 - \lambda L_m^3) x_{t_m+w-3}^{(3)} + \varepsilon_{t_m} \quad (9)$$

so that the response of y to $x^{(3)}$ remains non-seasonal.

The analogous multi-step model is written as:

$$y_{t_m} = \beta_0 + \lambda y_{t_m-h_m} + \beta_1 b(L_m; \theta) (1 - \lambda L_m^{h_m}) x_{t_m+w-h_m}^{(3)} + \varepsilon_{t_m}. \quad (10)$$

To estimate the MIDAS-AR model, the common procedure is to estimate the standard

MIDAS (the basic model), take the residuals $\widehat{\varepsilon}_{t_m}$ and estimate an initial value for λ , say λ_0 , where $\widehat{\lambda}_0 = \left(\sum \widehat{\varepsilon}_{t_m+w-h_m}^2\right)^{-1} \sum \widehat{\varepsilon}_{t_m} \widehat{\varepsilon}_{t_m+w-h_m}$. Then construct $y_{t_m}^* = y_{t_m} - \widehat{\lambda}_0 y_{t_m-h_m}$ and $x_{t_m+w-h_m}^{*(3)} = x_{t_m+w-h_m}^{(3)} - \widehat{\lambda}_0 x_{t_m-(h_m-w)-h_m}^{(3)}$. The estimator $\widehat{\theta}_1$ is obtained by applying nonlinear least squares to:

$$y_{t_m}^* = \beta_0 + \beta_1 b(L_m; \theta) x_{t_m+w-h_m}^{*(3)} + \varepsilon_{t_m}. \quad (11)$$

A new value of λ , $\widehat{\lambda}_1$, is obtained from the residuals of this regression. Then a new step is run, using $\widehat{\lambda}_1$ and $\widehat{\theta}_1$ as the initial values. In this way, the procedure gets the estimates and $\widehat{\lambda}$ and $\widehat{\theta}$ that minimize the sum of squared residuals.

2.3.3 The Unrestricted MIDAS model

Foroni, Marcellino and Schumacher (2012) study the performance of a variant of MIDAS which does not resort to functional distributed lag polynomials. In the paper, the authors discuss how unrestricted MIDAS (U-MIDAS) regressions can be derived in a general linear dynamic framework, and under which conditions the parameters of the underlying high-frequency model can be identified².

The U-MIDAS model based on a linear lag polynomial such as

$$\begin{aligned} c(L^m)\omega(L)y_{t_m} &= \delta_1(L)x_{1t_{m-1}} + \dots + \delta_N(L)x_{Nt_{m-1}} + \epsilon_{t_m}, \\ t &= 1, 2, 3, \dots \end{aligned} \quad (12)$$

where $c(L^m) = (1 - c_1L^m - \dots - c_cL^{mc})$, $\delta_j(L) = (\delta_{j,0} + \delta_{j,1}L + \dots + \delta_{j,v}L^v)$, $j = 1, \dots, N$.

Note that if we assume that the lag orders c and v are large enough to make the error term ϵ_{t_m} uncorrelated, then, all the parameters in the U-MIDAS model (12) can be estimated by simple OLS (while the aggregation scheme $\omega(L)$ is supposed known). From a practical point of view, the lag order v could differ across variables, and v_i and c could be selected by an information criterion such as BIC.

A simple approach to forecasting is to use a form of direct estimation and construct the forecast as

$$\widetilde{y}_{T_m^x+m|T_m^x} = \widetilde{c}(L^k)y_{T_m^x} + \widetilde{\delta}_1(L)x_{1T_m^x} + \dots + \widetilde{\delta}_N(L)x_{NT_m^x}, \quad (13)$$

where the polynomials $\widetilde{c}(Z) = \widetilde{c}_1L^m + \dots + \widetilde{c}_cL^{mc}$ and $\widetilde{\delta}_i(L)$ are obtained by projecting y_{t_m} on information dated $mt_m - m$ or earlier, for $t = 1, 2, \dots, T_m^x$. In general, the direct approach of (13) can also be extended to construct h_m -step ahead forecasts given information in T_m^x .

$$\overline{y}_{T_m^x+h_m|T_m^x} = \overline{c}(L^k)y_{T_m^x} + \overline{\delta}_1(L)x_{1T_m^x} + \dots + \overline{\delta}_N(L)x_{NT_m^x}, \quad (14)$$

²Koenig, Dolmas, and Piger (2003) already proposed U-MIDAS in the context of real-time estimation. However, they did not systematically study the role of the functional form of the lag polynomial.

where the polynomials $\bar{c}(Z)$ and $\bar{\delta}_i(L)$ are obtained by projecting y_{t_m} on information dated $mt - h_m$ or earlier, for $t = 1, 2, \dots, T_m^x$.

In the case of U-MIDAS, the autoregressive term can be included easily without any common factor restriction as in Clements and Galvao (2009).

Finally, Carriero, Clark and Marcellino (2012) use Bayesian techniques to estimate specifications similar to U-MIDAS models with several regressors and stochastic volatility, which can easily produce not only point but also interval and density forecasts. We refer to their paper for the technical details.

2.3.4 Extensions of the MIDAS model

Different extensions of the MIDAS models have been analyzed in the literature, to introduce the use of mixed-frequency data in specific applications or studies, in which there is a need to capture particular features. For example, some studies incorporate regime changes in the parameters or asymmetric reactions to negative or positive values of the explanatory variables.

In what follows, we provide a brief overview of the extensions of the MIDAS models discussed so far in the literature.

Multiple explanatory variables

To allow for the inclusion of several additional explanatory variables into the MIDAS framework, it is necessary to extend the basic model above as follows:

$$y_{t_m} = \beta_0 + \beta_1 b(L_m; \theta_1) x_{1,t_m+w-h_m}^{(m)} + \beta_2 b(L_m; \theta_2) x_{2,t_m+w-h_m}^{(m)} + \varepsilon_{t_m}. \quad (15)$$

In this case, we consider x_1 and x_2 as two different explanatory variables. The values of the theta parameters are assumed to take on independent values and are thus represented by two independent vectors for the parameters, which may have different lag lengths.

Obviously, the above specification may be extended to allow for the inclusion of more than two explanatory variables (or more than two lags), and for the presence of an autoregressive structure. The most general MIDAS linear regression model can then be written as

$$y_{t_m} = \beta_0 + \sum_{i=1}^K \sum_{j=1}^L b_{ij}(L_{m_i}; \theta) x_{t_m+w-h_m}^{(m_i)} + \varepsilon_{t_m}. \quad (16)$$

Within the more general framework, it is also possible to include explanatory variables at different frequencies, since each indicator is modelled with its own polynomial parameterization. As an example, quarterly GDP growth can be explained not only by monthly indicators but also by weekly financial variables, with the explanatory variables, therefore, sampled at two different frequencies.

Nonlinear MIDAS models

Ghysels, Sinko and Valkanov (2007) further generalize (16) to:

$$y_{t_m} = \beta_0 + f \left(\sum_{i=1}^K \sum_{j=1}^L b_{ij} (L_{m_i}; \theta) g \left(x_{t_m+w-h_m}^{(m_i)} \right) \right) + \varepsilon_{t_m}, \quad (17)$$

where the functions f and g can be either fully known or parameter dependent. This model is inspired by the EGARCH model, and can be useful especially in volatility applications and risk-return trade-off studies.

Asymmetric MIDAS models

Ghysels, Santa-Clara and Valkanov (2005) introduce the asymmetric MIDAS model given by:

$$y_{t_m} = \beta_0 + \beta_1 \left(\phi b(L_m; \theta^-) 1_{t_m-h_m}^- x_{t_m+w-h_m}^{(m)} + (2 - \phi) b(L_m; \theta^+) 1_{t_m-h_m}^+ x_{t_m+w-h_m}^{(m)} \right) + \varepsilon_{t_m} \quad (18)$$

where $1_{t_m-h_m}^+$ denotes the indicator function for $x_{t_m+w-h_m}^{(m)} \geq 0$ and $1_{t_m-h_m}^-$ for $x_{t_m+w-h_m}^{(m)} < 0$, and $\phi \in (0, 2)$ in order to ensure that the total weights sum up to one. This formulation allows for a different impact of negative and positive values of the regressor x . The value of ϕ controls the different weight put on positive and negative impacts. Allowing for an asymmetric impact of the indicator is important in financial applications, especially in examining the asymmetric reaction of volatility in positive and negative return shocks.

Smooth Transition MIDAS models

Galvao (2007) proposes a new regression model which combines a smooth transition regression with a mixed data sampling approach:

$$y_{t_m} = \beta_{0,h_m}^{(m)} + \beta_{1,h_m}^{(m)} x_{t_m+w-h_m}^{(m)} \left[1 - G_{t_m+w-h_m} \left(x_{t_m+w-h_m}^{(m)}; \gamma, c \right) \right] + \beta_{2,h_m}^{(m)} x_{t_m+w-h_m}^{(m)} \left[G_{t_m+w-h_m} \left(x_{t_m+w-h_m}^{(m)}; \gamma, c \right) \right] + \varepsilon_{t_m} \quad (19)$$

where

$$G_{t_m+w-h_m} \left(x_{t_m}^{(m)}; \gamma, c \right) = \frac{1}{1 + \exp \left(-\gamma / \hat{\sigma}_x \left(x_{t_m+w-h_m}^{(m)}; \gamma, c \right) \right)} \quad (20)$$

The transition function is a logistic function that depends on the weighted sum of the explanatory variable in the current quarter.

The time-varying structure allows for changes in the predictive power of the indicators. This can be particularly relevant when one wants to use asset returns for forecasting macroeconomic variables, since changes in the predictive power of asset returns on economic activity may be related to business cycle regimes.

Markov-Switching MIDAS models

Guerin and Marcellino (2011) incorporate regime changes in the parameters of the MIDAS models. The basic version of the Markov-Switching MIDAS (MS-MIDAS) regression

model they propose is:

$$y_{t_m} = \beta_0 (S_{t_m}) + \beta_1 (S_{t_m}) B (L_m; \theta) x_{t_m+w-h_m}^{(m)} + \varepsilon_{t_m} (S_{t_m}) \quad (21)$$

where $\varepsilon_{t_m} | S_{t_m} \sim NID(0, \sigma^2 (S_{t_m}))$. The regime generating process is an ergodic Markov-chain with a finite number of states S_{t_m} .

These models allow also mixed-sample estimation of the probabilities of being in a given regime, which are relevant, for example, when one wants to predict business cycle regimes.

MIDAS with step functions

Forsberg and Ghysels (2007) introduce a MIDAS regression with step functions, where the distributed lag pattern is approximated by a number of discrete steps. To define this MIDAS regression, we consider the regressors $X(t_m, K_i) = \sum_{j=1}^{K_i} x_{t_m-j}^{(m)}$, which are partial sums of the high frequency variables. Then the MIDAS regression with M steps is:

$$y_{t_m} = \beta_0 + \sum_{i=1}^M \beta_i X(t_m, K_i) + \varepsilon_{t_m}. \quad (22)$$

This special case of MIDAS models can be reconnected to the U-MIDAS case we have analyzed in Section 2.3.3, in which the steps are the single individual observations.

Multivariate MIDAS models

Regression (16) can be generalized to multivariate specifications:

$$Y_{t_m} = B_0 + \sum_{i=1}^K \sum_{j=1}^L B_{ij} (L_{m_i}; \theta) X_{t_m+w-h_m}^{(m_i)} + \varepsilon_{t_m}, \quad (23)$$

where Y , ε and X are n -dimensional vector processes B_0 is an n -dimensional vector and B_{ij} are $n \times n$ matrices of polynomials. The main issue is how to handle parameter proliferation in a multivariate context. One approach is to consider all the off-diagonal elements controlled by one polynomial, while the diagonal elements by a second one. Of course, the restrictions may not be valid, and will be chosen depending on the application.

Considering multivariate MIDAS regressions allows to address Granger causality issues, avoiding temporal aggregation errors that can disguise or create spurious causality.

2.4 Mixed-frequency VAR

While so far, we have seen models which take into account mixed-frequency data in a univariate approach, we now focus on multivariate methods which jointly specify the dynamics of the indicators and of the variable to be explained. To exploit the information available in series released at different frequencies and jointly analyze them, there is

a growing literature which looks at mixed-frequency VARs, which aim to characterize the co-movements in the series and summarize the information contained in the mixed-frequency data.

Nowadays, in the literature, there are both classical and Bayesian approaches to estimate MF-VAR models. In what follows, we describe the main features of these two classes of estimation, following two of the most representative studies in the literature, Mariano and Murasawa (2010) for the classical approach and Schorfheide and Song (2011) for the Bayesian approach.

Classical framework

One of the most compelling approaches in the literature to deal with mixed-frequency time series at the moment is the one proposed by Zadrozny (1988) for directly estimating a VARMA model sampled at different frequencies, see also Harvey (1989). The approach treats all the series as generated at the highest frequency, but some of them are not observed. Those variables that are observed only at the low frequency are therefore considered as periodically missing.

Following the notation of Mariano and Murasawa (2010), we consider the state-space representation of a VAR model in a classical framework, treating quarterly series as monthly series with missing observations and taking GDP growth as an example. The disaggregation of the quarterly GDP growth, y_{t_m} , observed every $t_m = 3, 6, 9, \dots, T_m$, into the month-on-month GDP growth, $y_{t_m}^*$, never observed, is based on the following aggregation equation:

$$\begin{aligned} y_{t_m} &= \frac{1}{3} (y_{t_m}^* + y_{t_m-1}^* + y_{t_m-2}^*) + \frac{1}{3} (y_{t_m-1}^* + y_{t_m-2}^* + y_{t_m-3}^*) + \\ &\quad + \frac{1}{3} (y_{t_m-2}^* + y_{t_m-3}^* + y_{t_m-4}^*) \\ &= \frac{1}{3} y_{t_m}^* + \frac{2}{3} y_{t_m-1}^* + y_{t_m-2}^* + \frac{2}{3} y_{t_m-3}^* + \frac{1}{3} y_{t_m-4}^*. \end{aligned} \quad (24)$$

This aggregation equation comes from the assumption that the quarterly GDP series (in log levels), Y_{t_m} , is the geometric mean of the latent monthly random sequence $Y_{t_m}^*, Y_{t_m-1}^*, Y_{t_m-2}^*$. Taking the three-period differences and defining $y_{t_m} = \Delta_3 Y_{t_m}$ and $y_{t_m}^* = \Delta Y_{t_m}^*$, we obtain eq. (24).

Let for all t_m the latent month-on-month GDP growth $y_{t_m}^*$ and the corresponding monthly indicator x_{t_m} follow a bivariate VAR(p) process

$$\phi(L_m) \begin{pmatrix} y_{t_m}^* - \mu_y^* \\ x_{t_m} - \mu_x \end{pmatrix} = u_{t_m}, \quad (25)$$

where $u_{t_m} \sim N(0, \Sigma)$.

The VAR(p) process in eq. (25) together with the aggregation equation (24) is then cast in a state-space representation.

Assuming $p \leq 4^3$ and defining

$$s_{t_m} = \begin{pmatrix} z_{t_m} \\ \vdots \\ z_{t_m-4} \end{pmatrix}, \quad z_{t_m} = \begin{pmatrix} y_{t_m}^* - \mu_y^* \\ x_{t_m} - \mu_x \end{pmatrix},$$

a state-space representation of the MF-VAR is

$$s_{t_m} = F s_{t_m-1} + G v_{t_m} \quad (26)$$

$$\begin{pmatrix} y_{t_m} - \mu_y \\ x_{t_m} - \mu_x \end{pmatrix} = H s_{t_m} \quad (27)$$

with $\mu_y = 3\mu_y^*$ that holds, and $v_{t_m} \sim N(0, I_2)$ The matrices are defined as:

$$F = \begin{bmatrix} F_1 \\ F_2 \end{bmatrix}; \quad F_1 = \begin{bmatrix} \phi_1 & \dots & \phi_p & 0_{2 \times 2(5-p)} \end{bmatrix}; \quad F_2 = \begin{bmatrix} I_8 & 0_{8 \times 2} \end{bmatrix}, \quad (28)$$

$$G = \begin{bmatrix} \Sigma^{1/2} \\ 0_{8 \times 2} \end{bmatrix}; \quad H = \begin{bmatrix} H_0 & \dots & H_4 \end{bmatrix} \quad (29)$$

where H contains the lag polynomial

$$H(L_m) = \begin{bmatrix} 1/3 & 0 \\ 0 & 1 \end{bmatrix} + \begin{bmatrix} 2/3 & 0 \\ 0 & 0 \end{bmatrix} L_m + \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} L_m^2 + \begin{bmatrix} 2/3 & 0 \\ 0 & 0 \end{bmatrix} L_m^3 + \begin{bmatrix} 1/3 & 0 \\ 0 & 0 \end{bmatrix} L_m^4 \quad (30)$$

The state-space model consisting of equations (26) and (27) can be estimated with maximum-likelihood techniques or the expectation-maximization algorithm, where we have to take into account missing values due to publication lags and the low-frequency nature of the GDP. We illustrate the estimation and forecasting issues later on, in Section 3.1, where we review the problems related to ragged-edge data.

Bayesian framework

The estimation of MF-VAR model with Bayesian techniques has been recently considered as an alternative framework in the literature. One of the earliest studies on this is the paper by Chiu et al. (2011). In this paper, the authors develop a Gibbs sampling approach to estimate a VAR with mixed and irregularly sampled data. The algorithm they develop is a Gibbs sampler which iterates over the draws from the missing data and from

³For the sake of conciseness, we do not report the state-space representation for $p > 4$. Details for this case can be found in Mariano and Murasawa (2010).

the unknown parameters in the model. Under the assumption of a normally distributed error term, the algorithm allows for draws from Gaussian conditional distributions for estimating the missing data, and for draws from Gaussian and inverse Wishart conditional posterior distributions for the parameters in the model.

As an example for the Bayesian estimation of a MF-VAR, we present the algorithm developed by Schorfheide and Song (2011). The authors represent the MF-VAR as a state-space model, and use MCMC methods to conduct Bayesian inference for model parameters and unobserved monthly variables.

The state equation of the model is represented by the VAR(p) model written in the companion form:

$$z_{t_m} = F_1(\Phi) z_{t_m-1} + F_c(\Phi) + v_{t_m}, \quad v_{t_m} \sim iidN(0, \Omega(\Sigma)). \quad (31)$$

To write the measurement equation, the authors need to write the aggregation equation, which is in this case different from the one considered by Mariano and Murasawa (2010). In this case, the quarterly variable is seen as the three-month average of the monthly process, which in the previous notation is:

$$y_{t_m} = \frac{1}{3} (y_{t_m}^* + y_{t_m-1}^* + y_{t_m-2}^*) = \Lambda_{mz} z_{t_m}. \quad (32)$$

However, since y_{t_m} is observed only every third month, there is a need of a selection matrix that equals the identity matrix if t_m corresponds to the last month of the quarter and is empty otherwise. Therefore, the measurement equation can be written as

$$\begin{pmatrix} y_{t_m} \\ x_{t_m} \end{pmatrix} = M_{t_m} \Lambda_z z_{t_m}, \quad (33)$$

where M_{t_m} is the selection matrix. A Minnesota prior that shrinks the VAR coefficients toward univariate random walk representations is introduced to cope with the issue of dimensionality.

2.5 Mixed-frequency factor models

Closely related to the MF-VAR for their state-space representation, factor models have also been employed in the literature to handle data with different frequencies. These models have been utilized to extract an unobserved state of the economy and create a new coincident indicator, but also to exploit more information and obtain more precise forecasts. In what follows, we discuss the Mariano and Murasawa (2003) small scale mixed-frequency factor model, developed to extend the Stock–Watson coincident index for the US economy by combining quarterly real GDP and monthly coincident business cycle indicators. Interesting applications of a similar approach can be found in Frale et al. (2010, 2011). Then, we present an example of large scale mixed-frequency factor

model, as proposed by Giannone, Reichlin and Small (2008), whose aim is to bridge the information in a large monthly dataset with the forecast of a quarterly variable. As an extension to it, we present the mixed-frequency state-space framework as developed by Banbura and Rünstler (2011). Finally, based on Marcellino and Schumacher (2010), we analyze the approach that merges factor models and the MIDAS framework presented above.

2.5.1 Mixed-frequency small scale factor models

Factor models have a long tradition in econometrics and they are also appealing from an economic point of view. In fact, they decompose each time series under analysis into a common component, driven by few factors that represent the key economic driving forces, and an idiosyncratic component.

Mariano and Murasawa (2003) set up a static one-factor model for a small set of observable monthly and quarterly series, and derive its state-space representation.

Following their notation, consider a one-factor model for y_t^* , such that for all t_m ,

$$y_{t_m}^* = \mu^* + \Lambda f_{t_m} + u_{t_m} \quad (34)$$

$$\Phi_f(L) f_{t_m} = v_{t_m} \quad (35)$$

$$\Phi_u(L) u_{t_m} = w_{t_m} \quad (36)$$

$$\begin{pmatrix} v_{t_m} \\ w_{t_m} \end{pmatrix} \sim N \left(0, \begin{bmatrix} \Sigma_{vv} & 0 \\ 0 & \Sigma_{ww} \end{bmatrix} \right) \quad (37)$$

where $\Phi_f(\cdot)$ is a p th-order polynomial on \mathbb{R} and $\Phi_u(\cdot)$ is a q th-order polynomial on $\mathbb{R}^{N \times N}$. In order to have identification, we assume $\Lambda := [I, \Lambda_2]'$ and $\Phi_u(\cdot)$ and Σ_{ww} diagonal.

State-space representation

Assuming $p, q \leq 4$, for all t_m , and defining

$$s_t = \begin{pmatrix} f_{t_m} \\ \vdots \\ f_{t_m-4} \\ u_{t_m} \\ \vdots \\ u_{t_m-4} \end{pmatrix},$$

the state-space representation of the factor model is

$$s_{t_m+1} = F s_{t_m} + G v_{t_m} \quad (38)$$

$$y_{t_m} = \mu + H s_{t_m} \quad (39)$$

with $v_{t_m} \sim N(0, I_3)$, where

$$F = \begin{bmatrix} F_1 & F_2 \\ F_3 & F_4 \end{bmatrix}; \quad F_1 = \begin{bmatrix} \Phi_{f,1} \dots \Phi_{f,p} & 0_{1 \times (5-p)} \\ I_4 & 0_{4 \times 1} \end{bmatrix}; \quad F_2 = 0_{5 \times 10}; \quad (40)$$

$$F_3 = 0_{10 \times 5}; \quad F_4 = \begin{bmatrix} \Phi_{u,1} \dots \Phi_{u,q} & 0_{1 \times (5-q)} \\ I_8 & 0_{8 \times 2} \end{bmatrix}$$

$$G = \begin{bmatrix} \Sigma_{vv}^{1/2} & 0_{1 \times 2} \\ 0_{4 \times 1} & 0_{4 \times 2} \\ 0_{2 \times 1} & \Sigma_{ww}^{1/2} \\ 0_{8 \times 1} & 0_{8 \times 2} \end{bmatrix}; \quad H = \begin{bmatrix} H_0 \Lambda & \dots & H_4 \Lambda & H_0 & \dots & H_4 \end{bmatrix} \quad (41)$$

where $H(L_m)$ is defined as in equation (30).

In the estimation, Mariano and Murasawa (2003) cannot use the standard EM algorithm, since the measurement equation has unknown parameters. The procedure they followed is similar to the one described in Section 3.1.

The dynamic factor model as extended by Mariano and Murasawa (2003) is also used in Frale et al. (2011) to handle mixed frequency data, in order to obtain estimates of the monthly Euro area GDP components from the output and expenditure sides, to be later aggregated into a single indicator, called EUROMIND. Broadly speaking, GDP is disaggregated by supply sectors and demand components. For each of these sectors and components, timely and economically sensible observable monthly indicators are then selected and represented with a dynamic factor model, as described above. The single models are then linked together based on the composition of GDP.

2.5.2 Bridging with factors

We now discuss a large mixed frequency factor model as proposed by Giannone, Reichlin and Small (2008), which exploits a large number of series that are released at different times and with different lags. The methodology the authors propose relies on the two-step estimator by Doz et al. (2011). This framework combines principal components with the Kalman filter. First, the parameters of the model are estimated by OLS regression on the estimated factors, where the latter are obtained through principal components calculated on a balanced version of the dataset. Then, the Kalman smoother is used to update the estimate of the signal variable on the basis of the entire unbalanced panel.

The model

The dynamic factor model of Doz et al. (2011) is given by

$$x_{t_m} = \Lambda f_{t_m} + \xi_{t_m} \quad \xi_{t_m} \sim N(0, \Sigma_\xi) \quad (42)$$

$$f_{t_m} = \sum_{i=1}^p A_i f_{t_m-i} + B \eta_{t_m} \quad \eta_{t_m} \sim N(0, I_q) \quad (43)$$

Equation (42) relates the N monthly series x_{t_m} to a $r \times 1$ vector of latent factors f_{t_m} , through a matrix of factor loadings Λ , plus an idiosyncratic component ξ_{t_m} , assumed to be a multivariate white noise with diagonal covariance matrix Σ_ξ . Equation (43) describes the law of motion of the latent factors, which are driven by a q -dimensional standardized white noise η_{t_m} , where B is a $r \times q$ matrix ($r \leq q$). Hence, $\zeta_{t_m} \sim N(0, BB')$.

To deal with missing observations at the end of the sample, the authors use a two-step estimator. In the first step, the parameters of the model are estimated consistently through principal components on a balanced panel, created by truncating the data set at the date of the least timely release. In the second step, the Kalman smoother is applied to update the estimates of the factor and the forecast on the basis of the entire unbalanced data set (see Section 3.2.3 for more details on the estimation method).

The model is then complemented by a forecast equation for mean-adjusted quarterly GDP. The forecast is defined as the projection of the quarterly GDP growth on the quarterly aggregated estimated common factors:

$$\hat{y}_{t_q} = \alpha + \beta \hat{f}_{t_q}, \quad (44)$$

where \hat{f}_{t_q} is the quarterly aggregated correspondent of \hat{f}_{t_m} .

If we look at eq. (44), we see that this is exactly what we analyzed in Section (2.2) for the bridge equations. The framework can be interpreted as a large bridge model which uses a large number of variables and bridges monthly data releases with the forecast of the quarterly variable.

2.5.3 Factor models in a mixed-frequency state-space representation

Banbura and Rünstler (2011) extend the model of Giannone et al. (2008), by integrating a forecast equation for quarterly GDP. More specifically, they introduce the forecast of monthly GDP growth y_{t_m} as a latent variable, related to the common factors by the static equation

$$y_{t_m} = \beta' f_{t_m} + \varepsilon_{t_m}, \quad \varepsilon_{t_m} \sim N(0, \sigma_\varepsilon^2). \quad (45)$$

The quarterly GDP growth, y_{t_m} , is assumed to be the quarterly average of the monthly series:

$$y_{t_m} = \frac{1}{3} (y_{t_m}^* + y_{t_m-1}^* + y_{t_m-2}^*). \quad (46)$$

The innovations $\varepsilon_{t_m}, \eta_{t_m}, \xi_{t_m}$ are assumed to be mutually independent at all leads and lags.

Equations (42) to (46) can be cast in state-space form. y_{t_m} is constructed in such a way that it contains the quarterly GDP growth in the third month of each quarter, while the other observations are treated as missing.

State-space representation

The state-space representation, when $p = 1$, is:

$$\begin{bmatrix} x_{t_m} \\ y_{t_m} \end{bmatrix} = \begin{bmatrix} \Lambda & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} f_{t_m} \\ y_{t_m}^* \\ y_{t_m}^C \end{bmatrix} + \begin{bmatrix} \xi_{t_m} \\ \varepsilon_{t_m} \end{bmatrix} \quad (47)$$

$$\begin{bmatrix} I_r & 0 & 0 \\ -\beta' & 1 & 0 \\ 0 & -1/3 & 1 \end{bmatrix} \begin{bmatrix} f_{t_{m+1}} \\ y_{t_{m+1}}^* \\ y_{t_{m+1}}^C \end{bmatrix} = \begin{bmatrix} A_1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \Xi_{t_{m+1}} \end{bmatrix} \begin{bmatrix} f_{t_m} \\ y_{t_m}^* \\ y_{t_m}^C \end{bmatrix} + \begin{bmatrix} B\eta_{t_{m+1}} \\ 0 \\ 0 \end{bmatrix} \quad (48)$$

The aggregation rule (46) is implemented in a recursive way, by introducing a latent cumulator variable $y_{t_m}^C = \Xi_{t_m} y_{t_{m-1}}^C + \frac{1}{3} y_{t_m}^*$, where $\Xi_{t_m} = 0$ for t_m corresponding to the first month of the quarter and $\Xi_{t_m} = 1$ otherwise. The estimation of the model parameters follows Giannone, Reichlin and Small (2008).

2.5.4 Factor-MIDAS

It is possible to augment the MIDAS regressions with the factors extracted from a large dataset to obtain a richer family of models that exploit a large high-frequency dataset to predict a low-frequency variable.

While the basic MIDAS framework consists of a regression of a low-frequency variable on a set of high-frequency indicators, the Factor-MIDAS approach exploits estimated factors rather than single or small groups of economic indicators as regressors.

Marcellino and Schumacher (2010) propose alternative MIDAS regressions. In the standard MIDAS case, they follow Clements and Galvao (2008), while as a modification they evaluate a more general regression approach, labeled unrestricted Factor-MIDAS, where the dynamic relationship between the low-frequency variables and the high-frequency indicators is unrestricted, in contrast to the distributed lag functions as proposed by Ghysels et al. (2007). As a third alternative, they consider a regression scheme proposed by Altissimo et al. (2010), which considers only correlation at certain frequencies between variables sampled at high- and low- frequencies. This approach is called smoothed MIDAS, since the regression essentially eliminates high-frequency correlations.

The information set consists of a large set of stationary monthly indicators, X_{t_m} . The last observation is at time $T_m + w$, $w > 0$, allowing for at most $w > 0$ monthly values of the indicators that are earlier available than the lower-frequency variable to be estimated. X_{t_m} is modeled using a factor representation, where r factors F_{t_m} are estimated in order to

summarize the information in X_{t_m} . The estimated factors, \widehat{F}_{t_m} , are used in the projection for the quarterly-frequency variable.

We now describe in details the three alternative Factor-MIDAS approaches proposed by Marcellino and Schumacher (2010), assuming again that the target variable is GDP. These approaches are tools for direct multi-step now- and forecasting, thus each model is for a specific forecast horizon.

The Basic Factor-MIDAS approach

In the basic Factor-MIDAS approach the explanatory variables used as regressors are estimated factors. Assume for simplicity $r = 1$, so that there is only one factor \widehat{f}_{t_m} . The Factor-MIDAS model for forecast horizon h_q quarters with $h_q = h_m/3$ is

$$y_{t_q+h_q} = y_{t_m+h_m} = \beta_0 + \beta_1 b(L_m; \theta) \widehat{f}_{t_m+w}^{(3)} + \varepsilon_{t_m+h_m}, \quad (49)$$

where $b(L_m; \theta) = \sum_{k=0}^K c(k; \theta) L_m^k$ and $c(k; \theta) = \frac{\exp(\theta_1 k + \theta_2 k^2)}{\sum_{k=0}^K \exp(\theta_1 k + \theta_2 k^2)}$.

$\widehat{f}_{t_m}^{(3)}$ is skip-sampled from the monthly factor \widehat{f}_{t_m} . Every third observation starting from the final one is included in the regressor $\widehat{f}_{t_m+w}^{(3)}$, i.e. $\widehat{f}_{t_m+w}^{(3)} = \widehat{f}_{t_m+w}, \forall t_m + w = \dots, T_m + w - 6, T_m + w - 3, T_m + w$. Note that we take into account the fact that a monthly indicator is typically available also within the quarter for which no GDP figure is available. As described above in the MIDAS models, the exponential lag function provides a parsimonious way to consider monthly lags of the factors.

The model can be estimated using nonlinear least squares in a regression of y_{t_m} onto the factors $\widehat{f}_{t_m+w-h}^{(3)}$. The forecast is given by

$$y_{T_m+h_m|T_m+w} = \widehat{\beta}_0 + \widehat{\beta}_1 b(L_m; \widehat{\theta}) \widehat{f}_{T_m+w}^{(3)}. \quad (50)$$

The projection is based on the final values of estimated factors.

MIDAS regression can be generalized to more than one factor and extended with the addition of autoregressive dynamics. Details on factor estimation are provided in Section 3.2.

Smoothed MIDAS

A different way to formulate a mixed-frequency projection is proposed by Altissimo et al. (2010). The projection can be written as:

$$y_{T_m+h_m|T_m+w} = \widehat{\mu} + G \widehat{F}_{T_m+w} \quad (51)$$

$$G = \widetilde{\Sigma}_{yF}(h_m - w) \times \widehat{\Sigma}_F^{-1}, \quad (52)$$

where $\widehat{\mu}$ is the sample mean of GDP, G is the projection coefficient matrix, $\widehat{\Sigma}_F$ is the estimated sample covariance of the factors, and $\widetilde{\Sigma}_{yF}(k)$ is a particular cross-covariance with k monthly lags between GDP and the factors. $\widetilde{\Sigma}_{yF}(k)$ is not an estimate of the sample

cross-covariance between factors and GDP, but a cross-covariance between smoothed GDP and factors. The smoothing aspect is introduced in $\tilde{\Sigma}_{yF}(k)$ as follows. The estimated covariance between \widehat{F}_{t_m-k} and y_{t_m} is:

$$\widehat{\Sigma}_{yF}(k) = \frac{1}{T^* - 1} \sum_{t_m=M+1}^{T_m} y_{t_m} \widehat{F}_{t_m-k}^{(3)'} \quad (53)$$

where T^* is the number of observations available to compute the cross-covariances. Note the use of skip-sampled factors, since GDP is available only quarterly. Given $\widehat{\Sigma}_{yF}(k)$, the estimated cross-spectral matrix is

$$\widehat{S}_{yF}(\omega_j) = \sum_{k=-M}^M \left(1 - \frac{|k|}{M+1}\right) \widehat{\Sigma}_{yF}(k) e^{-i\omega_j k}, \quad (54)$$

at frequencies $\omega_j = \frac{2\pi j}{2H}$ for $i = -H, \dots, H$ using a Bartlett lead-lag window. The low-frequency relationship between \widehat{F}_{t_m-k} and y_{t_m} is obtained by filtering out cross fluctuations at frequency higher than a certain threshold π/q , using the frequency-response function $\alpha(\omega_j)$, defined as $\alpha(\omega_j) = 1, \forall |\omega_j| < \pi/q$ and zero otherwise. The autocovariance matrix $\tilde{\Sigma}_{yF}(k)$ reflecting low-frequency co-movements between \widehat{F}_{t_m-k} and y_{t_m} is obtained by inverse Fourier transform:

$$\tilde{\Sigma}_{yF}(k) = \frac{1}{2H+1} \sum_{j=-H}^H \alpha(\omega_j) \widehat{S}_{yF}(\omega_j) e^{i\omega_j k}. \quad (55)$$

Note that $\widehat{\Sigma}_{yF}(k)$ is a consistent estimator of the true cross-covariance, if the sample size is sufficiently large.

Unrestricted MIDAS

An alternative to the two previous models is the unrestricted lag order specification, as described in the above Section 2.3.3:

$$y_{T_m+h_m} = \beta_0 + D(L_m) \widehat{F}_{t_m+w}^{(3)} + \varepsilon_{t_m+h_m}, \quad (56)$$

where $D(L_m) = \sum_{k=0}^K D_k L_m^k$ is an unrestricted lag polynomial of order K .

$D(L_m)$ and β_0 are estimated by OLS. To specify the lag order in the empirical application, Marcellino and Schumacher consider a fixed scheme with $k = 0$ and an automatic lag length selection using the BIC.

3 Ragged-edge data

After having analyzed the various techniques to deal with mixed-frequency data, in this section we review the estimation methods that can handle ragged-edge data, that is datasets which are not balanced, due to the presence of missing values at the end of the sample for some indicators.

First, we discuss the issues of estimation and forecasting MF-VAR in the presence of missing observations due to publication lags and to the low-frequency nature of one variable. We follow Mariano and Murasawa (2010) in the discussion of estimation of the state-space form and forecasting with the use of Kalman filter and/or smoother.

Going further, we analyze issues related to factor estimation in the presence of ragged-edge data. Marcellino and Schumacher (2010) review three different methods to tackle it. First, the method proposed by Altissimo et al. (2010), who realign each time series in the sample in order to obtain a balanced dataset, and then estimate the factors with dynamic PCA. As an alternative, to consider missing values in the data for estimating factors Stock and Watson (2002) propose an EM algorithm together with the standard principal component analysis (DPCA). As a third method, Doz et al. (2011) propose a factor estimation approach based on a complete parametric representation of the large factor model in state-space form.

3.1 Estimating the MF-VAR with missing observations

As already anticipated in Section 2, Kalman filtering techniques can handle ragged-edge data and missing values due to publication lags and the low-frequency nature of a time series.

Estimation

The state-space representation of the mixed-frequency VAR model is described by equations (26) and (27), reported also here:

$$\begin{aligned} s_{t_m} &= F s_{t_m-1} + G v_{t_m} \\ \begin{pmatrix} y_{t_m} - \mu_y \\ x_{t_m} - \mu_x \end{pmatrix} &= H s_{t_m}. \end{aligned}$$

It can be estimated by maximum-likelihood even in the presence of missing observations due to publication lags and the low-frequency nature of GDP. However, as Mariano and Murasawa (2010) mention in their paper, when the number of parameters is large, the ML method can fail to converge.

In these cases, it is useful to implement the EM algorithm modified to allow for missing observations. Mariano and Murasawa (2010) consider the missing values as realizations

of some iid standard normal random variables, i.e.

$$y_{t_m}^+ = \begin{cases} y_{t_m} & \text{if } y_{t_m} \text{ is observable} \\ \zeta_{t_m} & \text{otherwise} \end{cases} \quad (57)$$

where ζ_{t_m} is a draw from a standard normal distribution independent of the model parameters.

The measurement equation is modified accordingly in the first two months of each quarter, where the upper row of H is set to zero and a standard normal error term is added, so that the Kalman filter skips the random numbers. Since the realizations of the random numbers do not matter in practice, the authors suggest to replace the missing values with zeros. Then, the EM algorithm is employed to obtain estimates of the parameters.

Estimation of latent monthly real GDP

Mariano and Murasawa (2010) use the Kalman smoother instead of the Kalman filter, because it uses more information and also simplifies the formulation of the state-space model. Although GDP growth for a particular month is not available, the smoother considers the monthly indicators available for the same quarter, so that nowcasting is also possible. For the months in which no observations are available also for the monthly indicators, the Kalman smoother acts exactly as the Kalman filter.

3.2 Estimating the factors with ragged-edge data

Factor forecasting with large, single-frequency data is often carried out using a two-step procedure. First, the factors are estimated and, second, a dynamic model for the variable to be predicted is augmented with the estimated factors. The same procedure can be used also in case of mixed-frequency data. As an alternative, factor estimation and forecasting can be conducted in a single step, in the context of a parametric model.

The literature provides various ways to estimate the factors, in case of balanced datasets. However, in the following subsections we describe the methods that can handle ragged-edge data, that is datasets which are not balanced due to the presence of missing values at the end of the sample for some indicators.

3.2.1 Vertical realignment of data and DPCA

Altissimo et al. (2010) provide a convenient way to solve the ragged-edge problem. They propose to realign each time series in the sample in order to obtain a balanced dataset. Assume that variable i is realized with k_i months of publication lags. Thus, given a dataset in period $T_m + w$, the final observation of time series i is for period $T_m + w - k_i$. Altissimo et al. (2010) propose to realign the series in this way:

$$\tilde{x}_{i,T_m+w} = x_{i,T_m+w-k_i}, \quad (58)$$

for $t_m = k_1 + 1, \dots, T_m + w$. Once applied to each time series, the result is a balanced dataset \tilde{X}_{t_m} , for $t_m = \max\{k_1\}_{i=1}^N + 1, \dots, T_m + w$.

Given this balanced dataset, Altissimo et al. (2010) propose dynamic PCA to estimate the factors. The two-step estimation procedure introduced by Forni et al. (2005) directly applies, since the dataset is balanced.

One of the main advantages of this method is the simplicity. A drawback is that the availability of data determines cross-correlations between variables. Moreover, data releases are not the same over time, so that dynamic correlations within the data change and, as a consequence, factors change over time. The same happens if factors are reestimated at a higher frequency than the one of the factor model, for example in the case of a monthly factor model reestimated several times within a month, in correspondence of new releases of the data.

3.2.2 EM algorithm and PCA

As an alternative, to consider missing values in the data for estimating factors, Stock and Watson (2002) propose an EM algorithm combined with the standard PCA. Call X_i the column i of the dataset X_{t_m} : not all observations are available, due to publication lags. The vector X_i^{obs} contains the observations available for variable i , a subset of X_i . More precisely, the relation between observed and not fully observed variables is

$$X_i^{obs} = A_i X_i, \quad (59)$$

where A_i is the matrix that tackles missing values. A_i is equal to the identity matrix, in case there are no missing values in the series. When an observation is missing at the end of the sample, the corresponding final row of the identity matrix is removed.

The EM algorithm consists in the following steps:

1. Provide an initial guess $\hat{X}_i^{(0)}, \forall i$. These guesses together with the fully available series provide a balanced dataset $\hat{X}^{(0)}$. With a balanced dataset, standard PCA gives initial monthly factors $\hat{F}^{(0)}$ and loadings $\hat{\Lambda}^{(0)}$.
2. E-step: an updated estimate of the missing observations for variable i is provided by the expectation of X_i conditional on X_i^{obs} , factors $\hat{F}^{(j-1)}$ and loadings $\hat{\Lambda}^{(j-1)}$ from the previous iteration

$$\hat{X}_i^{(j)} = \hat{F}^{(j-1)} \hat{\Lambda}_i^{(j-1)} + A_i' (A_i' A_i)^{-1} \left(X_i^{obs} - A_i \hat{F}^{(j-1)} \hat{\Lambda}_i^{(j-1)} \right). \quad (60)$$

We can recognize two components in the update: the common component from the previous iteration $\hat{F}^{(j-1)} \hat{\Lambda}_i^{(j-1)}$, and a low-frequency idiosyncratic component $X_i^{obs} - A_i \hat{F}^{(j-1)} \hat{\Lambda}_i^{(j-1)}$, distributed by the projection coefficient $A_i' (A_i' A_i)^{-1}$ on the high-frequency periods, see Breitung and Schumacher (2008).

3. M-step: repeat the E-step for each variable i , in order to obtain a balanced dataset. Reestimate all the factors $\widehat{F}^{(j)}$ and loadings $\widehat{\Lambda}^{(j)}$ by PCA. Go back to step 2 until convergence.

After convergence, the EM algorithm provides both the monthly factor estimates \widehat{F}_{t_m} and the estimates of the missing values of the time series, see also Angelini et al. (2006).

3.2.3 Large parametric factor model in state-space form

Doz et al. (2011) propose a factor estimation approach based on a complete representation of the large-factor model in state-space form. The full state-space model has the form

$$X_{t_m} = \Lambda F_{t_m} + \xi_{t_m} \quad (61)$$

$$\Psi(L_m) F_{t_m} = B \eta_{t_m}. \quad (62)$$

Equation (61) is the static factor representation of X_{t_m} . Equation (62) specifies a VAR structure of the factors, with lag polynomial $\Psi(L_m) = \sum_{i=1}^p \Psi_i L_m^i$. η_{t_m} is a q -dimensional vector that contains the orthogonal dynamic shocks that drive the r factors. The factors F_{t_m} represent the states, while ξ_{t_m} is the stationary idiosyncratic component which admits a Wold representation. The shocks driving the factors and the idiosyncratic components are assumed to be independent. If the X_{t_m} is of a small dimension, the model can be estimated by iterative maximum likelihood. If the dimension is large, iterative ML is infeasible, so Doz et al. (2011) propose a quasi-ML to estimate the factors. For a given number of factors, r , and dynamic shocks, q , the estimation follows the steps illustrated below:

1. Estimate \widehat{F}_{t_m} using PCA as an initial estimate. The estimation is based on the balanced part of the data, obtained by removing the values at the end of the sample that create the unbalancedness.
2. Estimate the loadings $\widehat{\Lambda}$ by regressing X_{t_m} on the factors estimated in the previous step, \widehat{F}_{t_m} . Estimate also the covariance of the idiosyncratic components $\widehat{\xi}_{t_m}$, denoted as $\widehat{\Sigma}_\xi$.
3. Estimate a VAR(p) on the factors \widehat{F}_{t_m} , obtaining $\widehat{\Psi}(L_m)$, and the residual covariance of $\widehat{\varsigma}_{t_m} = \widehat{\Psi}(L_m) \widehat{F}_{t_m}$, denoted as $\widehat{\Sigma}_\varsigma$.
4. To obtain an estimate of B , given the number of dynamic shocks q , apply an eigenvalue decomposition of $\widehat{\Sigma}_\varsigma$. Call M the $(r \times q)$ -dimensional matrix of the eigenvectors corresponding to the q largest eigenvalues, and call P the $(q \times q)$ -dimensional matrix with the largest eigenvalues on the main diagonal and zero otherwise. The estimate of B is $\widehat{B} = M \times P^{-1/2}$. All the parameters and coefficients in the system of

equations (61) and (62) are then fully specified. The model is cast into state-space form.

5. The Kalman filter or smoother yields new estimates of the monthly factors. The dataset used now is the unbalanced one, where T_m is the last observation available in the whole set of monthly series. The Kalman filter also provides estimates and forecasts for the missing values conditional on the model structure and properties of the shocks.

Note that the coefficients in the system have to be estimated from a balanced subsample of data, as in step 1 there is the need of a fully balanced dataset for PCA initialization. Nevertheless, in step 5 the factor estimation based on the Kalman filter applies to the unbalanced dataset. The solution is to estimate the coefficients outside the state-space model and avoid to estimate a large number of coefficients by iterative ML.

4 A comparison of the different methods

So far, we have seen that several methods have been proposed in the literature to deal with mixed-frequency data, possibly with a ragged edge structure. In general, there is an agreement on the fact that exploiting data at different frequencies matters for nowcasting and short term forecasting. We now try to summarize the advantages and disadvantages of the different methods, comparing their most important features.

Bridge equations are still one of the most used techniques, especially in short run forecasting, because they are pretty easy to estimate and interpret, and allow computing early estimates of the low-frequency variable. The drawback is that they are purely statistical models, where the regressors are included only because they contain timely updated information. Therefore, if the model that exploits the high-frequency information is misspecified, the error transmits to the bridge equation and to the forecasts that are obtained recursively.

A more sophisticated way to deal with data sampled at different frequencies is the state-space approach. Casting the model in state-space form has the advantage of jointly specifying the dynamics of the indicators and of the variable to be explained without imposing any a-priori restriction. Moreover, since the low-frequency series is seen as a high-frequency series with missing values, the use of the Kalman filter permits the estimation of these missing data. As shown in Bai, Ghysels and Wright (2011), the Kalman filter results to be the optimal filter in population, when ignoring parameter estimation errors and assuming that the model is correctly specified. Therefore, under these ideal circumstances, the state-space approach cannot be beaten by any other method. On the other side, there are also some drawbacks from the use of this approach: first of all, in most of the cases it is computationally complex, and the complexity increases with the number of variables involved, so that most of the time only small-scale models can be estimated.

Moreover, the state space approach requires the correct specification of the model in high frequency, which is even more complex than usual given the missing observations in the dependent variable.

An alternative way to deal with mixed-frequency data is the MIDAS approach. Even though in population, when the process is correctly specified, MIDAS is coarse than the optimal Kalman filter, it can be more robust in the presence of mis-specification. Moreover, the lag polynomials are based on a very small number of parameters, allowing the MIDAS models to be parsimonious, even though it is still not clear which is the best polynomial specification. Contrary to what stated for the state-space models, MIDAS models can be easily estimated by NLS. However, it is only possible to obtain a high frequency update of the expected low frequency realization, not an estimate of the missing values in the low frequency variable.

Both the state-space and the MIDAS approaches can be combined with a factor specification, in order to use the information in a large dataset, possibly with a ragged edge. Whether factor methods provide more precise estimates and forecasts than VARs or single equation methods is a matter for empirical investigation, since there is a trade-off between model complexity and extended information set.

5 An overview of empirical studies

In this section we review the empirical literature on forecasting with mixed-frequency and ragged-edge data, providing some examples of all the models and estimation methods outlined in the previous sections.

5.1 Bridge equations

Bridge equations have been one of the first methods employed in nowcasting the current state of the economy, making use of the monthly information available. Studies of this kind have been conducted for the nowcasts of different economies. A common finding of these studies is that the exploitation of intra-period information reduces the forecasting error in the majority of the cases. The applications concern both "supply-side" and "demand-side" models.

Looking at US data, Trehahn and Ingenito (1996) construct a model that predicts current quarter real GDP based on knowledge of nonfarm payrolls, industrial production and real retail sales, which have the advantage of being released at a monthly frequency by the middle of the following month. In order to produce a model that predicts real GDP, the authors rely on auxiliary models that generate forecasts of the indicator variables themselves. Their evidence shows that consumption data provide key information about current output, and that retail sales release allows to have a good forecast of contemporaneous consumption.

Stark (2000) presents evidence on the usefulness of conditioning quarterly model forecasts on monthly current-quarter data, in the case of the US economy. Starting by generating a one-step-ahead forecast from a quarterly Bayesian vector error correction model, the author then specifies a monthly statistical model for variables that are thought to carry information about each of the variables in the quarterly model and uses it to generate sequences of current-quarter quarterly-average forecasts from the monthly indicators. Once he has these quarterly-average monthly indicator forecasts, he forms updated estimates of the quarterly model's current-quarter forecast. The findings show that exploiting monthly information produces economically and statistically significant improvements, particularly large especially during periods of recession.

A study by Barhoumi et al. (2011) presents a model to predict French gross domestic product (GDP) quarterly growth rate. The authors employ the bridge equations to forecast each component of the GDP, and select the monthly explanatory variables among a large set of hard and soft data. They find that changing the set of equations over the quarter is superior to keeping the same equations over time. These models turn out to beat the benchmark in terms of forecasting performance.

Studies involving bridge equations can be found for many other countries. In particular, bridge models have been employed also for nowcasting Euro Area GDP growth. As an example, we consider Baffigi et al. (2004). In this paper, bridge models are estimated for aggregate GDP and components, both area-wide and for the main countries of the Euro Area. Their short-term performance is then assessed with respect to benchmark univariate and multivariate standard models, and a small structural model. The results shown in the paper are clear-cut: bridge models performance is always better than benchmark models, provided that at least some indicators are available over the forecasting horizon. As far as the type of aggregation is concerned, the supply-side approach (modelling aggregate GDP) performs better than the demand-side approach (aggregation of forecasts by national account components). The supply-side models highlight the role of industrial production and manufacturing surveys as the best monthly indicators. Looking at the demand-side models, from the different equations estimated in this paper, private consumption results well tracked by retail sales index, while the consumer confidence index plays a minor role; in the case of investment a major role seems to be played by survey variables.

Diron (2008) makes use of bridge equations with Euro Area data to provide an assessment of forecast errors, which takes into account data-revisions. Using four years of data vintages, the paper provides estimates of forecast errors for Euro Area real GDP growth in genuine real-time conditions and assesses the impact of data revisions on short-term forecasts of real GDP growth. Given the size of revision to indicators of activity, the assessment of reliability of short-term forecasts based on revised series could potentially give a misleading picture. Nevertheless, averaging across all bridge equations, forecasts of individual quarters tend to be similar whether they are based on preliminary or revised

data. More specifically, the RMSEs based on real-time and pseudo real-time exercises are quite similar and both smaller compared with AR forecasts of GDP, considered as the benchmark. The difference in forecast accuracy is significant according to Diebold and Mariano tests, highlighting that short-term forecasts based on bridge equations are informative. Moreover, the paper investigates the contributions of the various sources to the overall forecasting errors. Revisions to the monthly variables and to GDP growth account only for a small share of the overall forecast error, while the main sources are from extrapolation of the monthly indicators. The relative contributions of extrapolation and of revision of monthly indicators vary depending on whether the equations include hard data, in which case both sources are significant, or survey and financial variables, in which case these two sources tend to have a smaller weight.

In the context of nowcasting, it has become more common to exploit the information coming from a large set of variables. Therefore, recent studies combine the bridge models with factors, in what is called "bridging with factors". This new kind of model is related to the one described in Section 2.5.2 by Giannone et al. (2008). In Section 5.4.2, we will review these studies and compare the performance of this new kind of bridge with factors to the standard bridge models and other benchmarks.

5.2 MIDAS models

In the first applications, MIDAS models have been applied to financial data, investigating the relation between the conditional mean and the conditional variance of the stock market returns or future volatility, see Ghysels et al. (2005) as an example. Clements and Galvao (2008) are the first to apply MIDAS regressions to macroeconomic data. In the next paragraphs we will overview applications based on financial and/or macroeconomic data.

Ghysels, Santa-Clara and Valkanov (2005) study the intertemporal relation between the conditional mean and the conditional variance of the aggregate market return. In support of Merton's ICAPM, the authors find a positive significant and robust relation between risk and return. They also find that the MIDAS estimator is a better forecaster of the stock market variance than two other benchmark models: rolling window and GARCH estimators. The authors also focus on the asymmetric reaction of volatility to positive and negative shocks. They find that positive shocks have a bigger impact overall on the conditional mean of returns, are slower to be incorporated in the conditional variance, and are much more persistent, while negative shocks have a large initial, but temporary, effect on the variance of returns. Ghysels, Santa-Clara and Valkanov (2006) consider various MIDAS regressions to predict volatility in a parsimonious way with data at different frequency. They find that daily realized power is the best predictor of future increments in quadratic variation. Surprisingly, the direct use of high-frequency (5 minutes) data does not improve volatility prediction.

Ghysels, Rubia and Valkanov (2009) compare three different approaches of producing multi-period-ahead forecasts of volatility: iterated, direct and MIDAS. The comparison

is conducted out-of-sample using returns data of the US stock market portfolio and a cross section of size, book-to-market and industry portfolios, in terms of the average forecasting accuracy, using the MSFE. The direct approach provides the worst forecasts. Iterated forecasts are suitable for shorter horizons, while MIDAS forecasts perform well at long horizons.

Clements and Galvao (2008) introduce the use of MIDAS regressions in forecasting macroeconomic data. They also look at whether a mixed-data sampling approach including an autoregressive term can improve forecasts of US real output growth. They conduct a real-time forecasting exercise that exploits monthly vintages of the indicators and the quarterly vintages of the output growth, consistent with the timing of the releases of the different data vintages. The authors find that the use of within-quarter information on monthly indicators can result in a marked reduction in RMSE compared with the more traditional quarterly-frequency AR or AR distributed-lag models. Moreover, Clements and Galvao (2009) evaluate the predictive power of leading indicators for output growth up to one year, using MIDAS approach to combine multiple leading indicators in a parsimonious way. The results confirm that MIDAS is a useful instrument to improve forecasts. Moreover, they show that the use of real-time vintage data improves forecast performance and that the predictive power of the indicators is stronger when the aim is to forecast final data rather than first-released data, although first releases can generally be forecasted more accurately.

Foroni, Marcellino and Schumacher (2012) compare the performance of the MIDAS with functional distributed lags estimated with NLS to the one of the U-MIDAS, the unrestricted version analyzed in Section 2.3.3. In Monte Carlo experiments, they show that U-MIDAS generally performs better than MIDAS when mixing quarterly and monthly data. On the other hand, with larger differences in sampling frequencies, distributed lag-functions outperform unrestricted polynomials. In an empirical application on out-of-sample nowcasting GDP in the Euro area and the US using monthly predictors, they find a good performance of U-MIDAS for a number of indicators, albeit the results depend on the evaluation sample.

In the recently increasing literature, which is exploiting the availability of a huge number of financial time series on a daily basis to forecast macroeconomic time series, the empirical evidence in support of the use of high-frequency financial series is rather mixed. On the one side, it is useful to use this great amount of timely information, but on the other side there is a question on how to weight the daily observations and to filter these data, to get rid of the possible noise. Results from recent studies suggest that daily variables seem to have useful information for forecasting inflation and economic activity.

Among these studies, Ghysels and Wright (2009) propose methods for using asset price data to construct daily forecasts of upcoming survey releases, employing MIDAS regression models and a more structural approach based on the Kalman filter to estimate what forecasters would predict if they were asked to make a forecast each day, treating

their forecasts as missing data to be interpolated. Their aim is to obtain high-frequency measures of forecasters' expectations. The authors consider two surveys in their empirical work: the Survey of Professional Forecasters and the Consensus Forecasts, and use the daily asset prices to predict the upcoming releases of either of the two surveys. In an in- and out-of-sample forecasting exercise, both approaches (MIDAS and Kalman filter) perform better than the simple random walk benchmark forecasts.

Andreou, Ghysels and Kourtellos (2010) assess whether daily financial data can improve macroeconomic forecasting, employing MIDAS regression models. They forecast US quarterly inflation rate and economic growth using a dataset including daily, monthly and quarterly indicators. An important advantage of the MIDAS model is that it can provide new forecasts as daily data become available. The authors find that on average daily financial predictors improve the forecasts of quarterly inflation and GDP relative to the AR benchmark model.

Monteforte and Moretti (2010) present a mixed frequency model for daily forecasts of Euro area inflation in real-time. The model they use allows to combine a monthly core inflation estimated from a dynamic factor model with daily financial market variables, which provide timely information on the recent shocks. They compare the results of this mixed-frequency model with standard univariate and multivariate models with monthly data, and also with the forecasts implied by financial market expectations extracted from future contracts. In both cases, the mixed frequency approach shows a superior predictive power.

5.3 Mixed-frequency VAR models

As we have seen in Section 2.4, studies on MF-VAR models have been conducted both in a classical and in a Bayesian context. We now outline the main empirical studies conducted in both frameworks.

Mittnik and Zadzorny (2005) evaluate a Kalman-filtering-based maximum-likelihood estimation method for forecasting German real GDP at monthly intervals. They estimate a VAR(2) model of quarterly GDP and up to three monthly indicator variables (industrial production, current and expected business conditions). They find that in general monthly models produce better short-term GDP forecasts, while quarterly models produce better long-term GDP forecasts.

Mariano and Murasawa (2010) apply the MF-VAR method to construct a new coincident indicator, i.e. an estimate of monthly real GDP. What they find is that the coincident index based on the VAR model is close to the one obtain by a factor model, and they both track well quarterly real GDP, although they are quite volatile.

Kuzin, Marcellino and Schumacher (2011) compare the MF-VAR, as presented in Mariano and Murasawa (2010), with the MIDAS approach to model specification in the presence of monthly and quarterly series. MIDAS leads to parsimonious models, while MF-VAR does not restrict the dynamics but suffers from the curse of dimensionality.

The authors argue that it is difficult to rank the different approaches a priori, so they compare their performance empirically, considering an AR process as a benchmark. The two approaches tend to be more complementary than substitutes, since the MF-VAR performs better for longer horizons, whereas MIDAS for shorter horizons. Looking at the relative MSE of the different models with respect to the benchmark, the mixed-frequency models perform relatively well, especially when forecast combinations are adopted.

Similar evidence is found by Foroni and Marcellino (2012) in their paper which focuses on different methods proposed in the literature to deal with mixed-frequency and ragged-edge datasets. The authors discuss the performance of the different methods on now- and forecasting the quarterly growth rate of the Euro Area GDP and its components, using a very large set of monthly indicators. They also find that MF-VAR outperforms the MIDAS approach only at longer horizons.

Ghysels (2011) introduces a different MF-VAR representation, in which he constructs the MF-VAR process as stacked skip-sampled processes. In this paper, the author characterizes explicitly the mis-specification of a traditional low frequency VAR and the consequent misspecification in the impulse response functions. Moreover, since the MF-VAR specified in this way can also characterize the timing of information releases, he shows how Choleski factorizations are a more natural tool for impulse response analysis because the elements in the vector represent a sequence of time events. As another contribution, he studies a Bayesian approach which accommodates the potentially large set of parameters to be estimated.

One of the earliest studies on Bayesian estimation of MF-VAR is the paper by Chiu et al. (2011). In this paper, the authors develop a Gibbs sampling approach to estimate a VAR with mixed and irregularly sampled data. The focus of the paper is on the parameter estimation. In an exercise with simulated data, the authors show that taking into account mixed-frequency data allows to obtain smaller root mean squared errors for all the parameter estimates regardless of the sample size and of the correlation between the variables of the system. These results find confirmation also in the two empirical examples, conducted with data respectively at monthly and quarterly and weekly and monthly frequency, for which the authors compare the posterior distributions of the parameters and the impulse response functions.

Another study by Viefers (2011) reconsiders the estimation of a MF-VAR as in Mariano and Murasawa (2010). First, the author makes use of the Bayesian MCMC algorithm to simulate and estimate the model, and second he extends the MF-VAR to allow for regime switching. In his model, the inference is based on the joint posterior density of all the unknowns. The findings of the simulation study suggest that inference on the latent series and the regime processes is fairly precise, although there is a more pronounced imprecision in the estimation of the VAR parameters. In the empirical exercise, the author considers monthly and quarterly data for the US economy. The results on the regime probabilities show a relative high ability to identify the same recession dates provided by the NBER,

although the probabilities tend to be more erratic and much worse in the most recent years.

Schorfheide and Song (2011) conduct a forecasting exercise on US data exploiting MF-VAR models. The goal of their paper is to study the extent to which the incorporation of monthly information improves the forecasts compared to models based on quarterly aggregated data. The analysis is conducted for 11 US variables, of which 3 observed at quarterly level, in a real-time context. The authors find that the monthly series provide important information in the short run, with significant RMSE reductions obtained with the mixed-frequency model. Moreover, the more intra-quarter information is available, the increasing the improvements.

5.4 Factor models

5.4.1 Applications of small-scale factor models

Small-scale factor models have been frequently employed in the literature to construct a coincident indicator, which is able to track the development of the economy in real-time. In what follows, we describe the main studies which employ small-scale factor models which extract an index and provide, in some of the cases presented, short-term forecasts of the real GDP growth.

As described also in Section 2.5.1, Mariano and Murasawa (2003) propose a new coincident index of business cycles that relies on both monthly and quarterly indicators. Stock and Watson (1989) construct a coincident index by applying maximum likelihood (ML) factor analysis to the four monthly coincident indicators. Mariano and Murasawa extend the Stock–Watson coincident index by including quarterly real GDP and compare the turning point detection performance of the two indexes. What they find is that the behavior of the common component is quite different from monthly real GDP, and more generally that the behavior of the common factor depends on the choice of the component indicators and therefore the monthly real GDP and the common factor component can have different turning points.

A different application of the Mariano and Murasawa model can be found in Frale et al. (2011). This paper proposes a new monthly indicator of the euro area economic conditions, EUROMIND, based on tracking real GDP on a monthly basis. The construction of this new monthly indicator of GDP is carried out indirectly through the temporal disaggregation of the value added by supply sectors from the output side and at the same time through the temporal disaggregation of the main components of the demand from the expenditure side. The two estimates are combined with optimal weights reflecting their precision. Therefore, the indicator is based on information at both monthly and quarterly level, modelled with a dynamic factor specification cast in state-space form, where computational efficiency is achieved by implementing univariate filtering and smoothing procedures, which also allows to handle the ragged-edge problem and other data irregu-

larities in a unified framework. The authors find satisfactory results in the application of the model to the sectorial data, while the results are less convincing on the expenditure side. In a second paper, Frale et al. (2010) introduce a modification in the model which consists of the introduction of a second common factor, capturing the contribution of the survey variables as coincident indicators. What they find is that the second factor loads significantly on the survey variables for the industry sector and for exports. Moreover, they also attempt to isolate the news content of each block of series by using a real-time database: the analysis of the revisions in the data indicates that the contribution of surveys is not negligible.

Camacho and Perez-Quiros (2010) construct a different coincident indicator of the Euro area economy, the so-called Euro-STING indicator, which evolves accordingly to the Euro area dynamics and it is also based on an extension of the dynamic factor model described in Mariano and Murasawa (2003). The authors accommodate the GDP releases (flash, first and second estimate) in a statistical model to examine the impact of preliminary announcements and data revisions in the accuracy of real time forecasting. They assume that monthly growth rates of quarterly series and monthly growth rates of hard indicators have a direct relation with the common factor, which represents the common component that drives the series dynamics. However, they model the relation of the common factor with the soft indicators in a different way, and precisely they relate the level of the soft indicators considered with the year-on-year common growth rate, written as the sum of current values of the common factors and the last eleven lagged values. In their empirical application, they deal with a relatively small number of indicators. What they find is that exploiting information within each quarter their model improves upon the accuracy of preliminary announcements in forecasting GDP and forecasting uncertainty decreases during the forecasting period. Moreover, not only hard indicators are useful in forecasting GDP but also business surveys are relevant, especially in the months when real activity data are not available yet due to publication lags.

What Camacho and Perez-Quiros (2010) do for the Euro area is closely related to the empirical work done by Evans (2005) for the US, who applies a model that allows for variable reporting lags and temporal aggregation to a wide range of US macroeconomic data releases. The author models the growth in GDP as the quarterly aggregate of an unobserved daily process and specifies the relationship between GDP, data releases on GDP growth and on other macroeconomic variables in such a way to accommodate the different timing of data releases. By writing the model in state-space form (similar to Mariano and Murasawa (2003) but accommodating for a more complex timing of the releases), Evans (2005) obtains a real-time estimate of GDP on a daily basis as a product of the Kalman filter applied to estimate the model. The results seem promising, showing that within quarter data releases contain useful information for real-time estimation and forecasting of GDP. However, gaps between the real-time estimates and ex-post GDP data remain both persistent and significant.

Another extension of the small-scale factor model of Mariano and Murasawa (2003) has been analyzed by Marcellino, Porqueddu and Venditti (2012), to investigate business cycle dynamics and for forecasting GDP growth at short-term horizons in the euro area. While so far the parameters of the model have been considered as constant, the authors consider time variation in the variance of the shocks, and they generalize the setup of Mariano and Murasawa (2003) to allow for continuous shifts in the volatility of the shocks both to the common components and to the single indicators. To do so, they model volatility shifts as independent random walks. Moreover, differently from the other studies, the model is estimated with a Bayesian technique, using a Gibbs sampling procedure. The authors use the model to evaluate the impact of macroeconomic releases on point and density forecast accuracy and on the width of forecast intervals, and they show how their setup allows to make a probabilistic assessment of the contribution of releases to forecast revisions. From a pseudo out of sample forecasting exercise, they find that stochastic volatility contributes to an improvement in density forecast accuracy.⁴

5.4.2 Bridging with factors

As we have seen in Section 2.5.2, Giannone et al. (2008) provide a framework that formalizes the updating of the nowcast and forecast of output and inflation as data are released throughout the month and that can be used to evaluate the marginal impact of new data releases on the precision of the now- and forecasts as well as the marginal contribution of different groups of variables. The framework they propose is adapted from the parametric dynamic factor model in state-space form proposed by Doz, Giannone and Reichlin (2011) that helps handling ragged-edge data. They extract monthly factors and use them in a state-space framework to forecast monthly GDP. The authors construct pseudo intra-month vintages according to a stylized data release calendar. As a new block of information is released, the factors are reestimated and the nowcast updated. The main finding is that information matters: the precision of the signal increases monotonically within the month, with the release of new data. Timeliness of the release and quality matters for decreasing uncertainty. Banbura et al. (2012) present a similar but updated and extended application that confirms this finding.

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. The authors consider the framework proposed by Giannone et al. (2008), but they also extract the factors following Forni et al. (2005), using generalized principal components which allow to take into account the ragged-edge structure of the dataset. In their paper, they focus on the Euro Area as a whole as well as on single Euro Area countries. The results obtained for the Euro Area countries show that models that exploit timely monthly releases fare better than

⁴Aastveit, et al (2011) develop a nowcasting system that combines forecasts from VAR models, bridging equations, and factor models. They also focus on density forecasts, for which their approach works well.

quarterly models, and among those, factor models do generally better than averages of bridge equations.

5.4.3 Factor models in a mixed-frequency state space representation

Enriching the model proposed by Giannone et al. (2008), Banbura and Rünstler (2011) and Banbura and Modugno (2010) make use of a large state-space model allows for joint estimation of GDP and the factors in a single framework.

Banbura and Rünstler (2011) develop measures for understanding the importance of an individual series in the forecasts: they derive the weights of the series in the forecast and use them to calculate their contributions to the forecasts. Moreover, they assess the gains in forecast precision due to certain series by measuring the increase in uncertainty once the series have been removed from the explanatory variables. Banbura and Rünstler (2011) use a factor model which implements the common factors as unobserved components in a state-space form, and integrate the monthly factor model and a forecast equation for quarterly GDP in a single state-space representation, using a mixed-frequency setup. The authors confirm the finding of the importance of intra-quarter information, showing that a quarterly AR model is clearly outperformed by the factor model. Moreover they find evidence that differences in the timeliness of data releases can have strong effects on the optimal weights of individual series in the forecast and on their contribution to forecast precision.

Banbura and Modugno (2010) extend the analysis in Banbura and Rünstler (2011) by augmenting the dataset by short history indicators and quarterly series. Moreover, they allow the model to have an AR(1) idiosyncratic component. In their pseudo real-time exercise, they recursively forecast the Euro area GDP on the basis of a large monthly dataset. Compared to the previous data employed in Banbura and Rünstler (2011), they introduce short-history indicators, the Purchasing Managers' Surveys, available only from mid 1997. The results obtained including these new short history monthly indicators are similar to the results obtained without including them, therefore it seems that these additional series do not improve the precision of the projections. Also with the explicit modelling of the serial autocorrelation of the idiosyncratic component the results do not improve significantly.

Moreover, Angelini et al. (2008) provide an out-of-sample evaluation of the method presented by Banbura and Rünstler (2011), and compare the forecasting performance of this approach to the one obtained by pooling the forecasts from different selected bridge equations. In order to evaluate the impact of new data releases on current GDP nowcast throughout the quarter, they update the model two times per month, measuring the accuracy of the forecasts computed using the information available at each date. The results they find show that the factor model forecast tracks GDP more accurately, most likely due to the fact that the factors take into account the information content of cross correlations across series.

5.4.4 Factor-MIDAS

Marcellino and Schumacher (2010) propose to merge factor models with the MIDAS approach, which allows to now- and forecast low frequency variables as GDP exploiting information in a large set of higher-frequency indicators. They consider three different MIDAS approaches - basic, smoothed and unrestricted - and the three alternative factor estimation methods that can account for unbalanced datasets, explained in Section 3.2, to have a total of nine Factor-MIDAS approaches. They then focus on German GDP and conduct now- and forecasting of quarterly GDP growth with a large set of timely monthly economic indicators. To relate Factor-MIDAS to the methods from the existing literature, the authors introduce two other approaches in the empirical comparison: a single-frequency factor model based on quarterly aggregated data and the integrated state-space approach by Banbura and Rünstler (2011). In terms of empirical results, MIDAS with exponentially distributed lag functions performs similarly to MIDAS with unrestricted lag polynomials. In most of the cases, the simplest MIDAS with one lag of the factors is the best performing. Autoregressive dynamics plays a minor role. As far as the choice of the factor estimation technique is concerned, there are not significant differences among the different estimation methods. All Factor-MIDAS nowcasts can improve over quarterly factor forecast based on time-aggregated data, while the results compared to the state-space approach are less clear-cut and depend on the forecast horizon.

Kuzin, Marcellino and Schumacher (2012) discuss nowcast pooling versus nowcasting of GDP growth for several countries, with single models in the presence of model uncertainty, with mixed-frequency and ragged-edge data. The nowcasts are based on MIDAS regressions with few indicators and Factor-MIDAS based on large datasets. The authors compare the performance of many alternative specifications with respect to alternative estimation methods, number of factors, indicators selected for MIDAS, the role of autoregressive dynamics. In their empirical analysis, they show that indicator models tend to outperform factor models in this ex-post evaluation. It is much more difficult to beat the benchmark when the models are selected based on information criteria or past performance. As a method to avoid the specification search, all the nowcasts and forecasts can be pooled together, using different selection schemes. This approach yields additional gains with respect to the factor specification based on past performance, in particular when all single-indicator and all Factor-MIDAS forecasts are combined together using inverse MSE weights, providing full support in favour of pooling for nowcasting and short-term forecasting.

6 Conclusions

In this paper we reviewed the literature concerning estimation and forecasting with mixed-sampling frequency and ragged-edge data. At the moment, temporal aggregation is still the predominant technique in the empirical applications: all data are sampled at the same

lower-frequency. In filtering the data so that the variables have all the same frequency, potentially useful information is discarded. Empirical studies show that mixed-frequency data matter, the use of the procedures that allow taking different frequencies and the timeliness of the data into account improve the forecasts.

One of the early approaches to deal with mixed-sampling frequency is bridge equations, still very common in central banks, where a dynamic equation is estimated between the low-frequency variable and time-aggregated indicators. Separate high-frequency models provide forecasts of the high-frequency indicators, and these forecasts are then aggregated and plugged into the bridge equation. Empirical studies with bridge models show that the exploitation of intra-period information reduces the forecasting error in the majority of the cases. Bridge equations are a useful instrument especially in nowcasting, since the more information becomes available, the more accurate they are in terms of RMSE.

A second strand of research is based on mixed frequency regressions, where a low-frequency variable is explained by high-frequency indicators using parsimonious distributed lag models. MIDAS models are in general restricted to a limited set of variables, and estimated via NLS. Different weighting functions have been used in the literature, but which one is better to use is not clear-cut and depends on the specific analysis. Initially MIDAS models were applied to financial data, investigating stock market returns or future volatility, but recently MIDAS regressions have been employed to forecast macroeconomic variables, providing promising results for short-term forecasting.

Another, line of research relies on the state-space framework, in connection with both factors and VAR representations. The state-space setup treats the low-frequency variable as a high-frequency series with missing observations. The use of the Kalman filter allows real-time filtering, i.e., taking quarterly economic activity explained by monthly indicators as example, it is possible to obtain an estimate of GDP growth in each month. However, because of the intensive computation required by this framework and a relatively high number of parameters to be estimated, only models with few variables can be implemented. Within the class of factor models, different factor estimation techniques are described in the literature to handle the ragged-edge problem.

Recently, mixed frequency factor models and MIDAS regressions have been merged into Factor-MIDAS, which allows to forecast a low-frequency variable, exploiting the monthly information available in large datasets and handling unbalanced data as typical in real-time. Various factor estimation methods have been employed, without significant differences in their forecasting performance. Evidence shows that taking into account higher frequency information and exploiting the most recent observations pays off: Factor-MIDAS outperform quarterly factor forecasts based on time-aggregated data.

In summary, there is consensus that exploiting data at different frequencies matters, but it is not clear which method is superior. State-space models are a system approach and allow the estimation of the missing high-frequency data thanks to the use of the Kalman filter. Within this class of models, different ways to estimate VAR parameters

or factors taking into account the unbalancedness of the datasets have been proposed, but the differences don't seem to be so pronounced. At the same time, the state-space approaches generally work only if the number of variables in the model is quite low, due to a dramatic increase in the number of parameters and associated complexity of the estimation. MIDAS models appear to be more robust to mis-specification compared both to bridge equation models and state-space approaches, and computationally simpler.

References

- [1] Aastveit, K., Karsten R., Jore A.S. (2011), Nowcasting GDP in Real-Time: A Density Combination Approach, mimeo, Norges Bank.
- [2] Altissimo, F., R. Cristadoro, M. Forni, M. Lippi, and G. Veronese (2010): "New Eurocoin: Tracking economic growth in real time", *The Review of Economics and Statistics*, 92(4), 1024-1034.
- [3] Andreou, E., E. Ghysels, and A. Kourtellis (2010): "Should macroeconomic forecasters use daily financial data and how?", University of Cyprus WP, 09-2010.
- [4] Angelini, E., J. Henry, and M. Marcellino (2006): "Interpolation and backdating with a large information set", *Journal of Economic Dynamics and Control*, 30(12), 2693-2724.
- [5] Angelini, E., G. Camba-Méndez, D. Giannone, G. Rünstler, and L. Reichlin (2008): "Short term forecasts of Euro Area GDP growth", Working Paper Series 949, European Central Bank.
- [6] Baffigi, A., R. Golinelli, and G. Parigi (2004): "Bridge models to forecast the Euro area GDP", *International Journal of Forecasting*, 20(3), 447-460.
- [7] Bai, J., E. Ghysels, and J. Wright (2011): "State space models and MIDAS regressions", *Econometric Reviews*, forthcoming.
- [8] Banbura M, Giannone D, Reichlin L. 2011. Nowcasting. In *Oxford Handbook on Economic Forecasting*, Clements MP, Hendry DF (eds). Oxford University Press: Oxford.
- [9] Banbura M, Giannone D, Modugno, M., Reichlin L. 2012. Now-casting and the real time data flow, CEPR WP 9112.
- [10] Banbura, M. and M. Modugno (2010): "Maximum likelihood estimation of factor models on data sets with arbitrary pattern of missing data", Working Paper Series 1189, European Central Bank.

- [11] Banbura, M., and G. Rünstler (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.
- [12] Barhoumi, K., S. Benk, R. Cristadoro, A. Den Reijer, A. Jakaitiene, P. Jelonek, A. Rua, G. Rünstler, 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.
- [13] Barhoumi, K., O. Darné, L. Ferrara, and B. Pluyaud (2011): "Monthly GDP forecasting using bridge models: application for the French economy", *Bulletin of Economic Research*, forthcoming.
- [14] Bencivelli, L., Marcellino, M. and Moretti, G.L., (2012), "Selecting predictors by Bayesian model averaging in bridge models", *Banca d'Italia WP*.
- [15] Breitung, J. and C. Schumacher (2008): "Real-time forecasting of German GDP based on a large factor model with monthly and quarterly data", *International Journal of Forecasting*, 24(3), 386-398.
- [16] Camacho, M. and G. Perez-Quiros (2010): "Introducing the EURO-STING: Short Term Indicator of Euro Area Growth", *Journal of Applied Econometrics*, 25(4), 663-694.
- [17] Carriero, A., Clark, T. and Marcellino, M. (2012), "Real-Time Nowcasting with Large Bridging Equations with Time-Varying Volatility", *mimeo*.
- [18] Clements, M. P., and A. B. Galvao (2008): "Macroeconomic forecasting with mixed-frequency data: Forecasting US output growth", *Journal of Business and Economic Statistics*, 26, 546-554.
- [19] Clements, M. P., and A. B. Galvao (2009): "Forecasting US output growth using leading indicators: An appraisal using MIDAS models", *Journal of Applied Econometrics*, 24(7), 1057-1217.
- [20] Diron, M. (2008): "Short-term forecasts of Euro Area real GDP growth. An assessment of real-time performance based on vintage data", *Journal of Forecasting*, 27(5), 371-390.
- [21] 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.
- [22] Chiu, C., B. Eraker, A. Foerster, T. Kim, and H. Seoane (2011): "Bayesian Mixed Frequency VAR's", *mimeo*.

- [23] Evans, N. D. (2005): "Where Are We Now? Real-Time Estimates of the Macroeconomy", *International Journal of Central Banking*, 1(2).
- [24] Forni, M., M. Hallin, M. Lippi and L. Reichlin (2005): "The Generalized Dynamic Factor Model: One-Sided Estimation and Forecasting", *Journal of the American Statistical Association*, 100, 830-840.
- [25] Foroni, C., and M. Marcellino (2012): "A Comparison of Mixed Frequency Approaches for Modelling Euro Area Macroeconomic Variables", *EUI ECO Working Papers*, 2012/07.
- [26] Foroni, C., M. Marcellino, and C. Schumacher (2012): "U-MIDAS: MIDAS regressions with unrestricted lag polynomials", *CEPR Discussion Papers*, 8828.
- [27] Forsberg, L., and E. Ghysels (2007): "Why do absolute returns predict volatility so well?", *Journal of Financial Econometrics*, 5(1), 31-67.
- [28] Frale, C., M. Marcellino, G. Mazzi, and T. Proietti (2010): "Survey data as coincident or leading indicators", *Journal of Forecasting*, 29(1-2), 109-131.
- [29] Frale, C., M. Marcellino, G. Mazzi, and T. Proietti (2011): "EUROMIND: A monthly indicator of the Euro Area economic conditions", *Journal of the Royal Statistical Society, Series A*, 174, 439-470.
- [30] Galvao, A. B. (2007): "Changes in Predictive Ability with Mixed Frequency Data", *Queen Mary University WP*, 595.
- [31] Ghysels, E. (2011): "Mixed frequency vector autoregressive models and the consequences of ignoring high frequency data", mimeo.
- [32] Ghysels, E., A. Rubia, and R. Valkanov (2009): "Multi-period forecasts of volatility: direct, iterated and mixed-data approaches". *University of North Carolina WP*.
- [33] Ghysels, E., P. Santa-Clara, and R. Valkanov (2004): "The MIDAS touch: MIXed DATA Sampling regression models", mimeo, Chapel Hill, N.C.
- [34] Ghysels, E., P. Santa-Clara, and R. Valkanov (2005): "There is a risk-return trade-off after all", *Journal of Financial Economics*, 76(3), 509-548.
- [35] Ghysels, E., P. Santa-Clara, and R. Valkanov (2006): "Predicting volatility: Getting the most out of return data sampled at different frequencies", *Journal of Econometrics*, 131(1-2), 59-95.
- [36] Ghysels, E., A. Sinko, and R. Valkanov (2007): "MIDAS regressions: Further results and new directions", *Econometric Reviews*, 26(1), 53-90.

- [37] Ghysels, E., and J. Wright (2009): "Forecasting professional forecasters", *Journal of Business and Economic Statistics*, 27(4), 504-516.
- [38] 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.
- [39] Guerin, P. and M. Marcellino (2011): "Markov-switching MIDAS models", *CEPR Discussion Papers*, 8234.
- [40] Harvey, A. C. (1989): "Forecasting structural time series models and the Kalman filter", Cambridge University Press.
- [41] Ingenito, R. and B. Trehan, 1996: "Using monthly data to predict quarterly output", *Economic Review*, Federal Reserve Bank of San Francisco, pages 3-11.
- [42] Koenig, E., S. Dolmas, and J. Piger (2003): "The use and abuse of real-time data in economic forecasting", *Review of Economics and Statistics*, 85(3), 618-628.
- [43] Kuzin, V., M. Marcellino, and C. Schumacher (2011): "MIDAS vs. mixed-frequency VAR: Nowcasting GDP in the Euro Area", *International Journal of Forecasting*, 27, 529-542.
- [44] Kuzin, V., M. Marcellino and C. Schumacher (2012): "Pooling versus model selection for nowcasting with many predictors: An application to German GDP", *Journal of Applied Econometrics*, forthcoming.
- [45] Lanning, S. (1986): "Missing observations: A simultaneous approach versus interpolation by related series", *Journal of Economic and Social Measurement*, 14(1), 155-163.
- [46] Marcellino, M. (1998): "Temporal disaggregation, missing observations, outliers, and forecasting: a unified non-model based procedure", *Advances in Econometrics*, 13, 181-202.
- [47] 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.
- [48] Marcellino, M., M. Porqueddu, and F. Venditti (2012): "Short-term GDP forecasting with a mixed frequency dynamic factor model with stochastic volatility", mimeo.
- [49] 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.

- [50] 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.
- [51] Mittnik, S., and P. A. Zadrozny (2005): "Forecasting quarterly German GDP at monthly intervals using monthly Ifo business conditions data", in *Ifo survey data in business cycle and monetary policy analysis*, ed. by J.-E. Sturm, and T. Wollmer-shÄauser, pp. 19-48. Heidelberg: Physica-Verlag.
- [52] Monteforte, M. and G. Moretti (2010): "Real-time forecasts of inflation: the role of financial variables", Banca D'Italia WP, 767.
- [53] Schorfheide, F., and D. Song (2011): "Real-time forecasting with a mixed frequency VAR", mimeo.
- [54] Stark, T. (2000): "Does current-quarter information improve quarterly forecasts for the U.S. economy?", Working Papers 00-2, Federal Reserve Bank of Philadelphia.
- [55] Stock, J., and M. Watson (1989): "New indexes of coincident and leading indicators," *NBER Macroeconomics Annual*, 4, 351-394.
- [56] 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.
- [57] Viefers, P. (2011): "Bayesian inference for the mixed-frequency VAR model", DIW Discussion Papers, 1172.
- [58] Wallis, K. (1986): "Forecasting with an econometric model: The 'ragged-edge' problem", *Journal of Forecasting*, 5, 1-13.
- [59] Wohlrabe, K. (2009): "Forecasting with Mixed-frequency Time Series Models, Ph.D. dissertation, University Munich.
- [60] Zadrozny, P. A. (1988): "Gaussian-Likelihood of continuous-time ARMAX models when data are stocks and flows at different frequencies", *Econometric Theory*, 4(1), 108-124.