European University Institute

DEPARTMENT OF ECONOMICS

EUI Working Papers

ECO 2009/32
DEPARTMENT OF ECONOMICS

MIDAS VS. MIXED-FREQUENCY VAR:
NOWCASTING GDP IN THE EURO AREA

Vladimir Kuzin, Massimiliano Marcellino and Christian Schumacher

EUROPEAN UNIVERSITY INSTITUTE, FLORENCE
DEPARTMENT OF ECONOMICS

MIDAS vs. Mixed-Frequency VAR: Nowcasting GDP in the Euro Area

VLADIMIR KUZIN,
MASSIMILIANO MARCELLINO
and
CHRISTIAN SCHUMACHER

EUI Working Paper **ECO** 2009/32

ISSN 1725-6704

Printed in Italy
European University Institute
Badia Fiesolana
I – 50014 San Domenico di Fiesole (FI)
Italy
www.eui.eu
cadmus.eui.eu

MIDAS vs. mixed-frequency VAR: Nowcasting GDP in the Euro Area[1]

Vladimir Kuzin
DIW Berlin

Massimiliano Marcellino
European University Institute, Università Bocconi and CEPR

Christian Schumacher
Deutsche Bundesbank

[1]This paper represents the authors' personal opinions and does not necessarily reflect the views of the Deutsche Bundesbank. We are grateful to the two Referees and seminar participants at University of Bielefeld, in particular Harry Haupt, for helpful comments and discussions. The codes for this paper were written in Matlab. Some functions were taken from the Econometrics Toolbox written by James P. LeSage from `www.spatial-econometrics.com`. Further codes were kindly provided by Arthur Sinko from `www.unc.edu/~sinko/midas.zip`. E-Mail: Vladimir Kuzin: `vkuzin@diw.de`; Massimiliano Marcellino: `massimiliano.marcellino@eui.eu`; Christian Schumacher: `christian.schumacher@bundesbank.de`.

Abstract

This paper compares the mixed-data sampling (MIDAS) and mixed-frequency VAR (MF-VAR) approaches to model specification in the presence of mixed-frequency data, e.g., monthly and quarterly series. MIDAS leads to parsimonious models based on exponential lag polynomials for the coefficients, whereas MF-VAR does not restrict the dynamics and therefore can suffer from the curse of dimensionality. But if the restrictions imposed by MIDAS are too stringent, the MF-VAR can perform better. Hence, it is difficult to rank MIDAS and MF-VAR a priori, and their relative ranking is better evaluated empirically. In this paper, we compare their performance in a relevant case for policy making, i.e., nowcasting and forecasting quarterly GDP growth in the euro area, on a monthly basis and using a set of 20 monthly indicators. It turns out that the two approaches are more complementary than substitutes, since MF-VAR tends to perform better for longer horizons, whereas MIDAS for shorter horizons.

JEL Classification Codes: E37, C53

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

1 Introduction

The development of econometric models based on mixed frequency data has attracted considerable attention recently. In particular, the mixed-data sampling (MIDAS) approach proposed by Ghysels, Santa-Clara and Valkanov (2004) and Ghysels, Sinko and Valkanov (2007) has proven useful for different forecasting purposes. MIDAS can be regarded as time-series regressions that allow the regressand and regressors to be sampled at different frequencies, where distributed lag polynomials are used to ensure parsimonious specifications. Whereas MIDAS has been initially used for financial applications, e.g. Ghysels, Santa-Clara and Valkanov (2006), it has been employed to forecast macroeconomic time series, in particular quarterly GDP with monthly indicators, in recent applications by Clements and Galvão (2008, 2009), Marcellino and Schumacher (2008), and Wohlrabe (2009).

In this paper, we compare the MIDAS approach to a mixed-frequency VAR (MF-VAR) model as proposed by Zadrozny (1988), Mittnik and Zadrozny (2005) and Mariano and Murasawa (2007). The MF-VAR is a VAR operating at the highest sampling frequency of the time series to be included in the model. Low-frequency variables are interpolated according to their stock-flow nature implying specific time-aggregation schemes. The high-frequency VAR together with the time-aggregation restriction can be cast in state-space form and estimated by maximum likelihood. In this framework, the Kalman filter can tackle missing values at the end of the sample, and take into account the mixed-frequency nature of the data.

Compared to single-equation MIDAS, MF-VAR is a system approach that jointly explains indicators and predictant without imposing a-priori restrictions on the dynamics. This can be an advantage when few variables are modelled, their dynamics is limited, and the VAR provides a good approximation to the data generating process (DGP). Otherwise, MIDAS can represent a more robust forecasting device. In addition, due to its single equation specification, a direct forecasting approach is preferable for MIDAS, while an iterated scheme is a more natural choice for the MF-VAR since it is cast in state-space form and iterated forecasts are directly provided by the Kalman filter. For a discussion of direct versus iterated forecasting see, e.g., Marcellino, Stock and Watson (2006) and Chevillon and Hendry (2005).

It is difficult to rank the MIDAS and MF-VAR approaches based purely on theoretical considerations since, as mentioned, their relative merits depend on the DGP. Therefore, their performance is better assessed in specific economic applications, and in this paper we focus on nowcasting and forecasting quarterly euro area GDP growth using a set of monthly indicators, a relevant issue also from the economic policy perspective.

In our application, we compare various specifications of MIDAS and MF-VAR models with single indicators, as well as combinations of these models. In addition, we take into account the different availability of monthly indicators that emerges from different statistical publication lags. The nowcast and forecast comparison is based on the relative mean-squared errors (MSE) at different horizons, and the analysis is conducted recursively, in a pseudo real-time way.

Our main finding is that in the case of euro area GDP growth, the two approaches are more complementary than substitutes, since MF-VAR tends to perform better for longer horizons, whereas MIDAS for shorter horizons.

The paper proceeds as follows. Section 2 provides a description of the MIDAS and MF-VAR approaches, as well as a discussion of their relative advantages. Section 3 presents the empirical results on nowcasting and forecasting quarterly euro area GDP growth with a set of monthly indicators. Section 4 summarizes our main findings and concludes.

2 Nowcasting quarterly GDP with ragged-edge data

In this paper we focus on quarterly GDP growth, which is denoted as y_{t_q} where t_q is the quarterly time index $t_q = 1, 2, 3, \ldots, T_q^y$ with T_q^y as the final quarter for which GDP is available. GDP growth can also be expressed at the monthly frequency by setting $y_{t_m} = y_{t_q} \forall t_m = 3t_q$ with t_m as the monthly time index. Thus, GDP growth y_{t_m} is observed only in months $t_m = 3, 6, 9, \ldots, T_m^y$ with $T_m^y = 3T_q^y$. The aim is to nowcast or forecast GDP h_q quarters ahead, or $h_m = 3h_q$ months ahead, yielding a value for $y_{T_m^y+h_m}$.

Nowcasting means that in a particular calender month, we do not observe GDP for the current quarter. It can even be the case that GDP is only available with a delay of two quarters. In April, for example, Euro Area GDP is only available for the fourth quarter of the previous year, and a nowcast for second quarter GDP requires $h_q = 2$. Thus, if a decision maker requests an estimate of current quarter GDP, the forecast horizon has to be set sufficiently large in order to provide the appropriate figures. For further discussion on nowcasting, see e.g. Giannone et al. (2008).

In this Section we assume, for the sake of exposition, that the information set for now- and forecasting includes one stationary monthly indicator x_{t_m} in addition to the available observations of GDP. The time index t_m denotes a monthly sampling frequency of x_{t_m} for $t_m = 1, 2, 3, \ldots, T_m^x$, where T_m^x is the final month for which an observation is available. Usually, T_m^x is larger than $T_m^y = 3T_q^y$, as monthly observations for many relevant macroeconomic indicators are earlier available than GDP observations. The forecast for GDP is denoted as $y_{T_m^y+h_m|T_m^x}$, as we condition the forecast on information available in month T_m^x, which also includes GDP observations up to T_q^y in addition to the indicator observations up to T_m^x with $T_m^x \geq T_m^y = 3T_q^y$.

2.1 The MIDAS approach

To forecast quarterly GDP using monthly indicators, we rely on the mixed-data sampling (MIDAS) approach as proposed by Ghysels and Valkanov (2006), Ghysels et al. (2007), and Andreou et al. (2009a). The MIDAS regression approach is a direct forecasting tool.

The dynamics of the indicators and joint dynamics between GDP and the indicators are not explicitly modelled. Rather, MIDAS directly relates future GDP to current and lagged indicators, thus yielding different forecasting models for each forecast horizon, see e.g. Marcellino, Stock and Watson (2006) as well as Chevillon and Hendry (2005) for detailed discussions of this issue in the single-frequency case.

The forecast model for forecast horizon h_q quarters with $h_q = h_m/3$ is

$$y_{t_q+h_q} = y_{t_m+h_m} = \beta_0 + \beta_1 b(L_m, \boldsymbol{\theta}) x^{(3)}_{t_m+w} + \varepsilon_{t_m+h_m}, \tag{1}$$

where $w = T^x_m - T^y_m$ and $b(L_m, \boldsymbol{\theta})$ is some lag polynomial

$$b(L_m, \boldsymbol{\theta}) = \sum_{k=0}^{K} c(k, \boldsymbol{\theta}) L^k_m \tag{2}$$

with the monthly lag operator L_m defined as $L_m x_{t_m} = x_{t_m-1}$. In the MIDAS approach, quarterly GDP $y_{t_q+h_q}$ is directly related to the indicator $x^{(3)}_{t_m+w}$ and its lags, where $x^{(3)}_{t_m}$ is skip sampled from the monthly observations of x_{t_m} in the following way. The superscript three indicates that every third observation starting from the t_m-th one is included in the regressor $x^{(3)}_{t_m}$, thus $x^{(3)}_{t_m} = x_{t_m} \,\forall\, t_m = \ldots, T^x_m - 6, T^x_m - 3, T^x_m$. Lags of the monthly factors are treated accordingly, e.g. the k-th lag $x^{(3)}_{t_m-k} = x_{t_m-k} \,\forall\, t_m = \ldots, T^x_m - k - 6, T^x_m - k - 3, T^x_m - k$. In the time index of $x^{(3)}_{t_m+w}$, w is equal to the number of monthly periods, the monthly indicator is earlier available than GDP. Thus, we take into account that a monthly indicator is typically available within the quarter for which no GDP figure is available, see Clements and Galvão (2008, 2009).

One of the main issues in MIDAS approach is a parsimonious parametrization of the lagged coefficients $c(k, \theta)$. Since the regressors $x^{(3)}_{t_m}$ are observed at a higher frequency than y_{t_q}, an adequate modelling often requires inclusion of many lags into the regression equation, which easily leads to overparameterization in unrestricted linear case. Ghysels et al. (2007) discuss several non-linear weighting schemes for $c(k, \theta)$. The first scheme is exponential Almon lag and possesses the following form

$$c(k, \boldsymbol{\theta}) = \frac{\exp(\theta_1 k + \ldots + \theta_Q k^Q)}{\sum_{k=0}^{K} \exp(\theta_1 k + \ldots + \theta_Q k^Q)}. \tag{3}$$

This functional form is quite flexible and allows for various shapes with only a few parameters, see Ghysels et al. (2007) for a more detailed discussion. An alternative weighting scheme includes the Beta function exploiting its well-known flexibility in the presence of only two parameters. Obviously, there are also other possible weighting schemes available in the literature, i.e. step functions, however, a detailed discussion is beyond the scope of the current paper, see Ghysels et al. (2007) for further reading.

In the subsequent empirical applications we follow Clements and Galvão (2008, 2009) as well as Ghysels et al. (2005) and employ the Almon lag weighting scheme with two parameters $\theta = \{\theta_1, \theta_2\}$. The exponential lag function $b(L_m, \boldsymbol{\theta})$ with quadratic expansion provides a parsimonious and easy way to consider monthly lags of the indicators as we can allow for large K to approximate the impulse response function of GDP to the indicators. The longer the lead-lag relationship in the data is, the less MIDAS suffers from sampling uncertainty compared with the estimation of unrestricted lags, where the number of coefficients increases with the lag length.

The MIDAS model can be estimated using nonlinear least squares (NLS) in a regression of y_{t_m} onto $x^{(3)}_{t_m-k}$, yielding coefficients $\widehat{\theta}_1$, $\widehat{\theta}_2$, $\widehat{\beta}_0$ and $\widehat{\beta}_1$. The forecast is given by

$$y_{T^y_m+h_m|T^x_m} = \widehat{\beta}_0 + \widehat{\beta}_1 b(L_m, \widehat{\boldsymbol{\theta}}) x_{T^x_m}. \tag{4}$$

Note that MIDAS is h-dependent, and thus has to be reestimated for multi-step forecasts. The same holds when new statistical information becomes available. For example, each month, new observations for the indicator are released, whereas GDP observations are released only once a quarter. Thus, also w changes from month to month, which also makes re-estimation necessary.

Autoregressive MIDAS As an extension to the basic MIDAS approach, Clements and Galvão (2008) consider autoregressive dynamics in the MIDAS approach. In particular, they propose the model

$$y_{t_m+h_m} = \beta_0 + \lambda y_{t_m} + \beta_1 b(L_m, \boldsymbol{\theta})(1 - \lambda L^3_m) x^{(3)}_{t_m+w} + \varepsilon_{t_m+h_m}. \tag{5}$$

The autoregressive coefficient λ is not estimated unrestrictedly to rule out discontinuities of the impulse response function of $x^{(3)}_{t_m}$ on $y_{t_m+h_m}$, see the discussion in Ghysels et al. (2007), pp. 60. The restriction on the coefficients is a common-factor restriction to ensure a smooth impulse response function, see Clements and Galvão (2008). The AR coefficient λ can be estimated together with the other coefficients by NLS. As an AR model is often supposed to be an appropriate benchmark specification for GDP, the extension of MIDAS might give additional insights in which direction the other MIDAS approaches considered so far might be improved. Henceforth, we denote this approach as 'AR-MIDAS', whereas we denote MIDAS without AR terms just as 'MIDAS'.

2.2 The mixed-frequency VAR

In contrast to the MIDAS approach and in line with a conventional VAR model based on single-frequency data, the MF-VAR model specifies the joint dynamics of monthly GDP, which is obtained from quarterly GDP by time disaggregation, and the monthly indicator. Following the notation of Mariano and Murasawa (2003, 2007), the disaggregation of quarterly GDP growth y_{t_m} into unobserved month-on-month GDP growth $y^*_{t_m}$ is based on the aggregation relation

$$y_{t_m} = \frac{1}{3}y^*_{t_m} + \frac{2}{3}y^*_{t_m-1} + y^*_{t_m-2} + \frac{2}{3}y^*_{t_m-3} + \frac{1}{3}y^*_{t_m-4}, \tag{6}$$

which holds for $t_m = 3, 6, 9, \ldots, T^y_m$, because GDP is observed only every third month of each quarter. The aggregation assumption represents the flow nature of GDP and allows for a linear state-space representation, see Mariano and Murasawa (2003) or Giannone et al. (2008). The latent month-on-month GDP growth $y^*_{t_m}$ and the corresponding monthly indicator x_{t_m} are then assumed to follow a bivariate VAR(p) process

$$\mathbf{\Phi}(L_m)\begin{pmatrix} y^*_{t_m} - \mu^*_y \\ x_{t_m} - \mu_x \end{pmatrix} = \mathbf{u}_{t_m}, \tag{7}$$

with $\mathbf{\Phi}(L_m) = \sum_{i=1}^p \mathbf{\Phi}_i L^i_m$ and $\mathbf{u}_{t_m} \sim \mathrm{N}(\mathbf{0}, \mathbf{\Sigma})$.

State-space representation To obtain the state-space representation of the MF-VAR, we define the state vector

$$\mathbf{s}_{t_m} = \begin{pmatrix} \mathbf{z}_{t_m} \\ \vdots \\ \mathbf{z}_{t_m-4} \end{pmatrix}, \quad \mathbf{z}_{t_m} = \begin{pmatrix} y^*_{t_m} - \mu^*_y \\ x_{t_m} - \mu_x \end{pmatrix} \tag{8}$$

consisting of demeaned monthly GDP growth with mean μ^*_y, and the monthly indicator demeaned with μ_x, as well as their lags. Transforming (7) into companion form and combining the latter with the aggregation constraint (6), we obtain the corresponding state-space form as

$$\mathbf{s}_{t_m+1} = \mathbf{A}\mathbf{s}_{t_m} + \mathbf{B}\mathbf{v}_{t_m}, \tag{9}$$

$$\begin{pmatrix} y_{t_m} - \mu_y \\ x_{t_m} - \mu_x \end{pmatrix} = \mathbf{C}\mathbf{s}_{t_m}, \tag{10}$$

where $\mathbf{v}_{t_m} \sim \mathrm{N}(\mathbf{0}, \mathbf{I}_2)$, and $\mu_y = 3\mu^*_y$ holds. Our experience shows that the mean parameters μ_y and μ_x are often quite difficult to estimate in the state-space framework. For this reason, we work with demeaned series for estimation. The system matrices are

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{bmatrix}, \quad \mathbf{A}_1 = \begin{bmatrix} \mathbf{\Phi}_1 & \ldots & \mathbf{\Phi}_p & \mathbf{0}_{2\times 2(5-p)} \end{bmatrix}, \quad \mathbf{A}_2 = \begin{bmatrix} \mathbf{I}_8 & \mathbf{0}_{8\times 2} \end{bmatrix}, \tag{11}$$

$$\mathbf{B} = \begin{bmatrix} \mathbf{\Sigma}^{1/2} \\ \mathbf{0}_{8\times 2} \end{bmatrix}, \quad \mathbf{C} = \begin{bmatrix} \mathbf{H}_0 & \ldots & \mathbf{H}_4 \end{bmatrix}, \tag{12}$$

where matrix $\mathbf{C}$ contains the lag polynomial $\mathbf{H}(L_m) = \sum_{i=0}^{4} \mathbf{H}_i L_m^i$ that is defined as

$$\mathbf{H}(L_m) = \begin{bmatrix} 1/3 & 0 \\ 0 & 1 \end{bmatrix} + \begin{bmatrix} 2/3 & 0 \\ 0 & 0 \end{bmatrix} L_m + \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} L_m^2 + \begin{bmatrix} 2/3 & 0 \\ 0 & 0 \end{bmatrix} L_m^3 + \begin{bmatrix} 1/3 & 0 \\ 0 & 0 \end{bmatrix} L_m^4, \quad (13)$$

according to the aggregation constraint (6). For notational convenience, we consider only $p \leq 4$ for $\mathbf{A}$ and $\mathbf{B}$, however, the representation for $p > 4$ can be derived in a straightforward manner by modifying the state vector and system matrices accordingly.

Missing observations and estimation The state-space model consisting of (9) and (10) can be estimated with maximum likelihood techniques or the expectation-maximization (EM) algorithm, where we have to take into account missing observations due to publication lags and the low-frequency nature of GDP. We follow Mariano and Murasawa (2003, 2007) and first replace all missing values with zeros, where the missing values are assumed to be realizations of some iid standard normal random variable. Second, the signal equation (10) is also modified accordingly: for the first two months of each quarter, the upper row of $\mathbf{C}$ is set to zero and a standard normal error term is added, for details see Mariano and Murasawa (2003, 2007). Then, the EM algorithm is employed for parameter estimation.

Forecasting After estimation, the forecasting of GDP growth is done by means of the Kalman smoother. The application of the Kalman smoother ensures that all timely observations from the monthly indicator are taken into account. Whereas quarterly GDP is available up to $T_m^y = 3T_q^y$, we have monthly indicator observations up to T_m^x with difference in publication lag of $w = T_m^x - T_m^y$. Although GDP for a particular quarter is not available, the smoother considers the monthly indicator observations of the current quarter. Thus, both the MF-VAR and the MIDAS approach can consider timely within-quarter observations for nowcasting. For months without indicator observations, the Kalman smoother operates equally as the Kalman filter, as no updating step can be carried out. As the smoother is applied iteratively, we obtain iterative multi-step forecasts for the MF-VAR model, according to the definitions from Chevillon and Hendry (2005). As GDP growth has been recursively demeaned prior to estimation, the mean is added back to the MF-VAR forecast to obtain the final forecast, which can be compared to actual GDP growth.

2.3 Discussion of MIDAS and MF-VAR

Both the MF-VAR and the MIDAS approaches can tackle the mixed-frequency nature of the data, and both can exploit timely indicator observations that are also available at higher frequency than GDP. However, in general, there are marked differences between the two methods:

- MIDAS is a single-equation approach whereas MF-VAR is a system approach that explains both GDP and the indicator. As such, misspecification in one equation can affect estimation and forecast accuracy of the other model equations. However, forecasts of the monthly indicators can be of interest by themselves.

- MIDAS has a sparse parameterization, whereas MF-VAR suffers more from the curse of dimensionality. For example, with MIDAS using the Almon lag with quadratic expansion, adding a monthly variable to the predictors requires only 3 more coefficients (θ_1, θ_2, and β) to be estimated in the lag polynomial, whereas a VAR(p) with N variables requires N^2p coefficients of the VAR lag polynomial to be estimated. On the other hand, the MIDAS restrictions on the lag polynomial could be invalid, whereas the coefficients of VAR polynomials are estimated unrestrictedly.

- In terms of flexibility, the performance of MIDAS could also be improved by the inclusion of higher frequency information, e.g., in the form of daily financial data. For example, Ghysels and Wright (2008) find that daily financial information included in MIDAS models is useful to predict the median GDP growth forecast in the Survey of Professional Forecasters, see also Andreou, Ghysels and Kourtellos (2009b) and Wohlrabe (2009) for applications. On the other hand, including daily data in MF-VAR adds substantially to their computational complexity. In addition, a VAR specification could be not appropriate for high frequency data.

- MIDAS is a direct multi-step forecast device, in the sense that the left hand side variable coincides with the variable of interest from the forecasting point of view, so that the model changes with the forecasting horizon. Instead, MF-VAR provides iterative forecasts, in the sense that the same model is iterated forward to produce the forecast of interest. Thus, the long-lasting discussion of the relative merits of direct versus iterative forecasting also applies here. Marcellino, Stock and Watson (2006) and Chevillon and Hendry (2005) are recent contributions, see Bhansali (2002) for a survey. The literature shows that there are arguments in favour of both approaches and, generally, the direct approach seems to dominate only in case of substantial misspecification.[1]

- In Ghysels and Valkanov (2006) it is shown how the MIDAS can be regarded as an approximation to a general dynamic linear model, in their case a high-frequency VAR(1), where the low-frequency variable is a stock variable. Thus, in case the true high-frequency DGP behind the data is close to a VAR model, we can expect the MF-VAR to perform better than MIDAS, depending on the dimension and parsimony of the DGP.

[1] The direct and iterated approaches have been typically compared with single frequency data. In the case of mixed frequency data, a direct forecast could be based either on the low frequency data only or on a MIDAS regression. See Ghysels, Rubia and Valkanov (2009) for an interesting comparison of the two approaches in the context of volatility forecasting.

- Ghysels and Wright (2008) note that the Kalman filter, underlying the MF-VAR, allows for real-time filtering. In other words, with the Kalman filter it is possible to obtain an estimate of the expected value of GDP growth in each month, while with MIDAS one can only obtain a monthly update of the future expected quarterly realization.

This discussion suggests that we cannot expect one approach to be clearly superior than the other one for any DGP, and either approach could dominate in a specific empirical application. Therefore, the relative advantages of MIDAS and MF-VAR should be evaluated empirically on a case-by-case basis, and in the next Section we focus on a policy-relevant case, i.e., nowcasting and forecasting quarterly GDP growth in the euro area, on a monthly basis.

3 Now- and forecasting Euro Area GDP with MIDAS and MF-VAR

The empirical comparison will be carried out in a recursive pseudo real-time context. In subsection 3.1, we describe the design of the exercise, the data used and the specification of the models. In the subsequent sections, we present and discuss the empirical results.

3.1 Design of the nowcast and forecast comparison exercise

Data The dataset contains Euro Area quarterly GDP growth from $1992Q1$ until $2008Q1$ and about 20 monthly indicators until $2008M06$. In particular, we consider industrial production by sector, survey on consumer sentiment, and business climate, raw material price indices, car registrations, interest rates, and monetary aggregates. More information about the data can be found in Appendix A.

The dataset is a final dataset. It is not a real-time dataset and does not contain vintages of data, so that we cannot discuss the role of revisions on the relative forecasting accuracy here. However, we do not expect any major changes in the results from the use of real-time vintages, since the data revisions are typically small after 2000, see e.g. Marcellino and Musso (2008) for euro area GDP growth. Furthermore, many empirical findings such as Bernanke and Boivin (2003) and Schumacher and Breitung (2008) suggest that data revisions do not affect forecast accuracy considerably. However, we take into account another specific characteristic of multivariate data in real time, namely the different availability of variables due to publication lags. These differences in availability of data lead to certain patterns of missing values at the end of every recursive sample, and recent papers find that accounting for this rather than using artificially balanced samples has a considerable impact on forecast accuracy, see Giannone et al. (2008), Schumacher and Breitung (2008), for example. In our paper, to consider the availability of the data at the end of each subsample, we follow Giannone et al. (2008), Marcellino and Schumacher (2008), amongst others, and replicate the availability of data in pseudo real-time from a final vintage of data. When downloading the final data - the download date for the data used here was 11th July 2008 -, we observe the data availability pattern in terms of the missing values at the end of the data sample. For example, at the beginning of July 2008, we observe interest rates until June 2008, thus there is only one missing value at the end of the sample, whereas industrial production is available up to April 2008, implying three missing values. For each time series, we store the missing values at the end of the sample. Under the assumption that these patterns of data availability remain stable over time, we impose the same missing values pattern at each point in time of the recursive experiment. Thus, we shift the missing values back in time to mimic the availability of information as in real time.

Nowcast and forecast design To evaluate the performance of the models, we carry out recursive estimation and nowcasting, where the full sample is split into an evaluation sample and an estimation sample, which is recursively expanded over time. The evaluation sample is between $1999Q2$ and $2008Q1$, providing 9 years for comparison. For each of these quarters, we want to compute nowcasts and forecasts depending on different monthly information sets. For example, for the initial evaluation quarter $1999Q2$, we want to compute a nowcast in June 1999, one in May, and April, whereas the forecasts are computed from March 1999 backwards in time accordingly. Thus, we have three nowcasts computed at the beginning of each of the intra-quarter months. Concerning the forecasts, we present results up to two quarters ahead. Thus, again for the initial evaluation quarter $1999Q2$, we have six forecasts computed based on information available in December 1998 up to information available in March 1999. Overall, we have nine projections for each GDP observation of the evaluation period, depending on the monthly information available to make the projection.

The estimation sample depends on the information available at each period in time when computing the now- and forecasts. Assume again we want to nowcast GDP for $1999Q2$ in June 1999, then we have to identify the time series observations available at that period in time. For this purpose, we exploit the ragged-edge structure from the end of the full sample of data, as discussed in the previous subsection. For example, for the nowcast GDP for $1999Q2$ made in June 1999, we know from our full sample that at each period in time, we have one missing value for interest rates and three missing values of industrial production. These missing values are imposed also for the period June 1999, thus replicating the same pattern of data availability. We do this accordingly in every recursive subsample to determine the pseudo real-time final observation of each time series. To replicate the publication lags of GDP, we exploit the fact that in the Euro Area GDP of the previous quarter is available at the beginning of the third month of the next quarter. Note that we reestimate all forecast models recursively when new information becomes available, so that the estimated coefficients are allowed to change over time. For each evaluation period, we compute nine now- and forecasts depending on the available information. To compare the nowcasts with the realisations of GDP growth, the mean-squared error (MSE) is employed.

Lag length specification For estimating the MF-VAR model, a lag order determination is required. For this purpose, we apply the Bayesian information criterion (BIC) with a maximum lag order of $p = 4$ months. Experimenting with higher lag orders did not affect the main results, as the chosen lag lengths are usually very small with only one or two lags in most of the cases. Concerning the specifications of MIDAS and AR-MIDAS, we use a large variety of initial parameter specifications, and compute the residual sum of squares (RSS) from (1) and (5), respectively. The parameter set with the smallest RSS then serves as the initial parameter set for NLS estimation. The parameters of the exponential lag function are restricted to $\theta_1 < 5$ and $\theta_2 < 0$. The maximum number of lags chosen for MIDAS is $K = 12$ months. Again, experimenting with higher lag orders did not affect the main results.

3.2 Empirical results

Individual models Below, we present a selection of well-performing models for different now- and forecast horizons h_m with respect to the relative MSE. To compute the relative MSE in this application, the benchmark forecast is produced by an AR model of GDP growth, which was estimated by direct estimation recursively, and where lag length was specified by BIC.

Starting from the full set of indicators in Appendix A, we focus on those representatives from surveys, industrial statistics, financial data and others that have some information content for future GDP growth at least at some forecast horizons. In particular, we compare surveys on production expectations (abbreviated in the table by: prod exp), order books (ord book), and consumer confidence (cons conf). We also look at industrial production of capital goods producers (prod cap), the 3-month EURIBOR money intermarket rate (is3m), yields on 10-year government bonds (il10), as well as the HWWA industrial raw material price index (hwwi ind) and passenger car registrations (car pass). In Table 1, we provide results of MIDAS without autoregressive terms, AR-MIDAS and MF-VAR for each of these indicators. A general finding related to all indicators and forecast models is that the AR benchmark can only be outperformed by a few models and not for all forecast horizons. Furthermore, the results show that MIDAS without AR terms in most of the cases yields uninformative forecasts with relative MSE larger than one for almost all horizons. However, incorporating AR terms in AR-MIDAS can in many cases improve the forecast performance. Particularly for short horizons $h_m = 1, 2, 3$, the AR-MIDAS shown can provide relative MSE smaller than one. At larger horizons, the evidence is mixed, however.

Table 1: Forecasting performance for quarterly GDP growth of selected individual mixed-frequency models measured by MSE of the corresponding indicator relative to MSE of the benchmark

model		horizon h_m								
		1	2	3	4	5	6	7	8	9
prod exp	midas	0.94	0.81	0.91	1.01	0.94	1.00	1.02	0.92	1.05
	ar-midas	0.81	0.81	0.84	0.89	0.92	0.94	0.89	1.00	1.01
	mf-var	1.01	0.89	0.91	0.82	0.75	0.97	0.97	1.00	1.04
ord book	midas	1.25	1.01	1.06	1.12	1.00	1.02	1.15	1.16	1.12
	ar-midas	0.89	0.89	0.91	0.94	1.07	1.07	1.07	1.04	1.04
	mf-var	1.34	1.08	0.78	1.02	0.92	0.98	1.16	0.87	0.91
cons conf	midas	1.20	1.01	1.10	1.12	1.05	1.10	1.25	1.35	1.34
	ar-midas	0.90	0.90	0.92	0.94	1.42	1.39	1.35	1.34	1.34
	mf-var	0.87	0.75	0.99	1.03	0.99	1.08	1.07	0.96	1.00
prcap	midas	1.26	1.15	1.10	1.37	1.09	1.04	1.23	1.02	0.88
	ar-midas	0.88	0.91	1.05	1.15	0.93	0.94	1.09	0.93	0.93
	mf-var	1.45	1.24	1.23	1.16	1.02	1.03	1.04	0.94	0.93
is3m	midas	1.19	0.91	1.00	1.00	0.91	0.87	1.37	1.21	1.05
	ar-midas	0.88	1.00	0.96	0.96	0.93	0.87	1.13	1.01	1.06
	mf-var	0.89	0.90	0.93	1.03	0.93	0.96	1.00	0.92	0.92
il10	midas	1.47	1.01	1.03	1.10	1.38	1.01	1.26	1.14	0.94
	ar-midas	0.98	0.87	0.93	0.98	1.04	0.90	1.08	1.00	0.89
	mf-var	0.97	0.93	1.03	1.11	1.02	1.04	1.04	0.92	0.93
hwwa ind	midas	0.91	0.77	0.84	0.92	0.79	0.85	0.87	0.87	0.93
	ar-midas	0.97	0.72	0.75	0.79	0.78	0.86	0.92	0.92	0.97
	mf-var	1.17	0.99	1.05	1.07	0.95	1.02	1.02	0.91	0.93
carpass	midas	1.41	1.13	1.39	1.25	1.24	1.19	1.07	1.09	1.12
	ar-midas	0.87	0.89	0.89	0.94	0.92	0.90	1.01	0.93	0.98
	mf-var	1.24	1.21	1.22	1.21	1.05	1.04	1.04	0.93	0.93

Note: We use the recursively estimated AR model as benchmark forecast. The first two columns in the table include the indicator name and model type (MIDAS, AR-MIDAS or MF-VAR). For the meaning of abbreviations of the particular indicators, see Appendix A. Details on the forecasting exercise are reported in Section 3.1.

To illustrate the behaviour of AR-MIDAS further, we present in Figure 1 the estimated weights attached to the indicators and their lags in AR-MIDAS based on the full sample of data. Most of the weights decay quickly, and with a few exceptions (hwwa ind), there

Figure 1: Weights of AR-MIDAS exponential lag polynomial

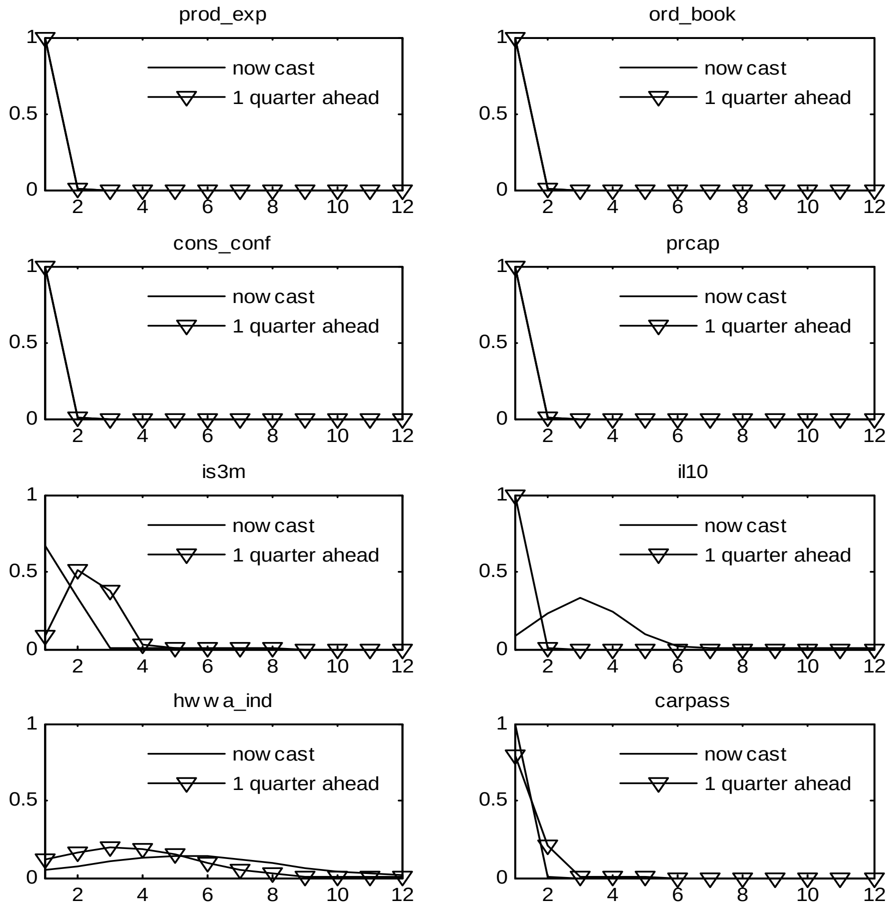

Note: The figure shows weights derived from estimated AR-MIDAS for the selected variables from table 1. The data used is full dataset. It corresponds to the nowcast situation (straight line) in the third month of a quarter, for which GDP is not available. The forecast one quarter ahead (line with triangle symbol) shows the weights corresponding to the forecast also made in third month of the final quarter. For the meaning of abbreviations of the particular indicators, see Appendix A.

is little mass at lag $k = K = 12$. In AR-MIDAS models with the surveys (prod exp, ord book, cons conf), weight is attached to the most recent monthly observation ($k = 0$) inducing a very parsimonious model. With respect to the other indicators, there is more variety in the estimated weights. Overall, the evidence on the estimated weights is in line with findings by Clements and Galvão (2009) for US data.

Concerning the relative performance of AR-MIDAS and MF-VAR, we cannot identify a clear winner from the results, as their relative ranking depends on the horizon and the indicator chosen. With respect to production expectations (prod exp), AR-MIDAS seems to dominate MF-VAR for most of the forecast horizons, whereas it is the other way round with respect to consumer confidence (cons conf). For the other cases, however, the results are more horizon-dependent. For short-horizons, MF-VAR forecasts are uninformative in most of the cases, whereas AR-MIDAS yields always relative MSE smaller than one. For long horizons, for example order book (ord book), MF-VAR tends to outperform AR-MIDAS.

Finally, it is interesting to note that there are no substantial differences between the indicators. In particular, we find representatives in all groups of indicators that can provide informative forecasts. In particular, hard indicators like industrial production as well as financial indicators such as interest rates turn out to have information content with respect to future GDP.

Average performance of models over all indicators The selection above concentrates on a selection of models only. To investigate the relative performance of MIDAS and MF-VAR further, we now follow Marcellino et al. (2006) and compare the relative performance of MIDAS and MF-VAR over the full set of indicators. For MIDAS, AR-MIDAS as well as MF-VAR, we compute the pairwise relative MSE of each model to the benchmark and average over all models within a class, see Table 2. On average, MIDAS

Table 2: Average relative MSE performance for forecasting quarterly GDP growth of mixed-frequency model classes against benchmark

	horizon h_m								
model	1	2	3	4	5	6	7	8	9
midas	1.29	1.04	1.10	1.13	1.01	1.03	1.07	0.98	0.98
ar-midas	0.98	0.93	0.93	0.96	0.96	0.99	1.03	0.99	1.00
mf-var	1.27	1.09	1.05	1.08	0.97	0.98	1.01	0.91	0.92

Note: The recursively estimated in-sample mean is used as benchmark forecast. The entries in the tables are obtained as follows: First, pairwise relative MSEs, defined as the MSE of a particular model relative to MSE of the benchmark, are calculated. Second, we take means over all models within a model class (MIDAS, AR-MIDAS or MF-VAR).

cannot do better than the benchmark for almost all the horizons. On the other hand, AR-MIDAS models are on average slightly better than the average up to horizons $h_m = 7$, however, only to a small extent. MF-VAR provides an average relative MSE smaller than one only for a few larger horizons, in particular for horizons $h_m = 8, 9$. This indicates that MF-VAR forecasts have information content for longer horizons than MIDAS, though the gains with respect to the benchmark are small. However, the AR-MIDAS models clearly outperform the MF-VAR approach for short nowcasting horizons. This is due to the more flexible dynamic specification of MIDAS, which can be particularly helpful at short horizons.

Finally, to relate MIDAS and MF-VAR directly, we compute the relative MSE of MIDAS to MSE of MF-VAR. We then average over all these relative MSEs, see Table 3. The ranking is very similar to that emerging from Table 2. For short horizons up

Table 3: Relative performance: (AR-)MIDAS vs. MF-VAR

	horizon h_m								
	1	2	3	4	5	6	7	8	9
midas									
mean	1.05	0.96	1.06	1.05	1.07	1.08	1.06	1.09	1.07
median	0.94	0.98	1.03	1.06	1.04	1.03	1.03	1.03	1.01
ar-midas									
mean	0.77	0.85	0.88	0.90	1.01	1.02	1.02	1.08	1.09
median	0.76	0.83	0.86	0.91	0.97	1.00	0.99	1.03	1.05

Note: The entries in the table are average relative MSEs, where MF-VAR models serve as benchmark for MIDAS and AR-MIDAS. They are computed as follows: First, for each single indicator, the MSE of MIDAS and AR-MIDAS forecasts is respectively divided by the corresponding MSE of the corresponding MF-VAR model. Second, means and medians over all relative MSE are computed.

to $h_m = 4$, AR-MIDAS has an average relative MSE smaller than one, and thus tends to outperform MF-VAR. MIDAS without AR component is in most of the cases worse than MF-VAR. For longer horizons MF-VAR tends to outperform both MIDAS and AR-MIDAS.

Forecast combinations The availability of many indicators and the possible presence of model misspecification and parameter instability suggest that combining forecast from alternative models could yield sizeable gains, since these are the conditions when the advantages from forecast pooling are maximized, see e.g. the review by Timmermann (2006). Clements and Galvão (2008) consider combinations of MIDAS models. A more detailed evaluation of pooling in the presence of a large, mixed-frequency dataset is undertaken in Kuzin, Marcellino and Schumacher (2009). Here we focus on a smaller set of variables chosen, following more closely e.g. Clements and Galvão (2008, 2009). We provide results for the mean, the median, and the weighted mean of the models of a particular class, where combination weights are obtained from the inverse MSE of the previous four-quarter performance of a model.

Below, we provide the relative MSE of the combinations to the benchmark (Table 4), as well as the relative MSE of the combination of MIDAS and MIDAS-AR with respect to the combined MF-VARs (Table 5). To investigate the relative performance of the forecast combinations against the individual models, we compute the percentiles of the forecast combinations with respect to all MSEs of individual models within a corresponding class, see Table 6. The figures in Table 6 represent the percentage of single indicator models that outperform the combined forecast.

Table 4: Relative MSE performance for forecasting quarterly GDP growth of model pooling within a given model class against benchmark

	horizon h_m								
	1	2	3	4	5	6	7	8	9
midas									
mean	1.01	0.85	0.91	0.95	0.88	0.92	0.97	0.90	0.93
weighted mean	0.99	0.81	0.84	0.91	0.79	0.85	0.90	0.83	0.88
median	1.04	0.91	0.98	1.00	0.93	0.96	1.01	0.93	0.93
ar-midas									
mean	0.87	0.83	0.80	0.87	0.87	0.87	0.91	0.92	0.93
weighted mean	0.87	0.83	0.81	0.87	0.84	0.86	0.90	0.88	0.91
median	0.90	0.90	0.91	0.93	0.91	0.89	0.94	0.92	0.91
mf-var									
mean	0.96	0.90	0.93	0.98	0.90	0.93	0.97	0.88	0.89
weighted mean	0.92	0.82	0.88	0.95	0.86	0.88	0.90	0.82	0.84
median	0.98	0.99	1.06	1.06	0.99	1.02	1.03	0.92	0.93

Note: The entries are obtained as follows: First, means, weighted averages based on past MSE performance and medians of all forecasts within a given class of models are computed. Second, the MSE of the combination is computed and finally divided by the MSE of the benchmark, the recursively in-sample sample mean.

Table 5: Relative MSE performance: Pooling of (AR-)MIDAS vs. pooling of MF-VAR

	horizon h_m								
	1	2	3	4	5	6	7	8	9
midas									
mean	1.04	0.95	0.97	0.97	0.97	0.98	1.00	1.03	1.04
weighted mean	1.07	0.98	0.96	0.96	0.92	0.97	0.99	1.00	1.04
median	1.06	0.92	0.92	0.94	0.94	0.95	0.98	1.02	1.00
ar-midas									
mean	0.91	0.92	0.86	0.89	0.96	0.93	0.93	1.05	1.04
weighted mean	0.94	1.00	0.92	0.92	0.97	0.98	1.00	1.07	1.08
median	0.92	0.91	0.86	0.88	0.91	0.87	0.91	1.01	0.97

Note: MF-VAR models serve as benchmark for MIDAS and AR-MIDAS. For further comments, see Tables 3 and 4.

Table 6: Quantiles of MSEs of pooled (AR-)MIDAS and MF-VAR forecasts

	horizon h_m								
	1	2	3	4	5	6	7	8	9
midas									
mean	0.12	0.09	0.15	0.08	0.17	0.21	0.18	0.26	0.34
weighted mean	0.11	0.06	0.06	0.00	0.04	0.11	0.08	0.05	0.08
median	0.13	0.16	0.22	0.14	0.27	0.24	0.27	0.33	0.34
ar-midas									
mean	0.13	0.12	0.11	0.10	0.11	0.11	0.09	0.17	0.15
weighted mean	0.12	0.12	0.11	0.10	0.09	0.11	0.08	0.07	0.12
median	0.28	0.33	0.33	0.27	0.17	0.18	0.16	0.18	0.11
mf-var									
mean	0.11	0.16	0.28	0.19	0.18	0.18	0.08	0.07	0.07
weighted mean	0.08	0.07	0.14	0.17	0.15	0.08	0.06	0.06	0.06
median	0.12	0.33	0.49	0.38	0.46	0.43	0.37	0.30	0.67

Note: We implement the pooling exercise as in Table 4 and then compute the quantiles of MSEs of pooled forecasts in the empirical distribution of all MSEs of individual indicators within a given class of models.

According to the results, most of the combinations do well relative to the benchmark. This results is relatively stable over all forecast horizons, with only few exceptions (MIDAS without AR terms and $h_m = 1$, MF-VAR and the median). Comparing Tables 2 and 4, we conclude that forecast combination is a useful method both in case of MIDAS and MF-VAR models, since the performance of forecast combinations relative to our benchmark is in most of the cases better then the average of all relative MSEs within a given class over all indicators. Following the direct comparison of the combinations in Table 5, AR-MIDAS seems to outperform MF-VAR at short forecast horizons, but its advantage seems not so pronounced as in Table 2, where only individual models were compared. Pooling of MF-VARs performs better at long forecast horizons ($h_m = 8, 9$).

The percentiles of the forecast combinations in Table 6 indicate that pooling is a useful alternative to forecasting with individual models. When the weighted mean is applied, several figures in Table 6 are clearly below 10% for MIDAS and MF-VAR models. However, even the forecast combinations cannot outperform all of the individual models. For example, in the case of pooling with weighted means for AR-MIDAS at $h_m = 1$, there are 12% individual models within the AR-MIDAS class with smaller MSE than the combination. But it should be considered that with a large set of indicators, it is natural to find that some of them perform particularly well. In addition, the analysis of Banerjee and Marcellino (2005) clearly indicates that the best leading indicators for euro area GDP growth change over time, and the pooled forecast can protect from this instability. With respect to the other weighting schemes, the results are worse than the weighted mean. For example, the median seems to work worse, in particular, when applied to MF-VAR models.

4 Conclusions

This paper considers MIDAS and MF-VAR as alternative forecasting methods suitable for now- and forecasting with mixed-frequency data that is also subject to different publication lags.

Theoretical arguments indicate that we cannot expect one approach to be clearly superior than the other. For example, MIDAS is a direct multi-step forecast approach, whereas MF-VAR provides iterative forecasts. MIDAS is more parsimonious than MF-VAR, but depends on certain distributed lag assumptions that might be too rigid. Thus, the relative performance of the two approaches will depend on the underlying unknown data generating process, and either MIDAS or the MF-VAR could dominate in a specific empirical application. Hence, we compare the alternative forecasting approaches empirically. In particular, we carry out a recursive comparison exercise in terms of now- and forecasting quarterly Euro Area GDP with a set of about twenty monthly indicators.

The main results are the following.

1. If we look at selected indicators, we find representatives of both MIDAS and MF-VAR classes of models that work well compared to the benchmark. However, the relative performance of MIDAS and MF-VAR differs with respect to the predictors and forecast horizons, and there seems to be no clear winner in terms of forecasting performance.

2. If we compare all the models pairwise with the same indicator and compute the average MSE over the whole set of models, we find that MF-VAR outperforms MIDAS and AR-MIDAS at long forecast horizons, whereas AR-MIDAS can do better at short horizons up to three months.

3. When the single MIDAS and MF-VAR forecasts are combined, there are advantages compared to the single-indicator models. In addition, pooled MF-VAR forecasts are better at longer horizons, and pooled MIDAS forecasts at shorter horizons.

Overall, the MF-VAR seems to be a reasonable competitor to MIDAS in macroeconomic datasets such as the one chosen here. More generally, it can be useful to consider both classes of models for forecasting specific variables of interest, and pooling can provide additional advantages.

A possible area of future research would be the use of other MIDAS weighting schemes, i.e. the exponential Almon lag with richer structure or the Beta polynomial suggested by Ghysels et al. (2007), especially for long horizon predictions. Further improvements in predictability could be also possible by exploiting daily data in the MIDAS approach, see Andreou, Ghysels, and Kourtellos (2009b) as an example for US data.

References

[1] Andreou, E., Ghysels, E., Kourtellos, A. (2009a), Regression Models with Mixed Sampling Frequencies, Journal of Econometrics, forthcoming.

[2] Andreou, E., Ghysels, E., Kourtellos, A. (2009b), Should macroeconomic forecasters look at daily financial data?, mimeo.

[3] Banerjee, A., Marcellino, M., Masten, I. (2005), Leading indicators for Euro area inflation and GDP growth, Oxford Bulletin of Economics and Statistics, 67, 785-813.

[4] Bernanke, B. S., Boivin, J. (2003), Monetary Policy in a Data-Rich Environment, Journal of Monetary Economics 50, 525-546.

[5] Bhansali, R. J. (2002), Multi-step forecasting, in: Clements, M. P., and Hendry, D. F. (eds.), A Companion to Economic Forecasting, 206–221.

[6] Chevillon, G., Hendry, D. F. (2005), Non-parametric direct multi-step estimation for forecasting economic processes, International Journal of Forecasting 21, 201-218.

[7] Clements, M. P., Galvão, A. (2008), Macroeconomic Forecasting With Mixed-Frequency Data: Forecasting Output Growth in the United States, Journal of Business & Economic Statistics 26, 546-554.

[8] Clements, M. P., Galvão, A. (2009), Forecasting US output growth using Leading Indicators: An appraisal using MIDAS models, Journal of Applied Econometrics, forthcoming.

[9] Ghysels, E., Rubia, A., Valkanov, R. (2009), Multi-period forecasts of volatility: direct, iterated and mixed-data approaches. Available at: http://www.unc.edu/~eghysels/papers/Var_9.pdf.

[10] Ghysels, E., Santa-Clara, P., Valkanov, R. (2004), The MIDAS touch: Mixed Data Sampling regression models, mimeo.

[11] Ghysels, E., Santa-Clara, P., Valkanov, R. (2005), There is a risk-return after all, Journal of Financial Economics 76, 509-548.

[12] Ghysels, E., Santa-Clara, P., Valkanov, R. (2006), Predicting volatility: Getting the most out of return data sampled at different frequencies, Journal of Econometrics 131, 59–95.

[13] Ghysels, E., Sinko, A., Valkanov, R. (2007), MIDAS Regressions: Further Results and New Directions, Econometric Reviews 26, 53-90.

[14] Ghysels, E., Valkanov, R. (2006), Linear Time Series Processes with Mixed Data Sampling and MIDAS Regression Models, mimeo.

[15] Ghysels, E., Wright, J. (2008), Forecasting Professional Forecasters, Journal of Business and Economic Statistics, forthcoming.

[16] Giannone, D., Reichlin, L., Small, D. H. (2008), Nowcasting GDP and Inflation: The Real-Time Informational Content of Macroeconomic Data Releases, Journal of Monetary Economics 55, 665-676.

[17] Kuzin, V., Marcellino, M., Schumacher, C. (2009), Pooling versus model selection for nowcasting with many predictors: An application to German GDP, Deutsche Bundesbank Discussion Paper, Series 1: Economic Studies, 03/2009.

[18] Marcellino, M., Musso, A. (2008), Real time estimates of the euro area output gap: Reliability and forecasting performance, mimeo.

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

[20] Marcellino, M., Stock, J. H., Watson, M. W. (2006), A comparison of direct and iterated multistep AR methods for forecasting macroeconomic time series, Journal of Econometrics 135, 499-526.

[21] Mariano, R., Murasawa, Y. (2003), A New Coincident Index of Business Cycles Based on Monthly and Quarterly Series, Journal of Applied Econometrics 18, 427-443.

[22] Mariano, R., Murasawa, Y. (2007), Constructing a Coincident Index of Business Cycles Without Assuming a One-Factor Model, Discussion Paper 2004-6, College of Economics, Osaka Prefecture University.

[23] Mittnik, S., Zadrozny, P. A. (2005), Forecasting German GDP at Monthly Frequency Using Monthly IFO Business Conditions Data, in: Sturm, J.-E., Wollmershäuser, T. (eds.), Ifo Survey Data in Business Cycle and Monetary Policy Analysis, Springer-Verlag, 19-48.

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

[25] Timmermann, A. (2006), Forecast Combinations, in: Elliot, G., Granger, C. W. J., Timmermann, A. (eds.), Handbook of Economic Forecasting, Vol 1, 135-196.

[26] Wohlrabe, K. (2009), Forecasting with Mixed-frequency Time Series Models, Ph. D. dissertation, University Munich.

[27] Zadrozny, P. A. (1988), Gaussian-Likelihood of countinuous-time ARMAX models when data are stocks and flows at different frequencies, Econometric Theory 4, 108-124.

A Euro Area dataset

This appendix describes the time series for the Euro Area economy used in the forecasting exercise. The whole data set for Euro Area contains 23 monthly time series over the sample period from $1992M1$ until $2008M6$. The time series cover broadly the following groups of data: industry statistics, surveys, financial data (interest rates, exchange rates, money stocks), and miscellaneous indicators, such as raw material price indices and car registrations. A complete list of variables is provided below, together with abbreviations used in the description of results in the main text.

The sources of the time series are the databases of the Bundesbank and the ECB. Original sources are the European Commission, the ECB, and the HWWA. Natural logarithms were taken for all time series except interest rates and the surveys. Stationarity was obtained by appropriately differencing the time series. All of the time series taken from the above sources are already seasonally adjusted, where this was necessary.

A.1 Industrial production

prind - production: total
prcap - production: capital goods industry
print - production: intermediate goods industry
prcons - production: consumer goods industry
prcs - production: construction sector

A.2 Surveys

indconf - business confidence industry
prodexp - business production expectations
ordbook - business order books
assstock - assessment of stocks of finished goods
consconf - consumer confidence

A.3 Interest rates, exchange rates, money stocks

is3m - money market rate, 3-months EURIBOR
il10 - yields on 10-year government bonds (GDP weights)
zdiff103 - yield spread: bond yields with 10 years minus 3 months EURIBOR
m1 - monetary aggregate M1
m3 - monetary aggregate M3
loans - loans
een - nominal effective exchange rate of the euro against the currencies of the EER-22 group
eer - real effective exchange rate of the euro against the currencies of the EER-22 group (on basis of consumer price index)

A.4 Raw material prices, car registrations

hwwa - HWWA raw material price index

hwwaind - HWWA raw material price index: industrial raw materials

hwwaenerg - HWWA raw material price index: energy industrial raw materials

carcomm - car registrations: new commercial

carpass - car registrations: new passenger cars