Forecasting with Mixed-frequency Time Series Models

Inaugural-Dissertation
zur Erlangung des Grades
Doctor oeconomiae publicae (Dr. oec. publ.)
an der Ludwig-Maximilians-Universität München

2008

vorgelegt von

Klaus Wohlrabe

Referent: Prof. Stefan Mittnik, Ph.D.
Korreferent: Prof. Dr. Kai Carstensen
Promotionsabschlussberatung: 04. Februar 2009

Danksagung

Während der Anfertigung dieser Dissertation habe ich direkt und indirekt viel Unterstützung erfahren und möchte mich dafür bedanken. Ohne diese Unterstützung wäre vieles vielleicht nicht, oder nicht so gut umzusetzen gewesen. Viele Menschen haben mir geholfen, diese Dissertation zu einem sehr guten Abschluss zu führen.

Zunächst möchte ich meinem Doktorvater Stefan Mittnik für seine Bereitschaft zur Betreuung, sein Interesse an dieser Doktorarbeit und seine Unterstützung danken. Des Weiteren danke ich Kai Carstensen vom ifo Institut für Unterstützung und Diskussionen zu verschiedenen Aspekten dieser Arbeit. Peter Zadrazony beantwortete mir alle Fragen bzgl. seines Modells. Seine Erörterungen haben die Arbeit sehr bereichert. Von ihm habe ich auch den Enthusiasmus für die untersuchte Fragestellung übernommen. Besonderen Dank an ihn. Nikolay Robinzonov unterstützte mich in der Verfeinerung und Neuentwicklung von MATLAB-Routinen, welche in dieser Arbeit genutzt wurden. Er war auch immer ein kritischer Begleiter.

Dirk Ulbricht verdanke ich eine intensive Diskussion der ganzen Arbeit. Er half mir viele Dinge aus einem anderen Blickwinkel zu sehen. Mit Johannes Mayr konnte ich einige Aspekte bzgl. der Prognose von Zeitreihen intensiv diskutieren.

Unzählige Kommentare habe ich auch von Gabriela Schütz, Marc Pioponik, Martin Schlotter, Anna Stangl, Stefan Lachenmaier, Christian Schumacher und den Teilnehmern an vielen internationalen Konferenzen erhalten. Die Deutsche Forschungsgemeinschaft unterstützte diese Doktorarbeit im Rahmen eines Projektes auch finanziell. Das ifo Institut gewährte mir ein intensives Forschungsumfeld. Beiden Institutionen möchte ich danken. Die Kollegen am ifo Institut und insbesondere der Abteilung Humankapital waren immer wichtige Ansprechpartner und boten Unterstützung in schwierigen Phasen.

Abschließend möchte ich meinen Eltern für ihre jahrelange Unterstützung und Begleitung meines Lebensweges danken. Zum guten Schluss möchte ich Grit danken. Sie hat mich mit unschätzbar viel Verständnis und Interesse für diese Arbeit begleitet. Für ihre Unterstützung, auch in sehr schwierigen Phasen, bin ich im tiefen Dank verbunden.

CONTENTS

LIST OF TABLES

LIST OF FIGURES

NOMENCLATURE

AR	autoregressive
ARCH	autoregressive conditional heteroscedasticity
GARCH	generalized autoregressive conditional heteroscedasticity
EM	Expectation maximization
GDP	Gross Domestic Product
GNP	Gross National Product
HF-VAR	high-frequency VAR
IP	Industrial Production
KF	Kalman Filter
LF-VAR	low-frequency VAR
m	frequency mixture
MF-VAR	mixed-frequency VAR
MIDAS	MIxed DAta Sampling
MSE	mean-squared error
NBER	National Bureau of Economic Research
NLS	Non-linear least squares
OLS	Ordinary Least Squares
RMSE	root-mean-squared error
US	United States
VAR	vector autoregressive
VARMA	vector autoregressive moving average
WF	weighting function

1. INTRODUCTION

Movements in economic and financial time series are closely watched by governments, central banks and companies. Among these series are stocks, market indexes, exchange rates, interest rates, survey data, and other aggregates of macroeconomic activity. Each time series possibly contains information about the current and the future unobserved states of the economy. Thus, they influence market expectations of economic agents. The knowledge of the current and future states affects interest rate decisions of central banks, asset price movements, investment decisions, private savings, and so forth. Moreover, as more and more time series are easily available electronically, the expectations and estimations of the current state of the economy are constantly updated.

One of the most important and comprehensive indicators, which is acknowledged to represent the unobserved state of the economy, is the Gross Domestic Product (GDP) which is recorded at quarterly intervals. Therefore it is of great interest to economic agents to estimate and forecast GDP. However, many other economic time series are sampled at a higher frequency. Survey data from the EU Commission, OECD or the Ifo Institute are sampled monthly. The same is true for other macroeconomic variables like inflation, unemployment and industrial production. Moreover, numerous economic variables , like stocks, the oil price, interest rates, and exchange rate are also available on a daily basis. Some evolve almost in real-time (on a minute basis). But such data are inherently 'noisy'. Furthermore, a chief difficulty with using multiple possible indicators is that they can, and usually do, provide conflicting signals; and there is no agreed-upon way for aggregating the statistics to give a single-valued answer.

The estimation of the current state and forecasting of the future state of macroeconomic variables, such as GDP, with multivariate time series models faces the general problem that the observations tend to be recorded and published at different frequencies. Therefore, virtually any attempt to estimate a multivariate economic time series model confronts the problem of efficiently using mixed-frequency data.

The literature and methodology so far have assumed that *all* processes are sampled at the same frequency. To ensure the same frequency, either the higher-frequency data are aggregated to the lowest frequency, or the lower-frequency data are interpolated to the highest frequency. In most empirical applications, the higher frequency is aggregated to the lower frequency by averaging, summing up, or by taking a representative corresponding value (for example, the third month of the quarter). Neither of these options is generally satisfactory. First, temporal aggregation destroys sample information. Aggregated processes entail less information, and such an information loss typically results in poorer predictability. Much is known about the adverse effect of temporal aggregation on prediction, see for example Lütkepohl (1987). Second, commonly used interpolation methods generally do not fully exploit the available sample information.

In addition to the frequency sampling problem, many economic time series are published with a delay and are subject to revision. For instance, the first estimate of GDP in Germany is released six weeks after the end of the quarter. The lack of a timely, comprehensive economic picture may mean that policy needs may be recognized only many months after a significant slowdown or an acceleration in the economy. This problem is especially important around business cycle peaks or troughs, where there may be only weak evidence that the economy is changing direction. The lack of timely information concerning macroeconomic aggregates is also important for understanding private sector behaviour, and in particular the behaviour of asset prices. When agents make trading decisions based on their own estimates of current macroeconomic conditions, they transmit information to their trading partners. This trading activity leads to the aggregation of dispersed information, and in the process affects the behaviour of asset prices. But many macroeconomic time series provide timely information about the current state of the economy. Some of them (like surveys) are not subject to revision, but are sampled at a higher frequency (for example monthly) compared to a target variable like GDP. So there is a potential need for models that combine data from different frequencies. This area of research evolved recently and is virtually unexplored.

Early examples of attempts to combine data from different frequencies were linkage models and bridge equations. In the former approach, forecasts are generated at different frequencies. The forecasts are combined to improve the forecasting accuracy of the lower-frequency time series. In contrast, bridge equations are essentially single-frequency time series models. The high-frequency data are forecasted up to the desired forecast horizon in a separate time series model. Finally, these forecasts are aggregated to the lower frequency and plugged into a lower-frequency time series model (as contemporaneous values). Bridge equations are especially useful for nowcasting: the forecast of the current period. These early attempts proved to be successful as they increase forecast accuracy in the short-run. But what about time series models, which are able to handle mixed-frequency time series models without any data transformations and forecast combinations? Is there something to gain in forecast accuracy?

From a theoretical point of view, Ghysels and Valkanov (2006) and Hyung and Granger (2008) showed that there are gains in terms of forecasting accuracy from considering mixtures of different frequencies without transformation of the data. Ghysels and Valkanov (2006) derive circumstances under which mixed-data sampling achieves the same forecasting efficiency as the hypothetical situation where all series are available at the highest frequency. However, the conditions behind this result cannot be verified empirically. In a Monte Carlo study, Ghysels and Valkanov (2006) demonstrate that the in-sample forecasting mean squared errors of mixed-frequency models are lower than temporally aggregated single-frequency time series models.

Currently there are two competing approaches in the literature to deal with mixed-frequency time series. The first one was proposed by Zadrozny (1988) for directly estimating a multivariate, continuous, autoregressive, moving average (VARMA) model with mixed-frequency time series. Zadrozny (1990, 2008) extended this idea to discrete time. The approach of the method is to assume that the model operates at the highest frequency in the data. All variables are assumed to be generated, but not necessarily observed, at this highest frequency, and thus can be used to produce forecasts of any variable at this frequency. Variables which are observed at a lower frequency are viewed as being periodically missing. For example, with quarterly and monthly data, the model is assumed to generate all variables at monthly intervals and each quarterly observation is assigned to the last month of a quarter, so that observations for the remaining months in the quarter are viewed as missing.

The second approach to handle time series sampled at different frequencies, which they term MIDAS (MIxed DAta Sampling), was proposed by Ghysels, Santa-Clara, and Valkanov (2004). MIDAS models specify conditional expectations as a distributed lag of regressors at some higher sampling frequencies. In practice the lowest frequency is regressed on the higher frequency. To avoid parameter proliferation, a weighting function is employed.

Practical applications of these two approaches are rather rare. Zadrozny (2008) shows that the forecasting performance for US Gross National Product (GNP) can be improved with the state-space VARMA model over an autoregressive benchmark model. Mittnik and Zadrozny (2005) find similar promising results for German GDP forecasts using the Business Climate Index of the Ifo Institute in Munich. Ghysels, Santa-Clara, and Valkanov (2006) predict volatility (5-minute frequency data) with various specifications of regressors. They show that MIDAS models outperform single-frequency benchmark models. Ghysels and Wright (2008) come to a similar conclusion when using daily financial data to make monthly and quarterly macroeconomic forecasts. Clements and Galvao (2008) obtain better forecasts for US output and inflation using the MIDAS approach compared to benchmark models (for example bridge models). Marcellino and Schumacher (2007) demonstrate how factor models, in combination with the MIDAS approach, can be used to improve short-run forecasts of German GDP.

Given these promising results it is interesting to ask whether these results hold in general, that is for any frequency mixture and time series length. More importantly, these two approaches have not been directly compared to date. Until now, there seems to be a 'peaceful coexistence' between both approaches. Articles dealing with one approach do not cite the other.

However, these two approaches are possible candidates to account for the problems stated before. The data need not be transformed and are able to account for any kind of high-frequency data. Are these models able to improve the forecasting accuracy of lower-frequency variables, like GDP, by using high-frequency data? We want to answer these questions in this thesis.

This dissertation contributes to the literature in several ways. First we outline all theoretical aspects concerning mixed-frequency data modelling. We are the first to present the different approaches in one review. So far the different approaches are presented quite disconnected in the literature. In a second step we review the literature that has dealt with forecasting with mixed-frequency data. This review presents empirical strategies and the success of mixed-frequency modelling approaches. Such a review is the first to attempt this.

In the main part of the thesis we compare the forecasting success of the two new mixed-frequency time series models: the mixed-frequency VAR and the MIDAS approach. Before the calculation of forecasts, a time series model needs to be specified. Current articles on mixed-frequency time series models neglect the model specification aspect. We are the first to investigate some specification issues for both model classes relevant for forecasting purposes.

As the mixed-frequency VAR operates at the highest frequency of the data, it is important to know how many lags should be included. Is there a problem of parameter proliferation or are models with few lags sufficient to obtain an improvement in forecasting accuracy? The lag selection problem also intrudes into the MIDAS model specification, but in a different way. Inclusion of further lags does not increase the number of estimated parameters. Due to large (high-frequency) data sets in financial econometrics, it is possible to include many lags. In contrast, in macroeconomics the trade-off of including more lags and shortening the estimating sample is more severe. Closely related to the lag selection problem is the question of whether the weighting function should be restricted (for example, ensuring declining weights) or not. Restrictions may be useful to ensure that more recent observations are given bigger weights than others are. In the literature, we find examples with restrictions and some without restrictions. But there is no theoretical or economic reasoning behind these choices.

Having demonstrated how to specify the model we want to systematically compare the forecasting performance of the two approaches in an extensive Monte Carlo study, since the two approaches have not been compared before. We will consider four data-generating processes to cover reasonable data structures. We allow both for homoscedastic and heteroscedastic errors in the data-generating process. The latter one especially is motivated by the fact that many economic time series show the existence of volatility clustering ((G)ARCH - generalized autoregressive conditional heteroscedasticity effects). Financial time series exhibit inherently volatility clustering. We investigate whether heteroscedastic data do influence the forecasting performance significantly compared with homoscedastic errors. We will focus on three mixtures that are predominant in macroeconomic forecasting: monthly-quarterly, weekly-quarterly, and quarterly-yearly. In addition to comparing the mixed-frequency approaches to each other, we investigate whether they have an advantage over single-frequency models. The final question in the Monte Carlo study is that we want know how the forecasting performance changes when larger time series are under investigation.

Eventually, we compare the two approaches in a case study using real data. We forecast German GDP growth with different indicators and different models. We focus especially on the nowcasting aspect. Does intra-quarterly information help to improve forecasting accuracy in the short and long run? Factor analysis, as a method to condense large information sets, will play a prominent role in our analysis. The second part of the empirical application is a nowcasting experiment. We investigate whether it is possible to track the German economy on a daily basis. In contrast to approaches that estimate a latent variable as the current state of the economy we are the first to track the economy almost in real-time by forecasting GDP growth on a daily basis. With this experiment we demonstrate how useful daily data are for forecasting GDP. We consider this as a starting point for future research in this area.

The thesis is structured as follows. Chapter 2 outlines the theoretical aspects of mixed-frequency data modelling, both data transformation and mixed-frequency time series models. Chapter 3 contains the literature review. In chapter 4 we introduce all aspects of the Monte Carlo study. In chapter 5 we demonstrate how to specify a mixed-frequency model for forecasting purposes. Chapter 6 contains the Monte Carlo forecasting study. In chapter 7 we use both approaches to forecast German GDP. We show how these models can help to increase forecasting accuracy with monthly and daily leading indicators. The latter is viewed as a nowcasting experiment. Finally we summarize and conclude.

2. MIXED-FREQUENCY DATA: THEORETICAL MODEL AND TRANSFORMATION ISSUES

Generally, most time series models employed for forecasting assume the same frequency for all data used. When time series are sampled at different frequencies one needs to transform them. Either the lower frequencies are interpolated to the highest frequency or the higher frequencies are aggregated to the least frequency. Strictly speaking, forecasting using mixed-frequency data is a two-step procedure. First, the data are transformed into a single frequency and second, the transformed data are plugged into a structural macroeconomic or time series model for forecasting. In this chapter we outline different transformation methods to obtain single-frequency data. Then we present the early approaches to combine data sampled at different frequencies, linkage models, and bridge equations. We keep the summary on data transformation short, as we focus on time series models able to handle mixed-frequency data and where no data transformation is necessary: the mixed-frequency VARMA approach and the MIxed DAta Sampling (MIDAS) approach. We provide details of the approaches, and specification and estimation issues. We also briefly compare these models in terms of practical forecasting without empirical investigation. The chapter starts with some details on notation and interpretation of mixed-frequency data.

2.1 Preliminaries

As we deal with mixed-frequency data throughout the thesis, first we want to outline the notation in order to avoid confusion. The basic time unit is denoted with t. With t we label the time unit of the lower frequency time series which has the range $t = 1, \dots, T$. To relate the difference frequencies to each other we will use the term frequency-mixture denoted as m. Between t and $t-1$ the higher-frequency variable is observed m times. To avoid cumbersome notation we will also use a time unit for the higher-frequency data, denoted as τ. The length of the higher frequency is then $\tau = 1, \dots mT$. The basic frequency mixtures investigated in more detail in this thesis are:

- $m = 3 \Rightarrow$ monthly-quarterly data, where 3 months constitute a quarter
- $m = 4 \Rightarrow$ weekly-monthly or quarterly-annual data, where 4 weeks (quarters) define a month (year)
- $m = 12 \Rightarrow$ weekly-quarterly data, where 12 weeks constitute a quarter.

This choice is based on typical mixtures that confront a researcher in empirical macroeconomics. Generally, m can take any value. Further possible mixtures are

- $m = 2 \Rightarrow$ biannual-annual data (not considered so far in the literature)
- $m = 20 \Rightarrow$ week-daily-monthly data (20 trading days per month)
- $m = 60 \Rightarrow$ week-daily-quarterly data (60 trading days per quarter)
- $m = 250 \Rightarrow$ week-daily-yearly data (250 trading days per year)

which are mixtures that can be found in financial economics. In the literature review we will present the empirical choices of m. This list can be extended by the practitioner to any (constant) mixture. The presence of two different frequencies in one time series model is represented by the notation $x_{t-i/m}$. Suppose, x_t denotes the March value of some monthly time series. Suppose that the target variable is sampled at quarterly intervals than m equals three. The February monthly value is then donated by $x_{t-1/3}$ (one month of observations is missing in the current quarter), the January value by $x_{t-2/3}$, the December by $x_{t-3/3} = x_{t-1}$ and so forth. A simple regression model is than given by

$$y_t = \alpha + \sum_{i=1}^{p} \beta_i x_{t-i/m} + \epsilon_t \tag{2.1}$$

This definition and notation of mixed-frequency data can be problematic (especially for higher frequency mixtures) in empirical applications for several reasons. First, it assumes that the data are equidistant over time. But this may not be the case, especially for very high-frequency data. For example, the three months constituting a quarter may not be of the same length. Abstracting from a leap year, in the first quarter, the month of January has 31 days, February 28 days and March 31 days. Furthermore the number of Mondays, Tuesdays etc. are not the same for each month. And there are moving holidays, like Easter. Such differences can have an impact on some macroeconomic variables such as retail sales or industrial production. In practical applications one can account for this problem, by seasonally and workday adjusting the data. But this is not done for every variable.[1] In general this aspect is ignored in practical macroeconomic forecasting.

Even more severe is the problem when mixing daily and monthly data. In financial applications a week is defined by five trading days. As months usually consist of 30 or 31 days it is likely that we have more than 20 trading days per month. On the other hand it is possible not to have enough trading days to define a week or a month. This occurs in months with many holidays as in December. On the one hand, one has too much information and has to discard some of it. On the other hand, there is too little information and one needs to extrapolate or interpolate it. These problematic aspects have not been investigated in the literature so far. We will abstract from the information content problem in our Monte Carlo study by assuming equidistant

[1] Official seasonally adjusted time series account for the actual number of working days. Many econometric packages provide seasonal adjustment, but to account for the actual number of working days, the researcher needs to provide this information, but this is rarely done. Thus seasonal adjustment is a technical process unless we account for the distribution of frequencies used.

observations. In the empirical example we will outline some strategies to deal with this problem.

There are other possible notations for denoting mixed-frequency data in a time series model, for instance a superscript for the corresponding frequency. For instance, y_t^m denotes month t,[2] whereas y_t^q denotes quarter t. Equation (2.1) can then be rewritten as

$$y_t^q = \alpha + \sum_{i=1}^{p} \beta_i x_{t-i}^m + \epsilon_t. \tag{2.2}$$

2.2 Transforming Mixed-frequency Data

2.2.1 Interpolation

Interpolation is rarely used in applied econometric forecasting.[3] Interpolation assumes that the lower-frequency variable is interpreted on some higher frequency which exhibits missing observations. Lanning (1986) points out that economists facing missing data have basically two different ways to interpolate. One approach is to estimate the missing data simultaneously with the model parameters. A second way is a two-step procedure where in a first step the missing data, which can be independent of the economist's model, are interpolated. In a second step, the new augmented series is used to estimate the model. Based on simulations, Lanning (1986) suggests using the two-step approach, as the model parameters have larger variances in the simultaneous approach.

The simplest way to interpolate is to apply pure statistical interpolation methods as linear, quadratic, or cubic interpolation. But these methods do not account for possible intra-period variability of the higher frequency.[4]

[2] The superscript m denotes in this case the month and not the frequency mixture.

[3] Examples can be found in the next chapter. Interpolation is also sometimes termed 'disaggregation'.

[4] A recent chronology of interpolation (Meijering (2002)) contains 358 (mostly modern) citations.

A simple approach to recovering disaggregated values is to compute partial weighted averages of the aggregated series, see for example Lisman and Sandee (1964). A different approach is that the disaggregated values are those which minimize a specific loss function under a compatibility constraint with aggregated data, see for example Boot, Feibes, and Lisman (1967), Cohen, Müller, and Padberg (1971), and Stram and Wei (1986). A further constraint can be added. This involves the existence of a preliminary disaggregated series (related time series), so that the interpolation issue becomes how best to revise the data for them to be compatible with the aggregated data, see for example Friedman (1962), Denton (1971), Chow and Lin (1971) (later extend by Chow and Lin (1976)), Fernandez (1981), Litterman (1983), and Mitchell, Smith, Weale, Wright, and Salazar (2005).

Assuming the higher-frequency observations as missing, there is a huge literature on estimating such missing observations relying on state-space interpretations and Kalman filtering, see for example Harvey and Pierse (1984), Kohn and Ansley (1986), Nijman and Palm (1990), and Gomez and Maravall (1994). The Kalman filter uses the underlying serial dependence of the data in order to estimate conditional expectations of the missing observations. More recent state-space approaches are provided by Bernanke, Gertler, and Watson (1997), Liu and Hall (2001), Proietti (2006), Nieto (2007), and Angelini, Henry, and Marcellino (2006). In the latter article the usage of factor models is proposed for interpolation. Angelini, Henry, and Marcellino (2006) also conduct a Monte Carlo study to investigate the performance of different interpolation approaches. Given a large information set of higher-frequency variables, the factor approach performed best. Given only a few related time series, the approach by Chow and Lin (1971) cannot be outperformed. These results still hold when applied to a real data set (GDP and inflation) where some of the observations were dropped.

Cuche and Hess (2000) provide an overview of which interpolation approach should be used depending on data availability and assumptions on the data-generating process of the interpolated series.

2.2.2 Aggregation

In general, aggregation generates an information loss. The analysis of temporal aggregation starts with the seminal paper of Amemiya and Wu (1972). They demonstrate that if the original variable is generated by an AR model of order p, the aggregate variable follows an AR model of order p with MA residuals structure. Tiao (1972) and Amemiya and Wu (1972) study the issue of information loss due to aggregation. In a general multivariate framework, Lütkepohl (1987) contains a deep analysis of temporal (and contemporaneous) aggregation for VARMA models; it also examines the impact of temporal aggregation of the efficiency of forecasts. For a recent survey on temporal aggregation of single-frequency variables see Silvestrini and Veredas (2008).

A single time series

Temporal aggregation of the higher frequency variables to the lowest frequency is by far more common in applied work. 'Aggregation' can be interpreted in different ways depending on the definition of the variable in focus. Assuming a stock variable, the latest available value of the higher frequency can be used. Assuming monthly-quarterly data, one could use the first, second or the last ('stock-end') value to be representative of the whole quarter. Considering the last value of the quarter one could argue that information from the previous quarters is reflected in this value. Employing such an 'aggregation' scheme for flow variables is more a matter of convenience than being theoretically justified. It is more appropriate for 'small' frequency mixtures such as monthly-quarterly than for larger ones such as daily-quarterly, as not too much information is condensed in one value.

As already stated, aggregation in practice depends on the interpretation of the data. The standard aggregation method is averaging over one lower-frequency period

$$x_t = \frac{1}{m} \sum_{i=1}^{m} x_{t-i/m}. \tag{2.3}$$

This is even done for stock variables. The average defines the average information content over the low-frequency period. For flow variables the higher-frequency values are simply added

$$x_t = \sum_{i=1}^{m} x_{t-i/m}. \tag{2.4}$$

Many time series

During the 1980s and 1990s more and more economic variables have become available. The inclusion of too many variables increases estimation uncertainty, which can lead to a deterioration in forecasting accuracy. Thus, the principle of parsimony in econometrics prevents the inclusion of all possible indicators in one time series model. One way to exploit the whole available information set is to condense all time series into a few variables. One attempt is the construction of composite indicators (CI). These CIs can be constructed for several purposes, for example forecasting or, coincidentally, for describing the current state of the economy. One popular attempt is the Stock-Watson Experimental Coincident Index (XCI) developed by Stock and Watson (1989). Stock and Watson (1991) construct a coincident index by applying maximum likelihood factor analysis to four monthly coincident indicators. So far, these methods assume the same data frequency for all time series.

Mariano and Murasawa (2003) extend the model of Stock and Watson (1989) to allow for mixed-frequency data, especially the inclusion of quarterly data. The model is cast into state-space form and the likelihood is maximized with the Kalman Filter.[5] The suggested filtering algorithm is only an approximation. Proietti and Moauro (2006) avoid the approximation at the cost of moving to a non-linear model with a corresponding, rather tedious, non-linear filtering algorithm. Mariano and Murasawa (2003) extract a new coincident indicator using one quarterly and four monthly time series.

Another approach to condensing information is estimating factors from large data sets. Factor analysis has become popular in applied forecasting (see

[5] The idea of using the Kalman filter to account for mixed-frequency data in this context was also suggested by Nunes (2005) and Crone and Clayton-Matthews (2005).

the literature review in chapter 3). We want to outline the aspect of factor estimation with irregular and mixed-frequency data in more detail, as we will apply these approaches in our empirical forecasting exercise. For illustration purposes we start with the static factor model with single-frequency data.[6]

Let Y_τ be an $(N \times 1)$ dimensional vector of stationary time series with observations for $\tau = 1, \ldots, mT$, and we assume that the series have zero means. The variables in a factor model are represented as the sum of two mutually orthogonal components: the common and the idiosyncratic components. The common component is driven by a small number of factors common to all variables in the model. The idiosyncratic component is driven by variable-specific shocks. The factor model can be written as

$$Y_\tau = \Lambda F_\tau + \xi_\tau \tag{2.5}$$

where F_τ is a $(r \times 1)$ vector of factors and the $(N \times r)$ dimensional matrix Λ contains the factor loadings. The idiosyncratic components comprise the vector ξ_τ. The basic idea of factor models is that a small number of factors can explain most of the variance of the data. The factors can be estimated with the principal components approach. Let V be the $(N \times r)$ matrix of stacked eigenvectors $V = (V_1, \ldots, V_r)$ corresponding to the r largest eigenvalues of the $(N \times N)$ sample covariance matrix $\hat{\Sigma} = (mT)^{-1} \sum Y_\tau Y_\tau'$. The principal components estimator of the factors and the loading matrix is given by

$$\hat{F}_\tau = V' Y_\tau \tag{2.6}$$

$$\hat{\Lambda} = V \tag{2.7}$$

The asymptotic properties of the factor estimators are outlined in Breitung and Eickmeier (2006) building on Stock and Watson (2002a) and Bai (2003). Under mild assumptions on serial correlation, heteroscedasticity, and cross-correlation among idiosyncratic components, the asymptotic normal distribution of the factor estimates is established for $N, mT \to \infty$.

[6] The following exposition draws on Schumacher and Breitung (2008).

Stock and Watson (2002b) showed how the expectation maximization (EM) algorithm can be employed in factor analysis when data irregularities such as missing observations or mixed-frequencies are present. Before we present the idea of the EM algorithm we need to introduce a transformation matrix. Let Y_i^- a $T^- \times 1$ vector of observations for the variable i, that may contain missing values and the complete vector of realizations Y_i, where $T^- \leq mT$. It is assumed that the relationship between observable and complete data are given by the linear relationship

$$Y_i^- = S_i Y_i \tag{2.8}$$

where S_i is a known $(T^- \times mT)$ selection matrix that can tackle missing values or mixed frequencies.[7] For example, if all observations are available, the matrix S_i is an identity matrix. Mixed-frequency data are interpreted as missing values. Consider mixing monthly and quarterly data. For a stock variable the first two months are not available, whereas the last monthly observation in a quarter is equal to the quarterly published value. The selection matrix S_i is adjusted by elimination of the respective rows.

Given these preliminaries, the basic idea of the EM algorithm with mixed-frequency data proceeds as follows:

1. Create a data set sampled at the highest frequency in the data. Thus we produce an initial (naive) higher-frequency estimate of lower-frequency variables. Given the initial estimate, the factors and loadings at the highest frequency are estimated as described above in the single-frequency case.

2. **E**xpectation-Step: For each iteration j, given an initial estimate of the factors and loadings from the previous iteration $j-1$, compute an

[7] The idea of the selection matrix originates from the work by Jones (1980). The author estimated ARMA models with missing observation within a state-space framework with the Kalman filter. In the mixed-frequency modelling framework the idea was picked up by Greene, Howrey, and Hymans (1986), Zadrozny (1990), Stock and Watson (2002b), and Giannone, Reichlin, and Small (2008).

update of the highest-frequency or missing observations by the expectation of the $T \times 1$ X_i conditional on the observed data Y_i^- and the previous iteration factors and loading for variable i according to

$$\begin{aligned} \hat{Y}_i^{(j)} &= E(Y_i | Y_i^-, \hat{F}^{(j-1)}, \hat{\Lambda}_i^{(j-1)}) \\ &= \hat{F}^{(j-1)} \hat{\Lambda}_i^{(j-1)} + S_i'(S_i S_i')^{-1}(Y_i^- - S_i \hat{F}^{(j-1)} \hat{\Lambda}_i^{(j-1)}) \end{aligned}$$

3. Repeat step 2 for all series in the sample that contain missing values or that have to be transformed from the higher to the lower frequency.

4. **M**aximization-Step: The estimated highest-frequency observations are used to re-estimate the factors $\hat{F}_t^{(j)}$ and loadings $\hat{\Lambda}^{(j)}$ by an eigen decomposition of the covariance matrix $\hat{\Sigma}^{(j)} = (mT)^{-1} \sum \hat{Y}_\tau^{(j)} Y_\tau^{(j)\prime}$ according to Equation (2.5). The estimates of the factors and loadings enter step 2 above again until some convergence criterion is fulfilled.

The steps of the above EM algorithm provide estimates of the lower-frequency variables for the highest frequency in the data. Furthermore, estimates for missing observations are established. Schumacher and Breitung (2008) conduct a Monte Carlo study to investigate how well the EM algorithm can estimate monthly observations of GDP. The authors carry out two simulations. The first simulation addresses the estimation of monthly observations from quarterly and monthly data for different degrees of idiosyncratic noise in the data. In a second approach, the estimation of missing observations at the end of the sample is investigated. Considering the first point, Schumacher and Breitung (2008) find that the performance strongly depends on how informative are the data. The less informative the time series are with respect to the factors, the less precisely (in MSE terms) are the factors and the monthly observations estimated. The same conclusions can be drawn for the second part of the Monte Carlo study.

As we will see in later chapters, the presence of 'ragged-edge' data is prominent in extracting factors for nowcasting.[8] Besides the EM algorithm there

[8] The term 'ragged-edge' illustrates the imbalance at the end of a sample due to pub-

are other estimation methods available to handle missing observations at the end of the sample. A simple approach to solve the ragged-edge problem is provided by Altissimo, Cristadoro, Forni, Lippi, and Veronese (2006). They propose to realign each time series in the sample to obtain a balanced data set. Given the balanced data set, the standard approach to extract factors can be applied. A different approach is to build a parametric factor model in state-space form. Kapetanios and Marcellino (2006) estimate the factors using subspace algorithms, while Doz, Giannone, and Reichlin (2006) exploit the Kalman filter.

2.3 Bridge Equations and Linkage Models

2.3.1 Linkage Models

In the 1950s, large marcroeconometric models in different countries were developed to describe and forecast parts of National Accounts. Therefore they are based on quarterly frequencies. These models were designed primarily for short- and long-run forecasting and policy simulations. Their equations therefore generally have a solid theoretical base and exhibit desirable long-run equilibrium properties, but place less emphasis on short-run forecasting accuracy. Some of the variables are also available on a higher frequency (for example monthly). As higher frequencies potentially provide valuable (and timely) information regarding future economic development, it was sensible to use this additional information. In practice this was often achieved by adjusting the constant term of relevant equations (or by using a non-zero error term) to make the model outcome agree with the new information. Such a procedure to alter a model's solution is often referred to as 'judgemental modification' and is subject to criticism as being unscientific or ad hoc.[9] To avoid this problem the idea of combining quarterly (interpreted as

lication lags.

[9] Klein and Sojo (1989) provide a framework for adjusting these constant terms in a more scientific way. Their article is more general and contains practical advice on how to use high-frequency data (in bridge equations) in a large macroeconometric model.

long-run relationships) and monthly (short-run relationships) forecasts was introduced.

The basic idea is to model the higher frequency variables in separate (time series) models. The higher-frequency forecasts are combined with lower-frequency forecasts of the corresponding variables. The combined forecasts can be used to update the macroeconometric model.

The first approach to link quarterly and monthly information was proposed by Greene, Howrey, and Hymans (1986).[10] To illustrate the idea consider the reduced form of the lower-frequency (quarterly) variable

$$Y_t = PY_{t-1} + QX_t + V_t \tag{2.9}$$

where Y_t denotes the endogenous variables and X_t the exogenous variables of the system and $V_t \sim N(0, \Sigma_{VV})$ is the error term. The one-quarter-ahead forecast of Y_t is given by:

$$\hat{Y}_{t+1} = PY_t + QX_{t+1} \tag{2.10}$$

so that

$$Y_{t+1} = \hat{Y}_{t+1} + V_{t+1}. \tag{2.11}$$

Let $\tilde{Y}_{t+1}$ be the collection of H quarterly forecasts derived from the monthly model, where H denotes the number of common variables (which are available both at the higher and lower frequency). Because only some of the variables are available monthly, a 'selection' matrix denoted by θ picks out the elements of Y_t corresponding to $\tilde{Y}_{t+1}$

$$\theta Y_{t+1} = \tilde{Y}_{t+1} + W_{t+1} \tag{2.12}$$

where W_{t+1} is a $H \times 1$ vector of disturbances with mean zero and covariance matrix Σ_{WW}. The difference between the quarterly and the monthly forecast

[10] As large macroeconometric models were predominant at this time, they used the term 'outside information' for monthly information.

models is given by

$$Z_{t+1} = \tilde{Y}_{t+1} - \theta \hat{Y}_{t+1} = \theta V_{t+1} - W_{t+1} \tag{2.13}$$

The optimal combined forecast of quarterly and monthly data is given by

$$\overline{Y}_{t+1} = \hat{Y}_{t+1} + K Z_{t+1} \tag{2.14}$$

where

$$\begin{aligned} K &= \Sigma_{ZV} \Sigma_{ZZ}^{-1} && (2.15) \\ \Sigma_{ZZ} &= \theta \Sigma_{VV} \theta' + \Sigma_{WW} - \theta \Sigma_{VW} - \Sigma_{WV} \theta' && (2.16) \\ \Sigma_{ZV} &= \theta \Sigma_{VV} - \Sigma_{WV}. && (2.17) \end{aligned}$$

Generally, this approach can be extended to any mixtures of frequencies.

Howrey, Hymans, and Donihue (1991) criticize the approach by Greene, Howrey, and Hymans (1986) in one important aspect. The forecasts are linked with quarterly aggregates of the monthly forecasts and therefore the approach could not take advantage of the information within the quarter. We want to sketch the alternative pooling approach of Howrey, Hymans, and Donihue (1991), which utilizes within-quarter information.

Let $\hat{Y}_{t+1} = [\hat{Y}_{t+1} | \hat{y}_{t+1} | \hat{y}_{t+1-1/3} | \hat{y}_{t+1-2/3}]$ the vector containing the forecasts from the quarterly ($\hat{Y}_{t+1}$) and the monthly ($\hat{y}_{t+1-i/3}$) models. For combining the forecasts from both frequencies, the estimate of the covariance matrix Φ of the forecast errors corresponding to this forecast vector is necessary.[11] As not all common variables satisfy the aggregation condition[12], it is assumed that the vector of actual and quarterly and monthly values, Y_{t+1}, is drawn from a (normal) distribution with mean $\hat{Y}_{t+1}$ and covariance matrix Φ. Additionally an aggregation condition between quarterly and monthly data is necessary

[11] We refer to the appendix of Howrey, Hymans, and Donihue (1991) for the derivation of this covariance matrix.

[12] As, for example not all information is available for the current quarter.

such that

$$QY_{t+1} = 0 \tag{2.18}$$

where Q is the aggregation matrix. From standard formulae for the multivariate normal distribution one can obtain the conditional mean $\overline{Y}_{t+1}$ and the conditional covariance matrix, Ψ, of Y_{t+1} given $QY_{t+1} = 0$

$$\overline{Y}_{t+1} = \hat{Y}_{t+1} - \Phi Q'(Q\Phi Q')^{-1}Q\hat{Y}_{t+1} \tag{2.19}$$

$$\Psi = \Phi - \Phi Q'(Q\Phi Q')^{-1}Q\Phi \tag{2.20}$$

The corresponding diagonal elements of Ψ and Φ indicate the expected improvement in forecast accuracy resulting from this pooling procedure.

Fuhrer and Haltmaier (1988) derive formulae for obtaining minimum variance for pooled forecasts at the disaggregated level. They prove that pooling at the disaggregated level produces the same aggregated pooled forecasts as pooling the two forecasts at the aggregate level. This result holds only in-sample and can deviate out-of-sample.

The paper by Rathjens and Robins (1993) provides a different interpretation of how to link monthly and quarterly information in one model. The previous approaches combined forecasts from different frequencies, whereas Rathjens and Robins (1993) point out the usefulness of within-quarter information and want to utilize it for producing quarterly information. Consider a time series y_τ which is sampled at a monthly frequency but forecasts are generated with aggregated quarterly data. In a univariate approach this variable is forecasted with an AR or an ARIMA model. Rathjens and Robins (1993) suggest the introduction of a new variable x_t which is defined as

$$x_t = y_t - \frac{1}{3}\sum_{i=1}^{3} y_{t-1-i/3}$$

that is, the difference between the third month of the quarter and the simple average of the quarter. The quarterly forecasts are then generated with an autoregressive model with exogenous variables (ARX) or an integrated autoregressive moving average model with exogenous variables (ARIMAX)

$$A(L)y_t = B(L)\epsilon_t + C(L)x_t, \tag{2.21}$$

where $A(L)$, $B(L)$ and $C(L)$ are lag polynomials of finite order. Indeed it makes no sense to apply this idea to univariate models which are sampled at the higher frequency and are aggregated to the lower frequency. Therefore it is better to employ the new variable x_t in multivariate lower frequency time series models. This approach cannot be used for nowcasting in a strict sense, that is using information of the current quarter. As the approach can use information only up to $t-1$.

2.3.2 Bridge Equations

Klein and Sojo (1989) describe a regression-based current quarter GDP forecasting system in which GDP components of the National Accounts are modelled individually. In general, Bridge Models (BM) can be seen as tools to 'translate' the information content of short-term indicators into the more coherent and complete 'language' of the National Accounts. BM are linear dynamic equations where the aggregate GDP or, alternatively, GDP components are explained by suitable short-term indicators. In fact, BM can be specified either as different equations for the main GDP components (namely, private consumption, government purchases of goods and services, fixed investment, inventory investment, exports, and imports), or as a single equation for the aggregate GDP. In the first case, the model is labelled 'demand-side' BM (where GDP is predicted by the National Accounts income-expenditure identity); in the second case it is labelled 'supply-side' BM (where GDP is forecast by a single bridge equation).

In contrast to large structural macroeconomic models, BM are not concerned with behavioural relations. The choice of the BM explanatory variables is based on the researchers' experience and several statistical testing procedures, rather than on causal (that is structural) relationships.

In general auxiliary equations forecast the higher-frequency variable up to the end of the lower-frequency period under consideration. For instance, if we are in January and we want to forecast the first quarter, we forecast the independent indicators up to March. Then the 'filled' period of the higher-frequency variable can be temporal-aggregated to the lower frequency. These values are then plugged into the lower-frequency bridge equation time series model. The indicators can be forecasted with any preferred model. In many cases an AR, ARIMA, vector autoregressive model (VAR) or Bayesian vector autoregressive model (BVAR) are used. In contrast, to use a specific forecast model one can also encounter a no-change forecast, where the latest available information is used to represent the information content of the current period.

Before building a bridge model, the selection of indicators is a crucial step. First, monthly indicators must be updated in a timely manner (published before the BM dependent variable is released). Second, indicators must be reliable; that is they should not be revised substantially after they are first published. Alternatively, real-time data could be used, that is only information available at the forecast horizon (first estimates or final data) is used to calculate the forecast. Finally, indicators must be related to the dependent variable of the BM.

Bridge equations are rather useful for short-term forecasting (nowcasting). Forecasting the indicators over a longer horizon would transmit larger forecasting errors into the primary forecasting model due to iterative forecasting uncertainty of the higher-frequency variable.

2.4 State-space Approaches

The first two subsection deal with different state-space representations of a mixed-frequency VAR(MA) model, where all variables are endogenous. The state-space approach by Evans (2005) estimates an unobserved state interpreted as growth rates. There is no modelling of dynamic relationships between the variables. Finally we outline the combination of factor models with state-space approaches, where the lower-frequency target variable is forecast at a higher frequency.

2.4.1 A Mixed-frequency VARMA Model

Zadrozny (1988) proposed a general approach to handle different frequencies in a continuous time series model. Zadrozny (1990, 2008) extended this idea to discrete time. We sketch this approach and give some modelling examples; for details we refer to the original literature. We consider the general sampled at the highest frequency, ARMA(p, q) model with n time series for a_τ as

$$A(L)a_\tau = B(L)e_\tau \tag{2.22}$$

for $\tau = 1, \ldots, mT$, where $A(L) = A_0 - \sum_{k=1}^{p} A_k L^k$, $B(L) = \sum_{k=0}^{q} B_k L^k$, L is the lag operator, and e_τ is an $mT \times 1$ unobserved, normally distributed, white noise disturbance vector with zero mean and constant covariance matrix, that is $e_\tau \sim N(0, \Sigma_e)$, where $\Sigma_e = \mathrm{E}(e_\tau e_\tau')$. We can partition the number of time series into stocks (n_1) and flows (n_2), where $n = n_1 + n_2$.

We assume that the data are adjusted for mean values and other possible fixed (regression) effects. We have to assume some restrictions to identify the model. First, we must assume that the model is complete, that is A_0 is non-singular. Second, we assume that there are no identities in the data, which implies a non-singular probability distribution of the data. To ensure this we impose the restriction $B_0 \Sigma_e B_0' > 0$. Third, there is redundancy among A_0, B_0 and Σ_e. We adopt the normalization $A_0 = I_n$, $B_0 =$ lower triangular, and $\Sigma_e = I_n$, where I_n denotes the $n \times n$ identity matrix.

The model is cast in state-space form. A state-space system consists of a state and observation equation. The law of motion for the unobserved state x is given by

$$x_\tau = F x_{\tau-1} + G e_\tau, \tag{2.23}$$

where the matrices F and G contain the corresponding AR and MA coefficients. The state vector x_τ is constructed for stocks and flows with $r \cdot n^*$

elements, where $r = \max(p, q+1, \nu)$ and $n^* = n_1 + 2n_2$. The parameter ν denotes the maximum lag in any flow aggregation.[13] We define a $m \times 1$ vector w_τ of potential observations of a_τ. By "potential" we define that the vector w_τ is not yet adjusted for the actual observed values in period τ.

The observation equation is constructed in two steps. Let ζ_τ be an unobserved vector of observation errors which has the same dimension as w_τ. That is we observe w_τ as $w_\tau = x_{11\tau} + \zeta_\tau$ and , or, equivalently as

$$w_\tau = \Delta x_\tau + \zeta_\tau \tag{2.24}$$

where

$$\Delta = \begin{bmatrix} I_n & 0 & 0 & 0 \end{bmatrix}.$$

We assume that $\zeta_\tau \sim N(0, \Sigma_\zeta)$, where $\Sigma_\zeta > 0$, and $\mathrm{E}(\zeta_\tau e'_\tau) = 0$ and $\mathrm{E}(\zeta_\tau x'_1) = 0$ for all τ. As before it is convenient to re-parameterize Σ_ζ to R, where R is a lower triangular matrix and satisfies $RR' = \Sigma_\zeta$. Furthermore, let y_τ denote the $m_\tau \times 1$ vector of values of w in period τ which are actually observed, where $m_\tau \leq m$. Therefore we have $y_\tau = \Lambda_\tau w_\tau$, where Λ_τ is the $m_\tau \times mT$ selection matrix which picks out the observed elements of w_τ. Combining (2.24) and $y_\tau = \Lambda_\tau w_\tau$ we finally get the observation equation

$$y_\tau = D_\tau x_\tau + v_\tau, \tag{2.25}$$

where $D_\tau = \Lambda_\tau \Delta$ and $v_\tau = \Lambda_\tau \zeta_\tau$. The disturbance vector v_τ obeys the same properties as ζ_τ: $v_\tau \sim N(0, \Sigma_{v_\tau})$, where $\Sigma_{v_\tau} = \Lambda_\tau \Sigma_\zeta \Lambda'_\tau$ and $\mathrm{E}(v_\tau e'_\tau) = 0$, for all τ.

Having outlined the theoretical model we want to illustrate some parts with examples to enhance the understanding. Consider a bivariate VARMA(1,1) model with monthly and quarterly data where one variable is a stock (for example the indicator) and the other a flow (for example GDP growth). The

[13] For example, in case of monthly-quarterly data the parameter ν takes the value 2.

matrices F and G from equation (2.23) are then given by

$$F = \begin{bmatrix} \alpha_{11} & \alpha_{12} & 0 & 1 & 0 & 0 \\ \alpha_{21} & \alpha_{22} & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad \text{and } G = \begin{bmatrix} \beta_{11} & 0 \\ \beta_{21} & \beta_{22} \\ 0 & 0 \\ \beta_{31} & \beta_{32} \\ \beta_{41} & \beta_{42} \\ 0 & 0 \end{bmatrix}$$

where α_{ij} denote the AR and β_{kl} denote the MA coefficients. If both variables are observed than

$$\Lambda_\tau = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

and thus the D_τ matrix in equation (2.24) is given by

$$D_\tau = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 \end{bmatrix}$$

If only one variable is observed than

$$\Lambda_\tau = \begin{bmatrix} 1 & 0 \end{bmatrix} \qquad \text{and } D_\tau = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0. \end{bmatrix}$$

In a second example consider three variables, sampled at weekly, monthly, and quarterly intervals all observed as stocks. The matrices F and G are then given by

$$F = \begin{bmatrix} \alpha_{11} & \alpha_{12} & \alpha_{13} & 1 & 0 & 0 \\ \alpha_{21} & \alpha_{22} & \alpha_{23} & 0 & 1 & 0 \\ \alpha_{31} & \alpha_{32} & \alpha_{33} & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad \text{and } G = \begin{bmatrix} \beta_{11} & 0 & 0 \\ \beta_{21} & \beta_{22} & 0 \\ \beta_{31} & \beta_{32} & \beta_{33} \\ \beta_{41} & \beta_{42} & \beta_{43} \\ \beta_{51} & \beta_{52} & \beta_{53} \\ \beta_{61} & \beta_{62} & \beta_{63} \end{bmatrix}$$

Note the subtle difference to the previous example. The dimension of the matrix F is the same but due to the flow variable in the previous sample a

1 is put under each flow variable. If all variables are observed variables then

$$\Lambda_\tau = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad \text{and } D_\tau = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \end{bmatrix}$$

If only two variables (weekly and monthly) are observed then

$$\Lambda_\tau = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad \text{and } D_t = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \end{bmatrix}$$

If only one variable (weekly) is observed then

$$\Lambda_\tau = \begin{bmatrix} 1 & 0 \end{bmatrix} \quad \text{and } D_\tau = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

The Kalman filter (KF) provides a very powerful tool for maximizing likelihood functions. The KF has many possible implementations, see for example Anderson and Moore (1979) or Hamilton (1994). We use the Kalman Filter to compute the likelihood function with the innovation of a time series. Given the model, its parameters, and the data, the KF computes L_τ by iterating over the sampling times $\tau = 1, \ldots, mT$. At the start of iteration τ, $x_{\tau|\tau-1}, V_\tau$, and $L_{\tau-1}$ are given from the previous iteration. Given the values of these quantities, $L_{\tau-1}$ is updated with

$$M_\tau = \Sigma_{v\tau} + D_\tau V_\tau D'_\tau \tag{2.26}$$

$$\xi_\tau = y_\tau - D_\tau x_{\tau|\tau-1} \tag{2.27}$$

$$L_\tau = L_{\tau-1} + \ln \left| M_\tau \right| + \xi'_\tau M_\tau^{-1} \xi_\tau \tag{2.28}$$

where $|\cdot|$ denotes the determinant; $x_{\tau|\tau-1}$ and V_τ are updated with

$$K_\tau = F V_\tau D'_\tau M_\tau^{-1} \tag{2.29}$$

$$x_{\tau+1|\tau} = F x_{\tau|\tau-1} + K_\tau \xi_\tau \tag{2.30}$$

$$\Phi_\tau = F - K_\tau D_\tau \tag{2.31}$$

$$V_{\tau+1} = GG' + FV_\tau \Phi'_\tau. \tag{2.32}$$

Equations (2.26) to (2.32) are written in the order in which computations proceed. K_τ is called the Kalman gain matrix. Zadrozny (2008) discusses the conditions for numerical stability of this implementation of the KF.

To start the iterations, $x_{1|0}$ and V_1 must be specified, and, of course, $L_0 = 0$. As we deal with stationary data, the exact likelihood function is obtained when $x_{1|0} = \mu_x$, the unconditional mean of x, and $V_1 = \Sigma_x$, the unconditional covariance of x. When the data have been adjusted for means and fixed effects, $\mu_x = 0$ and Σ_x solves the (discrete-time, algebraic) Lyapunov equation

$$\Sigma_x - F\Sigma_x F' = GG' \tag{2.33}$$

When F is a stable matrix, (2.33) yields a unique, symmetric, positive semi-definite value of Σ_x. In sum, in the stationary case, the exact likelihood function is obtained when $x_{1|0} = 0$ and $V_1 = \Sigma_x$, where Σ_x solves Equation (2.33). For further approaches to the initialization problem see Ansley and Kohn (1985) or Durbin and Koopman (2001).

Given the estimated parameters we now outline how forecasts are generated by the mixed-frequency VARMA model. The state representation of a multivariate ARMA model for mixed-frequency data in Equations (2.23) and (2.25), and the white noise assumptions on their disturbance imply that

$$x_{\tau+k+1} = Fx_{\tau+k|\tau} \tag{2.34}$$

$$y_{\tau+k|\tau} = D_{\tau+k}x_{\tau+k|\tau} \tag{2.35}$$

for $k = 1, \ldots, K$, where $x_{\tau+k|\tau} = \mathrm{E}(x_{\tau+k|Y_\tau})$ and $y_{\tau+k|\tau} = \mathrm{E}(y_{\tau+k|Y_\tau})$. Let $\tau = 1, \ldots, mT_1$ the estimation period and let $\tau = mT_1+1, \ldots, mT_2$ denote the forecasting period. Set $\theta = \hat{\theta}$, Σ_ξ as prescribed, $x_{1|0} = 0$, and V_1 with (2.33). Given these values, iterate with the Kalman filter over $t = 1, \ldots, mT_1 - K$, to obtain $x_{\tau_1-K+1|mT_1-K}$. Given $x_{\tau_1-K+1|mT_1-K}$, iterate with (2.34) and (2.35), for $k = 1, \ldots K$. Using (2.30), update $x_{\tau_1-k+1|mT_1-k}$ to $x_{\tau_1-K+2|mT_1-K+1}$. Given $x_{\tau_1-K+2|mT_1-K+1}$, iterate with (2.34) and (2.35), for $k = 1, \ldots K$. Continue in this fashion for, $t = mT_1 - K + 2, \ldots, mT_2 - 1$ to obtain the desired forecasts.

Although the model is quite general in handling missing data, temporal aggregation, measurement errors, reporting delays, and revisions, Mittnik and Zadrozny (2005) and Chen and Zadrozny (1998) note that the model performs poorly or not at all on large models with many parameters. As an alternative Chen and Zadrozny (1998) propose an extended Yule-Walker equation method to estimate VARs with mixed frequencies. The authors proposed an optimal three-step linear instrumental variable method using GMM estimation techniques. The proposed approach can handle larger models compared to Kalman filter implementations. The illustration of the method is outside the scope of this thesis and therefore we omit it. Chen and Zadrozny (1998) conduct a small Monte Carlo study comparing the Kalman filter approach with the Yule-Walker implementation. Using simulated data, based on coefficients obtained from the data set employed in Zadrozny (1990), average coefficients, biases of the estimates, standard deviations of coefficients estimates and the root mean-squared errors of the coefficient estimates are compared. In general, the Yule-Walker approach yields similar results to the maximum-likelihood Kalman filter estimation. The results are based on a simulated mixed-frequency VAR(2). Higher-order models to investigate the outlined advantage of the Yule-Walker approach are not estimated. However, this approach has not been applied since the paper by Chen and Zadrozny (1998).

So far we have assumed that the data are stationary. Seong, Ahn, and Zadrozny (2007) extend the mixed-frequency VARMA approach to allow for cointegration relationships between variables, thus also for non-stationary variables. The model uses the same framework as outlined here. This approach is outside the scope of this thesis.

2.4.2 Another Representation of a Mixed-frequency VAR Model

The Zadrozny (1990) representation of a mixed-frequency VARMA is quite general and can handle all possible aspects of mixed-frequency data (revisions, publication lags etc.). Hyung and Granger (2008) present a different VAR representation to handle mixed-frequency data which they call the linked ARMA model. There are two important aspects to note. First, Hyung and Granger (2008) do not cite any paper of Zadrozny. And second, the authors speak always of ARMA models, but the whole analysis is based on a VAR representation. The model is not suitable for estimating general VARMA models, as the important identification aspect is not considered. We outline the basic model set-up and compare it to the approach by Zadrozny (1990).

Both approaches have in common that they assume that the model operates at the highest frequency. Furthermore both models are cast into a state-space system and are estimated with the Kalman filter. Suppose the data generating process for two stock variables follows a VAR(1) process. The state equation of the MF-VAR model by Hyung and Granger (2008) is given by

$$x_\tau = F x_{\tau-1} + \epsilon_\tau \tag{2.36}$$

where $F = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix}$, which contains the parameters to be estimated, and $E(\epsilon_\tau \epsilon'_\tau) = \begin{cases} Q & \text{for } \tau = t \\ 0 & \text{otherwise} \end{cases}$, where $Q = \begin{bmatrix} \sigma_{11} & \sigma_{12} \\ \sigma_{21} & \sigma_{22} \end{bmatrix}$.

The corresponding observation equation is given by

$$y_\tau = H_\tau x_\tau, \tag{2.37}$$

where $H_\tau = I_2$ if both variables are observed and $H_\tau = \begin{bmatrix} 1 & 0 \end{bmatrix}$ if only the higher-frequency variable is observed.

Let us now assume that the lower frequency variable is observed as a flow which is difference stationary (like GDP growth). Since the difference of the temporally aggregated flow variable is

$$\begin{aligned}\Delta x_{2,\tau} &= \frac{1}{3}(x_{2,\tau} + x_{2,\tau-1} + x_{2,\tau-2}) - \frac{1}{3}(x_{2,\tau-4} + x_{2,\tau-5} + x_{2,\tau-6}) \\ &= \frac{1}{3}(\Delta x_{2,\tau} + 2\Delta x_{2,\tau-1} + 3\Delta x_{2,\tau-2} + 2\Delta x_{2,\tau-3}\Delta x_{2,\tau-4})\end{aligned}$$

the corresponding state-equation is given by

$$\begin{bmatrix}\Delta x_{1,\tau}\\ \Delta x_{2,\tau}\\ \Delta x_{1,\tau-1}\\ \Delta x_{2,\tau-1}\\ \Delta x_{1,\tau-2}\\ \Delta x_{2,\tau-2}\\ \Delta x_{1,\tau-3}\\ \Delta x_{2,\tau-3}\\ \Delta x_{1,\tau-4}\\ \Delta x_{2,\tau-4}\end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0\\ a_{21} & a_{22} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0\\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0\\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0\\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0\\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0\\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0\\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0\\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0\\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0\end{bmatrix} \begin{bmatrix}\Delta x_{1,\tau-1}\\ \Delta x_{2,\tau-1}\\ \Delta x_{1,\tau-2}\\ \Delta x_{2,\tau-2}\\ \Delta x_{1,\tau-3}\\ \Delta x_{2,\tau-3}\\ \Delta x_{1,\tau-4}\\ \Delta x_{2,\tau-4}\\ \Delta x_{1,\tau-5}\\ \Delta x_{2,\tau-5}\end{bmatrix} + \begin{bmatrix}\epsilon_{1,\tau}\\ \epsilon_{2,\tau}\\ 0\\ 0\\ 0\\ 0\\ 0\\ 0\\ 0\\ 0\end{bmatrix} \tag{2.38}$$

and the corresponding observation equation

$$\Delta y_\tau = H_\tau \Delta x_\tau \tag{2.39}$$

where

$$H_\tau = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0\\ 0 & 1/3 & 0 & 2/3 & 0 & 1 & 0 & 2/3 & 1/3 & 0\end{bmatrix}$$

if both variables are observed and

$$H_\tau = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0\end{bmatrix}$$

if only the higher frequency variable is observed. The estimation is done with the Kalman filter and builds upon the idea of Zadrozny (1990). For the formulae we refer to Hyung and Granger (2008).

Compared to Zadrozny (1990), the approach of Hyung and Granger (2008) is more restrictive, for example it assumes that both variables have to be modelled as difference stationary processes, when one variable is observed as a flow. Furthermore the model assumes a strict relationship m to estimate the model and seems not to be able to handle any data patterns.

2.4.3 A Completely Different State-space Approach

Evans (2005) provides an extensive state-space model to estimate the current (unobserved) state of the economy. Essentially, the model claims to be able to handle reporting lags, temporal aggregation and mixed-frequency as in Zadrozny (1990). The main difference is, that the dynamic property of the model is not represented as a VAR system where all variables are endogenous. In contrast, the dynamics of the model centre on the behaviour of two partial sums, which define daily contributions to GDP growth in a specific quarter. Thus, the current estimate is updated as new information comes in.

The paper is in the spirit of the missing-observations approach of Harvey and Pierse (1984) for estimating missing observations in economic time series but it is far more general. The model of Evans (2005) is too extensive to lay out in this thesis. Furthermore the used notation is very cumbersome.

2.4.4 Factor Models and Mixed-frequency State-space Models

We have already outlined the importance of factor models in current macroeconomic forecasting. In addition to the static factor model outlined in section 2.2 we present now the dynamic factor model proposed by Doz, Giannone, and Reichlin (2006) and how the extracted factors can be used within a mixed-frequency state-space framework inspired by Mariano and Murasawa (2003). The following exposition follows closely Banbura and Runstler (2007) but is extended to a general mixed-frequency mixture m. We denote the higher-frequency time index τ and the lower-frequency index by t.

Consider a vector of n stationary high-frequency variables $x_\tau = (x_{1,\tau}, \ldots, x_{n,\tau})'$, $\tau = 1, \ldots, mT$, which have been standardized to mean zero and variance one.

The dynamic factor model is given by the equations

$$x_\tau = \Lambda f_\tau + \xi_\tau \qquad \xi_\tau \sim N(0, \Sigma_\xi), \tag{2.40}$$

$$f_\tau = \sum_{i=1}^{p} A_i f_{\tau-i} + \zeta_\tau \tag{2.41}$$

$$\zeta_\tau = B\eta_\tau \qquad \eta_\tau \sim N(0, I_q). \tag{2.42}$$

The second equation describes the law of motion for the latent factors f_τ, which are driven by a q-dimensional standardized white noise η_τ, where B is a $r \times q$ matrix, where $q \leq r$. Hence $\zeta_\tau \sim N(0, BB')$. We assume that the stochastic process for f_τ is stationary.

For the purpose of forecasting the lower frequency variable y_t, we introduce a latent interpretation of y_t, $\hat{y_\tau}$, which is related to the common factors by the static equation

$$\hat{y}_\tau = \beta' f_\tau. \tag{2.43}$$

In the m^{th} period the forecast for the lower frequency variable $\hat{y}_t$ is evaluated, as the average of the higher-frequency series

$$\hat{y}_t = \frac{1}{m} \sum_{i=1}^{m} \hat{y}_\tau \tag{2.44}$$

and defines the forecast error $\epsilon_t = y_t - \hat{y}_t$. We assume that ϵ_t is distributed with $\epsilon_t \sim N(0, \sigma_\epsilon^2)$. The innovations ξ_τ, ζ_τ, and ϵ_t are assumed to be mutually independent at all leads and lags.

Equations (2.40) to (2.44) are cast in a state-space form. As proposed by Zadrozny (1990) the higher-frequency variables observed within the current period are assumed to be missing. The state and observation equations are given by (for $p = 1$)

$$\begin{bmatrix} x_\tau \\ y_t \end{bmatrix} = \begin{bmatrix} \Lambda & 0 & 1 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} f_\tau \\ \hat{y}_\tau \\ \hat{y}_t \end{bmatrix} + \begin{bmatrix} \xi_\tau \\ \epsilon_t \end{bmatrix} \tag{2.45}$$

$$\begin{bmatrix} I_r & 0 & 0 \\ -\beta' & 1 & 0 \\ 0 & -\frac{1}{m} & 1 \end{bmatrix} \begin{bmatrix} f_{\tau+1} \\ \hat{y}_{\tau+1} \\ \hat{y}_{t+1} \end{bmatrix} = \begin{bmatrix} A_1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \Xi_{\tau+1} \end{bmatrix} \begin{bmatrix} f_\tau \\ \hat{y}_\tau \\ \hat{y}_t \end{bmatrix} + \begin{bmatrix} \zeta_{\tau+1} \\ 0 \\ 0 \end{bmatrix} \tag{2.46}$$

The aggregation rule (2.44) is implemented in a recursive way in equation (2.46), as from $\hat{y}_t = \Xi_\tau \hat{y}_{t-1} + \frac{1}{m}\hat{y}_\tau$, where $\Xi_\tau = 0$ for τ corresponding to the first high-frequency period of the lower frequency (for example the first month of a quarter) and $\Xi_\tau = 1$ otherwise. As a result, expression (2.44) holds for every m^{th} period of each lower-frequency period. The estimation of the model parameters $\theta = (\Lambda, A_1, \dots A_p, \beta, \Sigma_\xi, B, \sigma_\epsilon^2)$ is discussed in Doz, Giannone, and Reichlin (2006).

Aruoba, Diebold, and Scotti (2008) claim to move the state-space dynamic factor framework close to its high-frequency limit, and hence to move statistically-rigorous conditions analysis to its high-frequency limit. The approaches so far in the literature were modelled at monthly intervals as the highest frequency.[14] The basic idea is essentially the same as outlined before. Aruoba, Diebold, and Scotti (2008) describe a dynamic one-factor model evolving on a daily basis. The following model is considered

$$y_t^i = c_i + \beta_i x_t + \sum_{k=1}^{K} \delta_{ik} w_t^k + \sum_{n=1}^{N} \gamma_{in} y_{t-nD_i}^i + u_t^i \tag{2.47}$$

where y_t^i is the $i-th$ daily economic or financial variable at day t, which depends linearly on the unobserved economic state x_t (which follows an AR(p)

[14] Although the models are generally suitable to generate higher-frequency latent variables, Aruoba, Diebold, and Scotti (2008) were the first to demonstrate this within the state-space dynamic factor framework on a daily basis. Evans (2005) is the notable exception but the author does not estimate an unobserved economic state but the actual long-run US GDP growth on a daily basis.

process) and possibly on various exogenous variables w_t. The parameter D_i is a number that links the frequency to the observed y_t^i. Aruoba, Diebold, and Scotti (2008) demonstrate how stock and flow variables and a trend can be modelled within the given framework. The model is cast into a state-space framework and estimated with the Kalman filter. Aruoba, Diebold, and Scotti (2008) also outline the general problem with the approach operating at a daily basis. In their example four variables at different frequencies are used to extract the unobserved state of the economy on a daily basis. Due to the flow nature of some variables the authors have 94 state variables and more than 16,000 daily observations. One evaluation of the likelihood takes about 20 seconds. Therefore, the model becomes intractable for many variables and many factors.

2.5 Distributed Lag Models

In general distributed lag models are given by

$$y_t = \beta_0 + B(L)x_t + \epsilon_t,$$

where $B(L)$ is some finite or infinite lag polynomial operator, usually parameterized by a small set of hyperparameters.[15] In general, distributed lags models assume the same frequency. The following models are not distributed lag models in a strict sense, but they regress the lower frequency on the higher frequency variables.

First we outline the quite confusing approach of Abeysinghe (1998) of how to use mixed-frequency data in one model. Abeysinghe (1998) considers the following model

$$y_t = \beta_0 + \beta_1 x_t + \lambda y_{t-1/m} + u_t \tag{2.48}$$

where $t = 1, \ldots T$. Given this notation there is no real mixed-frequency, as x and y are sampled at the same frequency t. Instead, an artificial unobserved

[15] See for example the survey by Dhrymes (1971) or the textbook by Judge, Griffith, Hill, Lütkepohl, and Lee (1985), among others.

higher-frequency part of the dependent variable y is introduced ($y_{t-1/m}$). The author states that 'if y is observed annually and x is observed quarterly then $m = 4$', suggesting that the variable x is observed at a higher frequency.[16] Abeysinghe (1998) proposes to reformulate (2.48) to get rid of the (artificial) fractional lags[17]

$$y_t = \alpha + \beta_1 z_t + \lambda^m y_{t-1} + \upsilon_t \tag{2.49}$$

where

$$z_t = \sum_{l=0}^{m-1} \lambda^l x_{t-l/m}, \quad \upsilon_t = \sum_{l=0}^{m-1} \lambda^l u_{t-l/m}, \alpha = \beta_0 \sum_{l=0}^{m-1} \lambda^l.$$

So z_t is the weighted sum of the higher-sampled observations during one basic time unit t. Note that (2.49) captures only m lags of the higher-frequency variable, but it can be extended to more lags with some extra modelling effort. Model (2.49) is more meaningful when the dependent variable is a stock variable. When it is applied to flow variables autocorrelation is introduced. Abeysinghe (2000) outlines that this autocorrelation is rather small and modifies (2.49) to account for that issue. The derivation of Abeysinghe (1998) resulting in equation (2.49) is just a different transformation of the higher-frequency variable to the lower frequency. Only the weighting is different to standard aggregation approaches. With the proposed transformation one cannot handle more than two frequencies. The model is non-linear in the parameters and can be estimated via non-linear least squares.

Koenig, Dolmas, and Piger (2003) suggest a very simple approach to utilize monthly data in quarterly regressions. They propose to regress the quarterly values on unrestricted monthly values. A general model is given by

$$y_t = \alpha_0 + \sum_{i=0}^{k} \sum_{i=0}^{n} \beta_{ij} x_{jt-i/m} + \epsilon_t \tag{2.50}$$

where k denotes the number of included higher-frequency (for example monthly)

[16] Abeysinghe (1998) uses a different notation for m. We have adjusted the quote to preserve the notation throughout the thesis.

[17] Lagging the whole equation by l/m, multiply it by λ^l and sum up over the range of m.

indicators and n the number of included lags. By choosing separate sum formulae, any frequencies can be included in the model. Koenig, Dolmas, and Piger (2003) choose n equal to 4 for quarterly-monthly data based on theoretical reasoning. Let y_t denote the logarithm of a quarterly variable and suppose that x_τ is a monthly coincident indicator such that $y_t = (x_t + x_{t-1/3} + x_{t-2/3})/3$ for all t. Then

$$\begin{aligned} y_t - y_{t-1} &= \{(x_t - x_{t-1/3}) + 2(x_{t-1/3} - x_{t-2/3}) + 3(x_{t-2/3} - x_{t-4/3}) \\ &\quad +2(x_{t-5/3} - x_{t-1}) + (x_{t-4/3} - x_{t-5/3})\}. \end{aligned}$$

Thus, the quarterly growth rate is a weighted average of five monthly growth rates in the coincident indicator. It is possible to restrict the regression (2.50) to the theoretical weights, but this should be based on a statistical test.

2.6 MIxed Data SAmpling (MIDAS)

2.6.1 The Basic Model Set-up

The MIxed DAta Sampling (MIDAS) model of Ghysels, Santa-Clara, and Valkanov (2004) is closely related to distributed lag models. The approach regresses the dependent (lower-frequency) variable y on a distributed lag of x which is sampled at a higher frequency. The basic MIDAS model for a single explanatory variable, and one-step ahead forecasting, is given by

$$y_t = \beta_0 + \beta_1 B\left(L^{1/m};\theta\right) x_{t-1}^{(m)} + \epsilon_t^{(m)} \tag{2.51}$$

where $B\left(L^{1/m};\theta\right) = \sum_{k=0}^{K} B(k;\theta)L^{k/m}$ denotes a weighting function, and $L^{k/m}x_{t-1}^{(m)} = x_{t-1-k/m}^{(m)}$ represents a fractional lag operator. Again, t indexes the basic time unit, and m is the frequency mixture. The multi-step analogue is given by

$$y_t = \beta_0 + \beta_1 B\left(L^{1/m};\theta\right) x_{t-h}^{(m)} + \epsilon_t^{(m)}. \tag{2.52}$$

To given an example, consider the basic MIDAS regression with $m = 3$ and $K = 6$

$$y_t = \beta_0 + \beta_1 B\left(L^{1/3};\theta\right) x_{t-1}^{(3)} + \epsilon_t^{(3)} \tag{2.53}$$

where $B\left(L^{1/3};\theta\right) = \sum_{k=0}^{6} B(k;\theta)L^{k/3}$, and $L^{k/3}x_{t-1}^{(3)} = x_{t-1-k/3}^{(3)}$, so that:

$$\begin{aligned} y_t &= \beta_0 + \beta_1 \Bigg(B(0;\theta)x_{t-1}^{(3)} + B(1;\theta)x_{t-1-1/3}^{(3)} + B(2;\theta)x_{t-1-2/3}^{(3)} \qquad (2.54)\\ &\quad + B(3;\theta)x_{t-2}^{(3)} \ldots B(6;\theta)x_{t-3}^{(3)} \Bigg) + \epsilon_t^{(3)}. \end{aligned}$$

Consider y_t as the first quarter GDP growth for 2007, x_{t-1} is then the December 2006, $x_{t-1-1/3}$ the November and so forth.

There are several possible finite and infinite polynomials $B(k;\theta)$. The first one is

$$B(k;\theta) = \frac{\theta_0 + \theta_1 k + \theta_2 k^2 \ldots \theta_p k^p}{\sum_{k=1}^{K}(\theta_0 + \theta_1 k + \theta_2 k^2 \ldots \theta_p k^p)} \tag{2.55}$$

which is related to an 'Almon Lag' polynomial, where the order Q is typically small. Ghysels, Santa-Clara, and Valkanov (2004) parameterize $B(k;\theta)$ as:

$$B(k;\theta) = \frac{\exp(\theta_1 k + \theta_2 k^2 \ldots \theta_p k^p)}{\sum_{k=1}^{K} \exp(\theta_1 k + \theta_2 k^2 \ldots \theta_p k^p)} \tag{2.56}$$

which is called exponential Almon lag weighting function. The rationale of using an exponential transformation is that it guarantees positivity of the weights and it has the desirable feature of "zero approximation errors" (see Ghysels and Valkanov (2006)). A last specification is the Beta function which has only two parameters

$$B(k;\theta_1,\theta_2) = \frac{f(\frac{k}{K},\theta_1;\theta_2)}{\sum_{k=1}^{K} f(\frac{k}{K},\theta_1;\theta_2)} \tag{2.57}$$

where

$$f(x,a,b) = \frac{x^{a-1}(1-x)^{b-1}\Gamma(a+b)}{\Gamma(a)\Gamma(b)} \qquad \Gamma(a) = \int_0^\infty e^{-x}x^{a-1}dx. \tag{2.58}$$

Both specifications ensure positive weights and they sum up to unity. The first fact is important for financial (volatility) forecasting and the second one for identification of the parameter β_1 in (2.51). Almon Lag specification is theoretically more flexible than the Beta lag, since it depends on more parameters, but its specification (2.56) is able to generate rich weighting schemes similar to the Beta function, see Ghysels, Sinko, and Valkanov (2007). Figure 2.1 displays some examples of the exponential Almon lag weighting function demonstrating the variety of weighting schemes.

Ghysels, Santa-Clara, and Valkanov (2004) suggest that MIDAS models can be estimated under general conditions via non-linear least squares (NLS), (quasi-)maximum-likelihood (MLE) or general method of moments (GMM). Ghysels, Santa-Clara, and Valkanov (2004) employ a spectral estimation method, proposed by Hannan (1963a) and Hannan (1963b). But this estimator is rather complicated for applied work.[18] The GMM estimator applies the continuum general method of moments estimator proposed by Carrasco and Florens (2000).

Ghysels and Valkanov (2006) prove that non-linear least squares is a consistent estimator for the model in (2.51). The dimension of the numerical optimization procedure to obtain the parameters β and θ can be reduced by concentrating the least squares objective function with respect to β. For a given θ, β can be obtained by the least squares formula:

$$\beta = \left(\sum_{t=h}^{T} x_{t-h}(\theta) x_{t-h}(\theta)' \right)^{-1} \left(\sum_{t=h}^{T} x_{t-h}(\theta) y_t \right) \tag{2.59}$$

where $x_{t-h}(\theta) = \left[1, B(L^{1/m}; \theta) x_{t-h}^{(m)}\right]'$ and $\beta = (\beta_1, \beta_2)'$. Andreou, Ghysels, and Kourtellos (2007) compare the aspect of unbiasedness and efficiency of NLS and least-squares (LS) estimators where the latter one involves temporal aggregation. They show that the LS estimator is always less efficient than

[18] The estimator is also used in Ghysels and Valkanov (2006) but not in any other applied work.

Fig. 2.1: Examples of exponential Almon lag weighting functions

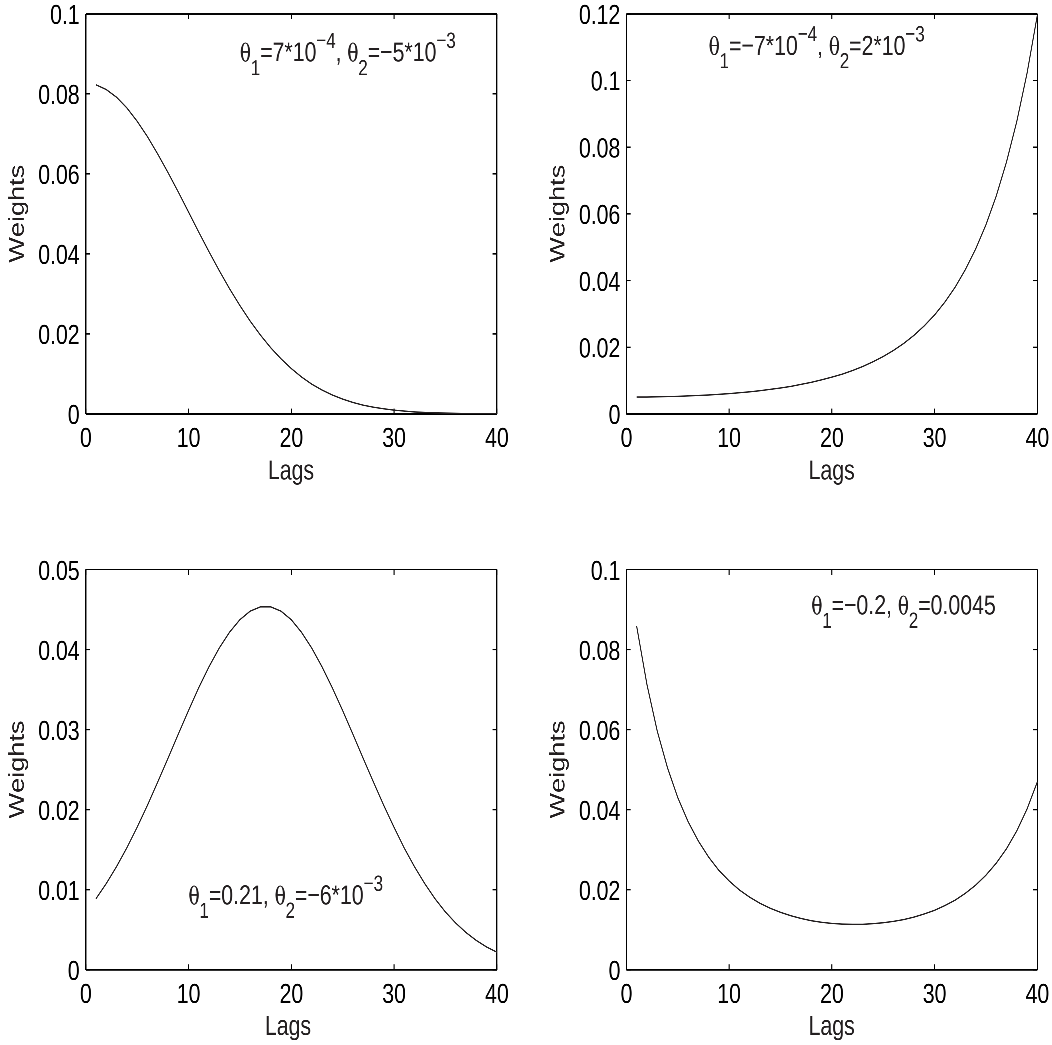

is the NLS estimator. In the current literature the NLS estimator is the preferred choice.

Ghysels, Santa-Clara, and Valkanov (2004) furthermore show that MIDAS regressions will always lead to more efficient estimation than the typical approach of aggregating all series to the least-frequent sampling. Furthermore the aggregation bias vanishes when some regressors are sampled more frequently.

2.6.2 Extensions of the MIDAS Approach

A natural extension of the MIDAS approach is to include autoregressive elements. Ghysels, Santa-Clara, and Valkanov (2004) show that efficiency losses can occur due to the introduction of lagged dependent variables. A naive extension by an AR term would result in a 'seasonal' polynomial, which can only be used if there are seasonal patterns in the explanatory variable. Clements and Galvao (2005) illustrate a solution to include and to consistently estimate the model with a simple autoregressive-distributive lag term. Adding a lower frequency of y_t, y_{t-1} to (2.51), results in:

$$y_t = \beta_0 + \lambda y_{t-1} + \beta_1 B\left(L^{1/m};\theta\right) x_{t-1}^{(m)} + \epsilon_t^{(m)} \tag{2.60}$$

Clements and Galvao (2005) suggest the introduction of autoregressive dynamics as a common factor

$$y_t = \beta_0 + \lambda y_{t-1} + \beta_1 B\left(L^{1/m};\theta\right)(1-\lambda L) x_{t-1}^{(m)} + \epsilon_t^{(m)}. \tag{2.61}$$

The multi-step analogue is given by

$$y_t = \beta_0 + \lambda y_{t-h} + \beta_1 B\left(L^{1/m};\theta\right)(1-\lambda L^h) x_{t-h}^{(m)} + \epsilon_t^{(m)}. \tag{2.62}$$

For estimation of the autoregressive MIDAS model (2.62), one takes the residuals ($\hat{\epsilon}_t$) of the standard MIDAS equation, and estimates an initial value for λ, say $\hat{\lambda}_0$, from $\hat{\lambda}_0 = \left(\sum \hat{\epsilon}_{t-h}^2\right)^{-1} \sum \hat{\epsilon}_t \hat{\epsilon_{t-h}}$. Then construct $y_t^* = y_t - \hat{\lambda}_0 y_{t-h}$

and $x^*_{t-h} = x_{t-h} - \hat{\lambda}_0 x_{t-2h}$, and the estimator of $\hat{\theta}_1$ is obtained by applying non-linear least squares to:

$$y^*_t = \beta_0 + \beta_1 B(L^{1/m}; \theta) x^*_{t-h} + \epsilon_t \tag{2.63}$$

A further variation of the MIDAS framework is to introduce a step function as in Forsberg and Ghysels (2007), which can be seen as a generalization of the heterogeneous autoregressive model introduced by Corsi (2003). Let us define $X_t(K, m) = \sum_{j=1}^{K} x^{(m)}_{t-j/m}$ as regressors, which are partial sums of the high-frequency $x^{(m)}$. The MIDAS regression with M steps is

$$y_t = \beta_0 + \sum_{i=1}^{M} \beta_i X_t(K_i, m) + \epsilon_t \tag{2.64}$$

The distributed lag pattern is approximated by a number of discrete steps. The more steps that are included, the less parsimonious is the model, which is one of the striking advantages of MIDAS regressions.

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

$$y_t = \beta_0 + \beta_1 \left(\phi B \left(L^{1/m}; \theta^+ \right) 1^+_{t-1} x^{(m)}_{t-1} + (2 - \phi) B \left(L^{1/m}; \theta^- \right) 1^-_{t-1} x^{(m)}_{t-1} \right) + \epsilon^{(m)}_t \tag{2.65}$$

where 1^i_{t-1} $(i = \{+, -\})$ denotes the indicator function, which takes the value 1 if $x_{t-1} > 0$ or $x_{t-1} < 0$, respectively. This formulation allows for a differential impact of positive and negative values of the explanatory variable x. The parameter ϕ is in the interval $[0, 2]$. This ensures that the sum of weights is 1 because the indicator functions are mutually exclusive of each of the positive and negative weight functions add up to 1. The coefficient ϕ controls the total weight between positive and negative impacts. A value of ϕ equal to one places equal weight on positive and negative impacts. Note that the parameters in the weighting function characterize the time profile of the weights from positive and negative shocks respectively.

A general univariate MIDAS regression involving more than two high-frequency variables is given by

$$y_{t+1} = \beta_0 + \sum_{i=1}^{K}\sum_{j=1}^{L} B_{ij}(L^{1/m_i};\theta)x_t^{(m_i)} + \epsilon_t \tag{2.66}$$

The case of $L > 1$ and $K = 1$ with $m_1 > 1$ corresponds to the case of having two or more polynomials for several time series sampled at the same frequency. In this case a mixture of polynomials is possible. This allows us to capture seasonal patterns or rich non-monotonic decay structures.

It is possible to extend the MIDAS regressions to semi- and non-parametric settings. Chen and Ghysels (2008) introduced semi-parametric MIDAS regressions that build upon the work of Linton and Mammen (2005) who propose the semi-parametric ARCH(∞) model.

$$y_{t+1} = \beta_0 + \sum_{i=1}^{K}\sum_{j=1}^{L} B_{ij}(L^{1/m_i,\theta} m(x_t^{(m_i)})) + \epsilon_t \tag{2.67}$$

where $m(\cdot)$ is an unknown function. Chen and Ghysels (2008) provide details on estimating such models. Furthermore it is shown that there is an efficiency gain compared to the single frequency case.

2.7 Comparison Between MIDAS and Mixed-frequency VAR

In this section we want to juxtapose the MIDAS model and the mixed-frequency VAR model (MF-VAR) without any empirical investigation. The approach to handle mixed-frequency data is completely different. On the one hand, the mixed-frequency VAR assumes that the model operates at the highest frequency; the corresponding values in the lower frequency are assumed to be periodically missing. On the other hand, the MIDAS approach regresses the lower frequency variable on the higher frequency variable. To avoid parameter proliferation a weighting function (in the sense of distributed lag models) is parameterized, where the weights add up to 1 to identify the parameters.

The MIDAS approach is very parsimonious in its basic specification where only four parameters have to be estimated independently of the number of included higher-frequency lags (for two variables). Nevertheless rich dynamic structures can be estimated due to the very flexible weighting functions. In contrast, for a bivariate VAR(p) $4p + 3$ parameters in the Zadrozny (1990) framework and $4p$ parameters in the Hyung and Granger (2008) framework have to be estimated. Thus, parameter proliferation can be a serious issue for higher-order VAR or by inclusion of more variables. Additionally a too-large frequency mixture (daily and yearly data) generates too many missing observations, which decreases the speed of the estimation dramatically. In contrast, the estimation of MIDAS models via NLS proceeds very fast.

Unlike the MIDAS approach, the mixed-frequency VAR can model feedback between the variables as all variables are endogenous. Furthermore it can interpolate the missing observations, that is estimated higher-frequency observations of the lower frequency variable.

One of the basic assumptions of the MIDAS model is, that the data are sampled equidistantly, that is the frequency-mixture m is fixed.[19] This assumption can be severe for larger frequency mixtures. Thus, publication lags, measurement errors, and aggregation issues are difficult to implement within the MIDAS framework. In contrast, the mixed-frequency VAR can handle any data pattern. The selection matrix in the Kalman filter can be flexible adjusted for the actual observed values at each time point.

A further difference between the approaches is the calculation of forecasts. Forecasts can be generated in two different ways: iterated (indirect or "plug-in") and directly. The iterated forecasts entail estimating an autoregression and then iterating upon that autoregression to obtain the multi-period forecast. The direct forecast entails regressing a multi-period-ahead value of the dependent variable on current and past values of the variable. The MIDAS approach generates a direct forecast, whereas the MF-VAR calculates

[19] Ghysels, Sinko, and Valkanov (2007) state that MIDAS models can handle unequally spaced data. They propose instead of using the lag operator $L^{1/m}$ to use an operator L^{τ}, where τ is real-valued instead of a rational number. But this approach has not been applied to real data sets so far.

iterative forecasts. To give an example, to forecast the first quarter of a year in December, the MIDAS approach calculates the first quarter forecast directly. The MF-VAR generates first the forecast for January and than iterates up to the March forecast. Choosing between iterated and direct forecasts involves a trade-off between bias and estimation variance. The iterated method produces more efficient parameter estimates than the direct method, but is prone to bias if the one-step-ahead model is wrongly specified. See Marcellino, Stock, and Watson (2006) for further details and references. The difference can become crucial when the frequency mixture is large. Forecasting quarterly time series with daily data, even only one-step ahead, can involve generating forecasts 60 to 90 steps ahead. For stationary transformations of variables, there exists a maximum horizon beyond which forecasts can provide no more information about the variable than is present in the unconditional mean. This point, called the 'content horizon' by Galbraith (2003), varies substantially across data series.

Finally, both approaches are able to augment the information set to account for intra-period information for lower frequency variable, that is information that becomes available in the m periods between t and $t+1$. For example, both models are able to include the February value of a specific indicator to forecast the first quarter of GDP.

The treatment of non-stationary data within the MIDAS framework remains unclear. We do not have found an statement concerning this issue. Within the state-space framework, Seong, Ahn, and Zadrozny (2007) offer a solution to deal with non-stationary data.

3. MIXED-FREQUENCY DATA AND MODELS: EMPIRICAL EVIDENCE

In this chapter we relate the approaches of data transformation and mixed-frequency time series modelling to the empirical evidence in the literature. We focus on those articles where an explicit comparison between single-frequency data or models and the mixed-frequency counterpart is made (although there are some exceptions); we omit structural investigations. We describe many articles in detail to outline the empirical strategy, and to see whether there is any advantage in using mixed-frequency data. We proceed in topical and chronological order, except for articles which are closely related. Each section starts with a short overview as a guideline through the review.

3.1 Transforming Mixed-frequency Data

In this section we leave out any temporal aggregation articles, as they are standard in applied forecasting. Concerning the interpolation aspect, we review early attempts to estimate structural macroeconomic models on a monthly basis. Recent attempts to elicit the unobserved state of the economy on a higher-frequency interval employ state-space models. We review these articles in the state-space section below.

Liu (1969) presents an early approach to combine data from different frequencies in one large structural macroeconomic model. With a related time series approach (estimated with OLS), the author uses monthly figures from a quarterly series of the US National Accounts for 1948-1964. The obtained monthly figures are used in a structural model for the US economy. The intention is to demonstrate the feasibility of a monthly model. The author investigates whether the higher monthly frequency may lead to more erratic estimations or misleading interpretations. He concludes that the model is feasible in any respect. However, forecasts are not generated from the model. A first forecast comparison of the structural monthly model by Liu (1969) to other (quarterly) structural models was conducted by Fromm and Klein (1973). They compare 11 structural econometric models of the United States in place at that time; two of them operate at annual and one at monthly intervals. There are difficulties in comparing the results from all the models (for example different simulation intervals). The model by Liu (1969) provides for the real GNP growth rate and the implicit price deflator a considerably lower RMSE compared with the other models. Liu and Hwa (1974) extend the monthly model in Liu (1969) and employ the related time series approach proposed by Chow and Lin (1971) to interpolate quarterly time series on a monthly interval. The monthly structural model of the United States (1954-1971) yields higher forecasting accuracy compared with two structural macroeconomic quarterly models in RMSE terms.[1]

Schumacher and Breitung (2008) apply the EM algorithm to obtain monthly factors from monthly and quarterly data to forecast German GDP.[2] The monthly factors are plugged in into a monthly VAR(p) (direct and indirect approach). The monthly GDP forecasts are then aggregated to quarterly values. The authors use both a real-time data set as well as final data. The data set consists of 41 monthly and 13 quarterly time series (1998-2005). In an out-of-sample forecasting exercise (one and two quarters ahead), the EM factor approach performs better on average than the AR(p) and the naive prediction no-change benchmark. Furthermore, there is almost no difference between real and final data vintages. This result is in contrast to the findings of Koenig, Dolmas, and Piger (2003) that the real-time aspect matters for US data.

[1] The Wharton Quarterly Mark II and the Quarterly model of the Bureau of Economic Analysis at that time.

[2] Another application of the EM algorithm can be found in Bernanke and Boivin (2003).

3.2 Nowcasting and Short-term Forecasts: Bridge Equations and Linkage Models

3.2.1 Linkage Models

Almost all articles reviewed in this section combine the forecasts from a large structural model of the US economy and from monthly forecasts. Shen (1996) and Liou and Shen (1996) are notable exceptions forecasting economic time series for Taiwan. Some results should be interpreted with care, as they are based on very few generated forecasts such as Fuhrer and Haltmaier (1988), and Howrey, Hymans, and Donihue (1991). Greene, Howrey, and Hymans (1986), Donihue and Howrey (1992), and Miller and Chin (1996) provide evidence to update the forecast of the current quarter as new information becomes available, whereas the others, Corrado and Greene (1988), Shen (1996), Liou and Shen (1996), and Rathjens and Robins (1993), generate forecasts when the quarter has elapsed. The updating of forecasts is in the spirit of the bridge equations, but the forecasts are generated via combination.

We outlined in the previous chapter the theoretical framework of how to combine forecasts from different frequencies by Greene, Howrey, and Hymans (1986). The authors also presented some empirical evidence from their modelling. Employing a small-scale version of the Michigan Quarterly Econometric Model consisting of 13 macroeconomic variables and corresponding equations, Greene, Howrey, and Hymans (1986) were the first to show empirically that ex post a gain in forecasting accuracy through combination is feasible. Four out of 13 variables (quarterly) are used as 'outside information', as these variables are also available monthly. They do not forecast the monthly variables in a separate model. Instead, as a new value (interpreted as information) becomes available, this is regarded as an updated forecast. As expected *a priori*, the authors show that as new information becomes available in a quarter, forecasting accuracy increases. Greene, Howrey, and Hymans (1986) were also the first to investigate the value of timely information for longer-horizon forecasts. They demonstrate that in-quarter information contain information even for a eight-quarter ahead forecast.

A further empirical example of linking monthly and quarterly forecasts was provided by Corrado and Greene (1988). Forecasts are generated for both a monthly and quarterly model for different target variables established at the Federal Reserve Board. The monthly forecasts are pooled, with the quarterly forecasts as 'add-factors'. The quarterly model's forecast errors are adjusted, conditional on the monthly information set. In the empirical application, the results are outlined for the linkage system in place at the Federal Reserve Board at that time. Several quarterly macroeconomic variables are forecast from 1972-1982. The authors show that the inclusion of monthly information reduces the RMSE compared with the pure quarterly model. The difference compared with Greene, Howrey, and Hymans (1986) is that forecasts are only generated after the whole quarter has elapsed.

Fuhrer and Haltmaier (1988) compare forecasts from a macroeconometric quarterly and monthly model from the Federal Reserve Board, as did Corrado and Greene (1988). Forecasts for eight US macroeconomic variables are only produced for the last quarter 1986 (the range of the data is 1972-1986). Therefore the results are not comparable to Corrado and Greene (1988). The authors distinguish three cases with different information structures (no information, one month, and two months of observations available for the current quarter). Fuhrer and Haltmaier (1988) do not compare their forecasts to realized values, but focus on the comparison of quarterly and monthly pooled forecasts. They show on a theoretical basis that these forecasts are as good in-sample as theoretically predicted.[3] Nevertheless, producing only one quarter forecast makes it difficult to interpret the results on a more general basis; it is natural to question the results based on such thin statistical evidence. In a second example, they forecast M1 for the United States at monthly and quarterly frequencies and calculate the RMSE (1983-1986). They find that pooling does not always provide lower RMSE compared with unpooled forecasts.

[3] See the previous chapter.

Howrey, Hymans, and Donihue (1991) combine monthly forecasts with forecasts from the Michigan Quarterly Econometric Model (MQEM) of the US economy (12 variables, 1954-1985). They provide both *ex post* as well as *ex ante* evidence from pooling forecasts from quarterly and monthly frequencies. The within-sample evidence demonstrates that the RMSEs of quarterly figures are reduced by linking them to monthly models, even when no monthly observations are available for the current quarter. The monthly forecasts are generated by a VAR(2) process. The out-of-sample data forecasting performances of the quarterly MQEM, the adjusted MQEM, and the combined quarterly-monthly model are compared. In the adjusted MQEM the constants are adjusted due to new monthly information. The results are ambiguous. First, as more information becomes available within a quarter, the lower is the corresponding RMSE. As two months of information are known in the quarter, the adjusted MQEM provides the lowest RMSE for six variables, the standard MQEM for two variables, and the pooled forecast for four variables. Thus, the inclusion of monthly information does not necessarily lead to lower RMSEs. We note that the results are based on only eight one-step-ahead forecasts. Therefore the results should be interpreted carefully.

Donihue and Howrey (1992) also combine monthly forecasts with forecasts from the Michigan Quarterly Econometric Model of the US economy. As with Howrey, Hymans, and Donihue (1991), the monthly forecasts are generated with a VAR model. The authors focus on inventory investment and show how the RMSE can be reduced by combining forecasts obtained at different frequencies. The forecast errors of other variables, only available quarterly, are also reduced. These results rely on 16 out-of-sample quarterly forecasts made between 1986 and 1990. They also provide evidence for improving forecast accuracy as more information becomes available during a quarter.

Shen (1996) employs the linkage approach to forecast 88 variables of the National Accounts for Taiwan. Basically these variables are forecast with a quarterly macro model serving as a benchmark. Thirteen variables are available on a monthly basis. These are forecast with either an ARIMA, VAR, or BVAR on a monthly basis. Shen (1996) shows that the combination of monthly and quarterly forecasts yields improved forecasts (in RMSE or MAE terms) compared with quarterly forecasts. This result holds even for the majority of variables available only on a quarterly basis. The author finds evidence that more 'outside' information does not necessarily improve forecasting accuracy. Liou and Shen (1996) provide essentially the same results as Shen (1996). The articles are in many ways similar. One difference is the evaluation sample which differs slightly. Another difference is the investigation of whether intra-quarterly information improves the accuracy of forecasts two quarters ahead. The answer is that there is a slight improvement. The authors also test for significant differences in forecasting performance. The employed test was proposed by Ashley, Granger, and Schmalensee (1980) but this tests for causality and not for equality of RMSEs.

In the previous chapter we outlined the idea by Rathjens and Robins (1993) of utilizing monthly data for generating quarterly forecasts. In the univariate framework they provide evidence that the inclusion of a variable reflecting within-quarter variance improves forecasting accuracy. They forecast 15 US time series which are available monthly at the quarterly frequency. However, the authors do not provide a comparison of the forecasts generated at the monthly level. The authors state that this approach is only useful for short-term forecasting. In a multivariate example the authors compare their approach to the macro model from Fair and Shiller (1990). Rathjens and Robins (1993) demonstrate that the within-quarter variance of the variable does contain information for forecasting the quarterly variable (US GNP).

Miller and Chin (1996) were first to combine monthly and quarterly forecasts from pure time series models. The authors estimate both for quarterly and monthly data vector autoregressions to obtain forecasts. The combination of the forecasts is just a pure linear combination of both forecasts. They forecast five quarterly economic time series (GDP, consumption of durable goods, federal government purchases, the civilian unemployment rate and the consumer price index). They provide evidence for the expected results, as the forecasting accuracy increases as new information within the quarter arrives.

Compared with a quarterly benchmark, the RMSEs are always lower.[4] In a second step Miller and Chin (1996) show how monthly information in the current quarter is useful for the forecast of the next quarter. They provide a conditional forecast, based on the current quarterly forecast. Monthly information provides useful data for the next quarter. In contrast, forecasting the monthly values up to the next quarter and then again combining monthly and quarterly forecasts does not help to decrease the RMSE. In a final section, they compare their results to the Blue Chip forecasts.[5] The within-quarterly forecasts are only marginally better than their Blue Chip competitor. This finding is similar to Trehan (1992).

3.2.2 Bridge Models

Summary

The knowledge of the current state of the economy ('nowcast') is especially important for central banks to conduct monetary policy. Therefore almost all articles in this section are written by authors affiliated either at the US federal system or the European Central Bank. Thus, the majority of articles forecasts US or Euro Area GDP. There are few applications to other single countries.

The basic goal is to improve the forecasting accuracy of the current quarter, as new higher-frequency information becomes available within the quarter. It is demonstrated that this exploitation of intra-period information reduces the forecasting error measures in almost all cases. The majority of applications are 'supply-side' bridge equations, where GDP is forecast by a single bridge equation. The articles by Parigi and co-authors (Parigi and Schlitzer (1995), Golinelli and Parigi (2007), Baffigi, Golinelli, and Parigi (2004)) employ 'demand-side' bridge equations where GDP is predicted by the National

[4] Miller and Chin (1996) test for significant improvements of the RMSE based on an approach proposed by Christiano (1989). As the limiting distribution of this test is not known, it has not been applied very often in empirical forecasting.

[5] The Blue Chip Survey is based on a panel of forecasts and is contained in a newsletter entitled, *Blue Chip Economic Indicators* published by Aspen Publishers.

Accounts income-expenditure identity. The choice of auxiliary equations, which forecast the higher-frequency variables, differs across applications.

US Data

Trehan (1989) is the first application of the bridge approach, to the best of our knowledge.[6] The author employs a two-step procedure. First, he selects predictors out of 16 variables, by minimizing an out-of-sample criterion. Eventually, the author employs monthly data for non-farm payroll employment, industrial production, and retail sales. The monthly variables are forecast with an AR and BVAR model. The quarterly forecasts are generated with an ARX model which includes the contemporaneous values of the predictors (no lagged values). The first target variable is annualized real GNP. As more information becomes available within a quarter, the RMSE decreases. Trehan (1989) compares the forecast to the Blue Chip forecast which is also released on a monthly basis. The model forecasts are slightly better than the Blue Chip forecasts in RMSE terms. Trehan (1989) also provides an early attempt at real-time forecasting by using preliminary data. Thus, in a second step, the target variable is the Advance GNP estimate. Here similar conclusions can be drawn, but now the model forecast is much better than the Blue Chip forecast.[7] In a last step, Trehan (1989) combines information from the model forecast and the advance real GNP to forecast final GNP. Some improvement can be seen. This does not hold for combinations with the Blue Chip forecast, where no improvement is found.

Trehan (1992) extends his previous study slightly. The evaluation period is different and the target variable is now GDP not GNP. The predictor variables as well as the basic results stay essentially the same. The only real extension is the investigation of whether the monthly indicators contain information beyond the current quarter. They do, but the improvement decreases with an increasing forecast horizon.

Trehan and Ingenito (1996) is a further extension of the article by Trehan

[6] But the author does not use this term.

[7] We note that the assessment is based only on four forecasts due to data availability restrictions.

(1989). They select the indicators out of 34 variables, based on various statistical measures. Finally, they employ non-farm payroll employment and real consumption data for forecasting GDP growth. In a recursive forecasting exercise they demonstrate that as new information becomes available, the RMSE declines for the current quarter. The monthly indicators are forecast within the AR, VAR and BVAR framework, where the last one works best. However, Trehan and Ingenito (1996) do not provide a comparison with other single-frequency data models.

Kitchen and Monaco (2003) extend the idea of Trehan and Ingenito (1996) to select monthly variables to relate them to forecasting quarterly GDP. But in contrast, they use monthly indicators in a separate model and combine the forecasts. Furthermore they do not forecast the independent variables. The actual, observed value within a quarter is assumed representative of the whole quarter. For data ranging from 1982 to 2003 they find the standard result that the RMSE declines with increasing information within the quarter.

Fitzgerald and Miller (1989) employ a simple bridge model to forecast US advance GNP. As monthly predictors, the authors use three measures of hours worked: the index of aggregate weekly hours of production or non-supervisory workers on private non-farm payrolls; the component of that series for goods-producing industries; and the component for service industries. These indicators are forecast with an ARX model. They conclude, that their simple approach performs better than the Bayesian vector autoregression employed by the Minneapolis Fed.

Braun (1990) is an example which departs from the pure time series approach and employs a theory-based model. He uses monthly preliminary labour market data to estimate the current quarter GNP. In a first approach the author employs an hour-to-output procedure with autoregressive components to obtain information for the whole quarter. In a second approach Braun (1990) estimated the unemployment rate with the help of Okun´s Law. The unemployment rate is then used to nowcast current GNP growth. The author also distinguishes between preliminary and final data. Both models, and a combination of them, do help to reduce the RMSE with each additional month within a quarter.

The paper by Klein and Park (1993) is written more in an anecdotal style. The authors outline their experience with the Pennsylvania system for forecasting national income and production accounts for the United States. The Pennsylvania model consists of a structural as well as a time series part. In the case of the latter, higher-frequency variables are forecast with an ARIMA model and aggregated (averaged) to a quarterly frequency. The aggregated variables are employed in a bridge equation to forecast national accounts variables. Furthermore Klein and Park (1993) demonstrate how to update current forecasts as new (high-frequency) information becomes available. They demonstrate their approach for the US national income and production accounts for only four quarters (1990:IV-1991:III). There is no comparison with realized values or with competitive models.

European Data

Bodo, Cividini, and Signorini (1991) use daily electricity data to nowcast Italian industrial production but they do not estimate daily production. They aim to forecast the current month before the actual release (40-55 days delay). The electricity data are aggregated and adjusted for actual trading days in a month. In the basic set-up the model fails to improve when compared with single benchmark models such as ARIMA and Holt-Winters algorithm. Thus, electricity data do not provide useful information in a single model to nowcast Italian industrial production. Nevertheless, in combining the benchmark models with the electricity data model they are able to yield the lowest RMSE when compared with all other models employed.

Parigi and Schlitzer (1995) employ the demand-side approach to forecast Italian GDP. They bridge ten different quarterly variables of the National Accounts ((non-)durable consumption, total consumption, investment in construction and equipment, inventory investment, total investment, exports, imports, and GDP) with different (leading) monthly indicators (survey data, financial indicators). These bridge variables are used to forecast Italian GDP in a rolling exercise; they do outperform quarterly ARX models. The authors employ the test proposed by Fair and Shiller (1990) to compare forecasting ability.

Rünster and Sédillot (2003) is an example of an extensive investigation of bridge models. Furthermore they are an example of attempting to account for the timely prediction of the target variable, that is accounting for the publishing date of the indicators. Their focus variable is the Euro area GDP (quarterly growth rate). As indicators of quantitative real activity indicators (such as industrial production), surveys, and composite indicators (as Euro-Coin). After investigating the predictive power at the single frequency using ARX models, Rünster and Sédillot (2003) compare the fit of the different indicators in the bridge equations.[8] The authors are the first who outline in detail how to forecast the indicators. The other papers in this chapter do not provide details about the fit and forecasting power of the auxiliary models. Rünster and Sédillot (2003) employ the following approaches: naive (no-change), ARIMA, a univariate and multivariate structural time series model and a VAR. There is no general advice on which approach is the best, it depends on the specific indicator. The main finding of the paper is that the bridge equations continue (compared with a full information set-up) to outperform naive and ARIMA forecasts when based on a limited number of monthly observations. Furthermore current quarter information also contains information for forecasts of the next quarter.

Baffigi, Golinelli, and Parigi (2004) employ both supply-side and demand-side bridge models to forecast GDP for the euro area GDP and the three main countries of the euro area (Germany, Italy, and France). The paper shows that national bridge models are better than the benchmark models. In addition, euro area GDP and its components are more precisely predicted by aggregating national forecasts.

Golinelli and Parigi (2007) is essentially an extension of Baffigi, Golinelli, and Parigi (2004). In addition to forecasting the Euro area, the countries of the G7 and the EU area are forecast with bridge equations. The authors employ a rolling procedure, where both the model specification and the dimension of

[8] The ARX model is interpreted as the bridge equation.

the estimation sample is kept fixed. The authors select a rather large set of indicators (between five and eight) for each country in an ARX model. The auxiliary models for indicators are an AR(12), an ARIMA model with an automatic model selection procedure (general-to-specific) and a VAR model. To test for differences in the RMSE, the test proposed by Giacomini and White (2006) is used. Independently of the countries and auxiliary models, if two months of the quarter are known, the bridge models yield statistically lower RMSE than the AR(5) benchmark model (except for France). Given that one month is known, the differences are only significant for Germany and Italy.

Diron (2006) analyses the predictive power of real-time vintage data and pseudo real-time data.[9] The target variable is the Euro area GDP growth, but the real-time data set is only available from 2001 to 2004. Thus, Diron (2006) notes that the results should be interpreted with caution. Seven indictors are employed (industrial production, retail sales, new car registrations, construction production, value added in services, EuroCoin and OECD survey data). The monthly indicators are forecast with an AR(6).[10] The structure of the eight bridge equations is kept constant and only the coefficients are re-estimated on a rolling basis. Diron (2006) states, that given the size of the revisions of the indicators, the assessment of reliability of short-term forecasts on revised series could potentially give a misleading picture. Nevertheless, by averaging across all bridge equations, forecasts of individual quarters tend to be similar whether they are based on preliminary or revised data. Finally, Diron (2006) investigates the sources of forecasting errors. The main sources are from the extrapolation of the monthly indicators. Meanwhile, revisions to the monthly variable and GDP growth account only for a small share of the overall forecasting errors.

---

[9] Pseudo real-time experiments are 'real-time' in the sense that they mimic the actual real-time situation faced by the forecasters in terms of the schedule of data releases and thereby of monthly indicators. However, pseudo real-time experiments do not reflect the genuine real-time situation, to the extent that they use current estimates of GDP and the monthly indicators that are a post-revision of the series.

[10] Diron (2006) also experimented with VARs and BVARs but failed to improve the forecast results. This stands in contrast to the findings in Rünster and Sédillot (2003).

The most recent application of bridge models is provided by Golinelli and Parigi (2008). The authors forecast Italian GDP growth in real-time by establishing a complete new real-time data set for the Italian economy. In contrast to most other papers employing the bridge approach, Golinelli and Parigi (2008) employ many versions of the general-to-specific modelling approach, depending on whether lagged depended variables are included, variables are differenced or in levels, inclusion of simultaneous and (or) lagged explanatory indicators and many more. The indicators are forecast with an AR(5) model to fill out the quarter. The quarterly forecasts are calculated using monthly data within a regression approach . The results are presented for one and four quarters ahead and are not based on available monthly vintages as in other bridge model applications. Compared with random walk with drift and ARIMA, not all general-to-specific approaches outperform the benchmark models. The difference between RMSEs is assessed by the statistical test of Giacomini and White (2006). Only bridge models with both lagged and simultaneous explanatory indicators, where the one-step-ahead predictions are obtained assuming that all of the simultaneous regressors are known (nowcast), produce significantly lower RMSE´s when compared with the random walk model. However, the relevance of indicators tends to vanish at longer forecasting horizons (four quarters). Finally, the authors test the rationality of the first GDP release with the test by Fair and Shiller (1990). For the Italian economy, the first GDP data releases appear to be rational forecasts of the final outcome, but not of the latest available data.

Special Application

Perez (2007) is a an exception in the application of bridge models. All other articles in this review forecast GDP or other parts of the National Accounts. Perez (2007) forecasts general government fiscal deficits in the overall euro area and for most of its members. The target variable is sampled at annual intervals, whereas the indicators (different between countries) are sampled at quarterly intervals. The auxiliary quarterly forecast model is chosen to maximize the forecast performance (a choice can be made between random walk, ARIMA and unobserved components). The forecasting model is a vector error correction model. As a benchmark it serves the official forecast of the EU and a random walk forecast. In addition to the single forecasts, the official forecast and the indicator forecasts are combined via the regression approach. In the out-of-sample forecasting exercise the results differ across countries. This approach provides the lowest RMSE. At least all of them are better than the random walk forecast. Perez (2007) concludes that existing intra-annual fiscal information should be used and included in the preparation of official estimates of government deficits.

3.3 State-space Approaches

The first subsection deals with applications that extract an unobserved state of the economy, as in Mariano and Murasawa (2003), Nunes (2005), Evans (2005), and Aruoba, Diebold, and Scotti (2008). These articles demonstrate the use of mixed-frequency data to estimate the current (unobserved) state of the economy. They are not designed a priori for forecasting purposes, but to detect (*ex post* and *ex ante*) turning points. The approaches by Giannone, Reichlin, and Small (2008) and Banbura and Runstler (2007) estimate GDP on a monthly basis within a state-space framework. These approaches are designed for forecasting and they account for the actual release date of the indicators.

In the third subsection we present the VAR approaches which estimate dynamic relationships between target and indicator variables. In contrast to the other approaches they do not account for actual release dates and do not employ higher frequencies than monthly intervals.

3.3.1 Extracting an Unobserved State of the Economy

Mariano and Murasawa (2003) extract a new coincident indicator out of one quarterly (Real GDP) and four monthly time series (employees on non-agricultural payrolls, personal income less transfer payments, index of industrial production, and manufacturing and trade sales). These are the same variables, except for GDP, used by Stock and Watson (1989). Therefore the extracted coincident index exhibits a strong correlation with the Stock-Watson experimental index. Mariano and Murasawa (2003) compare the turning point (NBER references) detection performance of their own and the Stock-Watson index. They investigate the possibility of turning point detection comparing it with the official NBER dates. The results are not convincing at all. One trough is detected earlier and some peaks later than the Stock-Watson index.

Nunes (2005) applies the same approach as Mariano and Murasawa (2003) for nowcasting the Portuguese economy. The author demonstrates how the CI (constructed out of five quarterly and six monthly indicators) with an estimated ARX model can significantly reduce the RMSE compared with an AR(1) benchmark model. Crone and Clayton-Matthews (2005) use three monthly and two quarterly series to extract economic indexes for the 50 US states. But there is no forecasting comparison.[11]

Evans (2005) was the first to track the status of an economy on a daily basis. He estimates log GDP and GDP growth of the United States based on three quarterly GDP series (advanced, preliminary, and final) and 18 monthly series from 1993 and 1999. The contribution of each variable to the daily GDP estimate is based on the actual release date. In addition, Evans (2005) also uses the expected values of GDP growth.[12] The state-space model is estimated with the Kalman filter and contains 63 parameters. Evans (2005) claims that despite the short sample the parameters are estimated with high precision. The real-time estimate of US GDP growth displays a good deal of high-frequency volatility. Furthermore, the gaps between the real-time estimates and *ex post* GDP data are on occasion both persistent and significant. The real-time estimates are on average lower than are the *ex post*

[11] Note that Nunes (2005) and Crone and Clayton-Matthews (2005) do not cite the article by Mariano and Murasawa (2003).

[12] The market expectations and the release data are obtained from International Money Market Service (MMS). MMS asked about forty money managers on the Friday of the week before the release day. Many earlier studies have used MMS data to construct proxies for the news contained in data releases (see for example Andersen, Bollerslev, Diebold, and Vega (2003)).

final GDP data. Evans (2005) suggests that the *ex post* data should not be viewed as a close approximation to what was known at the time. Finally, the model estimates reveal that the monthly data releases contain information that is useful for forecasting the future path of GDP. Nevertheless, Evans (2005) does not provide a standard forecasting exercise, where forecast and realized values are compared. Instead, the contributions of each economic variable to the real-time estimate variance are provided.

Aruoba, Diebold, and Scotti (2008) employ a one factor dynamic model to extract the unobserved state of the US economy. They claim that they provide a "call to action" to demonstrate the feasibility of measuring macroeconomic activity in real-time by inclusion of real high-frequency data. The article is an extension of Mariano and Murasawa (2003) but avoids approximations. Aruoba, Diebold, and Scotti (2008) use four variables: the yield curve (daily, stock), initial claims for unemployment insurance (weekly, flow), employees on non-agricultural payrolls (monthly, stock) and real GDP (quarterly, flow). The unobserved variables and the indicators follow an AR(1) process at their observational frequencies. As stated in the previous chapter the model, even with only four variables, is computationally very intensive. The extracted factor broadly coheres to the NBER business cycle dating chronology. The inclusion of high-frequency daily data does not really change the picture compared with monthly indicators, but it is available sooner. Aruoba, Diebold, and Scotti (2008) do not provide a quantitative assessment of forecasting performance for turning points or for point forecasts. But the elicited indicator is a coincident index and not a leading indicator.

3.3.2 Factor Models and State-space Forecasts

Giannone, Reichlin, and Small (2008) were the first to obtain factors and plug-them into the state-space frame work to generate forecasts of monthly GDP . The authors nowcast US GDP. They extract monthly static and dynamic factors from a real-time data set consisting of more than 200 time series.[13] As the ragged-edge data problem is strongly present in the data set, the authors apply the state-space approach of Doz, Giannone, and Reichlin (2006). The factors obtained are used within a state-space framework to forecast monthly GDP. The novelty of Giannone, Reichlin, and Small (2008) is that they demonstrate that as new information becomes available inside the quarter, the forecast uncertainty falls. The authors define a stylized calendar where in each month the releases are ordered to 15 release blocks. As the new block of information is released the factors are estimated and the nowcast of the current quarter is updated. The novelty is using exact calendar dates to track the forecast performance, whereas early approaches implicitly assumed that information arrives at end of each corresponding month. After the first month in a quarter the out-of-sample forecast is better than the random walk benchmark model, and furthermore they outperform the Survey of Professional Forecasters. The authors assess the impact of different releases on the forecast accuracy.[14]

Banbura and Runstler (2007) employ the same framework as Giannone, Reichlin, and Small (2008) but the focus is set on publication lags and not on real-time data issues. Furthermore, the target variable is the euro area GDP growth. Instead of assessing the contribution of single variables to the forecast performance, Banbura and Runstler (2007) investigate how the weights contributed to the recursive forecasts. The authors confirm the standard finding of using within-quarter information and show that the results obtained from monthly factor models are better than the autoregressive quarterly benchmark.

3.3.3 Mixed-frequency VAR

Zadrozny (1990) forecasts quarterly US GDP with monthly employment as

[13] This data set also contains quarterly time series, but in contrast to Schumacher and Breitung (2008) the authors linearly interpolate the quarterly figures to a monthly interval. The data set also contains daily data but these are averaged to monthly figures.

[14] The working paper version (Giannone, Reichlin, and Small (2006)) conducts the same analysis for US inflation with similar results.

a predictor. The data range from 1954 to December 1988 (with 1979 to 1988 as the evaluation period). The author estimates an MF-VARMA(1,1) to calculate forecasts.[15] Compared with a no-change and AR(1) benchmark for both variables, the RMSE is lower. Zadrozny (1990) also compares his forecasts with those ones generated by Roberds (1988) (only slightly better) and McNees (1986) (only better for the first quarter). Concerning the now-casting aspect, Zadrozny (1990) compares his results to Trehan (1989) where the MF-VAR is only better for one month ahead.

Mittnik and Zadrozny (2005) apply the mixed-frequency VAR to forecast German real GDP (1970-2003).[16] They investigate the properties of the approach in more detail. They state that forecasts are feasible if variables are in compatible cyclical form and not too many parameters have to be estimated. Real industrial production, Ifo Business Climate and Ifo Business expectations for the next six months are used as indicators. Mittnik and Zadrozny (2005) focus on a VAR(2) model with two to four parameters (GDP and industrial production are always included). In general they find that monthly models produce better short-term GDP forecasts, whereas quarterly models produce better long-term forecasts. The Ifo variables improve the quarterly short-term GDP significantly.[17]

Seong, Ahn, and Zadrozny (2007) conduct only a small forecasting exercise. Using trending data of US GDP and the consumer price index the authors estimate a monthly cointegrated VAR system and calculate 12 out-of-sample forecasts. Compared with single-frequency forecasts, the high-frequency model performs better for both variables. But this judgement is based on only four quarterly comparisons.

Hyung and Granger (2008) apply their linked-ARMA model to forecast US GNP. The only indicator is industrial production. The authors compare their approach with a standard AR model and the approach of Rathjens

[15] Zadrozny (1990) starts with an MF-VARMA(3,1) and reduces this model to an MF-VARMA(1,1) based on a corrected AIC and tests for white noise in the residuals.

[16] The data are filtered to account for the structural break due to the reunification of Germany.

[17] All model combinations are at least better than the no-change benchmark.

and Robins (1993) using a within-quarter information variable. The data sample ranges from 1947 to 2004 and the evaluation sample starts in 1990. Forecasting up to four quarters ahead, the results of the linked-ARMA models are rather poor. There are only very small improvements in the RMSE compared with the two benchmark models. An improvement in forecasting accuracy can only be gained by forecasting within the quarter (nowcasting). Assuming that one or two months of information of the quarter being forecast are known the RMSE of the linked-ARMA model is statistically significantly (Diebold-Mariano-Test) lower than the two benchmark models. The results by Hyung and Granger (2008) demonstrate that contemporaneous indicators (like industrial production) combined with a mixed-frequency model can improve the nowcast. As they do not use leading indicators, they cannot demonstrate whether their approach is useful for longer-horizon forecasts.

3.4 Distributed lag models

Abeysinghe (1998) applies his transformation approach given in equation (2.49) and forecasts the GDP of Singapore based on monthly external trade data. Compared with an AR(1) and a quarterly-frequency benchmark model, the RMSE can be reduced. We have to note that the article suggests that a fixed-estimates model is used to produce the forecasts. Thus the model is not updated with every recursion.

Koenig, Dolmas, and Piger (2003) apply the unrestricted distributed lag model to forecast current-quarter real GDP using monthly measures of economic activity. They follow Trehan (1992) and use annualized percentage changes in non-farm employment, industrial production, and real retail sales as indicators. The focus of the paper is more on the real-time aspect of the data, rather than on the mixed-frequency structure of the data. They compare three strategies for employing real-time data. Compared with a naive forecast, an autoregression and the Blue Chip consensus forecast, the use of real-time data vintages perform best in RMSE terms. Whereas for pseudo-real-time data and final data, the Blue Chip forecast is not outperformed.

3.5 MIxed DAta Sampling (MIDAS)

Most empirical applications of the MIDAS model focus on financial data. Furthermore there is often no explicit comparison between single- and mixed-frequency models but only on the exploration of the extension to the use of mixed-frequency data. As this approach is rather new we include those studies in our review. To guide the reader, we summarize the most important facts concerning model specification in a table. Table 3.1 summarizes for each application the employed weighting functions, whether the weighting function is restricted or not, the number of included lags, and the mixed-frequency data used. In general we can say that no article states whether an explicit model selection procedure is employed or not. From a personal perspective, it appears that the specification used is chosen ad hoc or is based on experience.

3.5.1 Financial Applications

Ghysels, Santa-Clara, and Valkanov (2005) investigate the inter-temporal relation between the conditional mean and the conditional variance of the aggregate stock market returns. They find a significant positive relationship suggested by the intertemporal capital asset pricing model of Merton (1973). This result stands in contrast to the previous works on this issue which find insignificant positive relationships or even a significant negative one. Ghysels, Santa-Clara, and Valkanov (2005) use monthly and daily market return data from 1928 to 2000. The MIDAS model is employed to estimate the conditional variance of monthly returns based on prior daily squared return data. The weighting function is parameterized as the exponential Almon lag with two parameters. They restrict these parameters to ensure declining weights. They allow for a maximum lag of 252 trading days (approximating a year).[18] The model is estimated via quasi-maximum likelihood. The authors compare their results with the rolling moving window approach of French, Schwert, and Stambaugh (1987) and GARCH-in-mean models. In the first

[18] The authors state that their results are not sensitive to the chosen lag length.

approach, the weights are constant and inversely proportional to the window length. French, Schwert, and Stambaugh (1987) fail to find a positive significant relationship. Furthermore their model is sensitive to the choice of the window lengths. The GARCH-in-mean model uses monthly instead of daily squared returns. Again, only a positive but insignificant relationship is identified. Then all three models forecast realized volatility. In general the MIDAS model forecasts, on average, more accurately and without bias. The GARCH approach exhibits a small upward bias. Finally the rolling window approach produces far more dispersed forecasts. Ghysels, Santa-Clara, and Valkanov (2005) conduct several robustness checks (for example asymmetric shocks) but in general a positive significant relationship between the conditional mean and the conditional variance is detected with the MIDAS approach.

Ghysels, Sinko, and Valkanov (2007) build upon the previous study. They use a different and shorter data set (Dow Jones Index, April 1993 to October 2003) but obtain their returns from a 5-minute price series. The authors investigate whether the results of Ghysels, Santa-Clara, and Valkanov (2005) still hold using different horizons ($h = \{5, 10, 15, 22\}$ days), different measures of volatility (additionally to squared returns, absolute returns, daily ranges, realized volatility, and realized power)[19] and employing the Beta weighting function.[20] For the weighting function they allow for only 30 lags of the higher frequency compared with 252 in Ghysels, Santa-Clara, and Valkanov (2005).[21] In general the results are similar to those found in Ghysels, Santa-Clara, and Valkanov (2005). Comparing the two different weighting functions the authors conclude that the Beta polynomial could be a better choice for higher frequency models, whereas the exponential Almon lag polynomial

[19] These extension build on results by Ghysels, Santa-Clara, and Valkanov (2006).

[20] We note that the authors are unclear about how the model was estimated. In the text they state that the MIDAS regressions are estimated using NLS (p. 73), whereas in Tables 1 and 2 (p. 74 and 75) they state that they employed quasi-maximum likelihood. As shown in chapter 2 both estimation approaches can be used to estimate MIDAS regressions.

[21] We presume this decision is due to the trade-off between included lags and observations available for estimation (see chapter 5). The authors do not go into detail on this issue.

could be a better choice for the lower frequency. There is no comparison with single-frequency models.

León, Nave, and Rubio (2007) extend the previous analysis on Mertons theoretical predictions to European equity markets. Starting with GARCH-type models the authors fail to find evidence of a positive significant relationship for the stock market indices in France, Germany, Spain, the United Kingdom and the stock index Eurostoxx50. Applying the MIDAS model, the results favour the positive risk-return relationship, except for the UK data. In a second step they apply the asymmetric MIDAS model and the results stay the same. Finally they find some evidence in favour of the two-factor inter-temporal capital asset pricing model using a bivariate MIDAS model.

Kong, Liu, and Wang (2008) investigate the risk-return trade-off for Chinese stock markets. The authors also compare the MIDAS with the GARCH and rolling windows approaches. The authors do not use mixed-frequency data but focus only on daily data. The MIDAS model is used to specify a parsimonious model. They fail to find any evidence in favour of the risk-return trade-off for the whole sample (1993-2005). The existence is found with a sub-sample (2001-2005). In forecasting the conditional variance, the GARCH forecasts seem to outperform the MIDAS approach but the authors do not provide forecast accuracy measures.

Ghysels, Santa-Clara, and Valkanov (2006) forecast future volatility with different regressors. Here, volatility is defined as quadratic variation. The following regressors are used: lagged quadratic variation, lagged returns, lagged absolute returns, daily ranges and realized power.[22] Volatility is measured at daily, weekly, bi-weekly, tri-weekly, and monthly frequency, whereas the forecasting variables are available at daily and higher frequencies. As a benchmark model they employ an ARFI(5,d) model (autoregressive fractional integrated) proposed by Andersen, Bollerslev, Diebold, and Labys (2003) (labelled ABDL). The ABDL model forecasts daily volatility and adds up these forecasts to obtain weekly volatility forecasts. In a strict sense it is

[22] Ghysels, Santa-Clara, and Valkanov (2006) also consider the log-version of all variables.

a single-frequency model. The main focus is forecasting the volatility of the Dow Jones and some individual stocks (1993-2003). As a weighting scheme the Beta function is employed and the number of included lags is restricted to $K = 50$.[23] The results can be summarized as follows: on average MIDAS performs better than the single-frequency benchmark model. This result holds both in-sample and out-of-sample. Longer-horizon forecasts do not necessarily produce worse out-of-sample performances (at least up to four weeks). The inclusion of high-frequency data (5-minute returns) does not necessarily lead to better volatility forecasts.

Ghysels, Sinko, and Valkanov (2007) extend the previous study. The authors adjust the realized volatility for 5 and 30 minutes frequency regressors for microstructure noise with formulae suggested by Hansen and Lunde (2006). Using the same data set they predict realized volatility at weekly, two-week, three-week, and monthly horizons. For two individual stocks they find that the noise-corrected volatility measures perform, on average, worse than the unadjusted volatility measure. As a possible explanation, Ghysels, Sinko, and Valkanov (2007) speculate that the noise for the 5-minute data is negligible compared to the signal. Or, it could be that the MIDAS regressions are more efficient in extracting the signal from the unadjusted, daily, realized volatility measures compared with the noise-corrected schemes.

Forsberg and Ghysels (2007) extend this strand of literature and investigate why absolute return forecasts volatility (measured as quadratic variation) so well. Employing tick-by-tick data for the S&P 500 index from 1985 and 2003 they compare different regressors using MIDAS and HAC regressions. They extend the basic models by allowing for jumps in the volatility process. Forsberg and Ghysels (2007) use daily regressors (calculated using 5-minute returns) to forecast one-day, and one to four weeks ahead. For the MIDAS regression they allow for a maximum lag of $K = 50$ days.[24] As the weighting function they use the Beta function with declining weights. In an in- and out-of-sample comparison between the two models and different regressors,

[23] The authors set $\theta_1 = 1$ to ensure a decaying weighting function.

[24] The authors claim that the results are not sensitive to varying the included lags.

the MIDAS model is on average better (in adjusted R^2 and MSE terms) compared with the HAR model. For both models realized absolute returns has the highest predictive power.

Chen and Ghysels (2008) apply the semi-parametric MIDAS model to investigate whether finer data sampling schemes change the shape of the news-impact curve introduced by Engle and Ng (1993) based on daily data. For the parametric part, two Beta polynomial specifications are used to accommodate intra-daily and daily memory decay. The data set consists of five-minute returns of (respectively) Dow Jones and S&P 500 cash and futures markets from 1993(96) to 2003. In a out-of-sample exercise they show that accounting for asymmetry improves forecasting accuracy.

3.5.2 Macroeconomic Applications

Clements and Galvao (2005) are the first to apply MIDAS regressions to macroeconomic data. Furthermore they are the first to forecast MIDAS with an autoregressive term. In their final published paper, Clements and Galvao (2008) focus only on the MIDAS-AR model. They forecast quarterly US output growth with three monthly indicators (industrial production (IP), employment and capacity utilization (CU)).[25] The authors use real-time data which range from 1959-2005. They employ exponential Almon lag weighting functions with two parameters, restricted to ensure declining weights. In the paper it is not stated how many lags are included in the model.[26] They estimate the model with NLS. In a recursive forecasting exercise they compare the MIDAS-AR model with a quarterly AR(1), bridge equation and the mixed-frequency distributed lag model by Koenig, Dolmas, and Piger (2003). Furthermore Clements and Galvao (2008) implement within-quarter forecasts. They find that the use of monthly indicators (IP and CU) in the MIDAS regression, especially for short horizons (within-quarter), results in sizeable reductions in RMSE compared with single-frequency models. Com-

[25] In the working paper version Clements and Galvao (2005) also forecast US inflation.
[26] In personal correspondence, Michael Clements told me that they included 24 months (8 quarters).

paring MIDAS-AR with mixed-frequency distributed lag models and bridge equations there is little to choose.

Ghysels and Wright (2008) replicate the forecasts of others employing both the MIDAS and the Kalman filter approaches. As survey forecasts are infrequently published and often found to be stale, it would be interesting for policy makers to predict the upcoming survey releases at a higher frequency. Using survey data from the Survey of Professional Forecasters (SPF) and the Consensus Forecast (CF), the forecasts are generated via daily data (excess returns and the yield curve among others) for real GDP growth, CPI inflation, T-Bill and the unemployment rate. Using MIDAS (Beta weighing function either estimated, unrestricted, or assuming equal weights) the upcoming release is forecast, whereas in the state-space approach the forecast for a specific horizon at a specific day is interpolated by applying the Kalman filter, viewing these as missing data. In an in- and out-of-sample forecasting exercise, on average both approaches beat the simple random walk benchmark forecasts.

Marcellino and Schumacher (2007) were the first to combine factor models with the MIDAS approach. They apply the standard two-step procedure: first estimate the factors and then plug these into a specific time series model. A second focus of the paper is the ragged-edge data problem, that observations are not available for all time series at the end of the sample. The authors compare the realignment approach of Altissimo, Cristadoro, Forni, Lippi, and Veronese (2006), the EM algorithm outlined in chapter 2, and a parametric state-space factor estimator of Doz, Giannone, and Reichlin (2006). In addition to the standard MIDAS model[27] they employ the 'smoothed' MIDAS model and an unrestricted version of the MIDAS approach. Irrespective of the factor estimation approach and the employed MIDAS model, factors do provide valuable information for short-term forecasting (nowcasting). The results are interpreted relative to GDP variance. For longer forecasting horizons, the results are more ambiguous. These find-

[27] The authors use the following restrictions for the exponential Almon lag weighting function: $\theta_1 < 2/5$ and $\theta_2 < 0$, which generates either declining weights or a hump-shaped weighting function.

ings even hold when compared with quarterly single-frequency factor models, an AR(p), and a no-change benchmark. Including an autoregressive lag in the MIDAS approach does not improve the results. In general the unrestricted MIDAS model performed best in many cases.

Hogrefe (2008) compares single-frequency, mixed-frequency and interpolation approaches to forecast data revisions for US GNP. Building upon a real-time data set consisting of 13 quarterly and 30 monthly series, the author employs the quarterly-only single-frequency model, quarterly plus aggregated monthly data, the approach by Chow and Lin (1971), and the MIDAS approach. First, Hogrefe (2008) confirms a strong rejection of the hypotheses of data rationality found in similar single-frequency studies.[28] with the latter two approaches. Both models perform better in an out-of-sample forecasting exercise (first, second and last revision), where the MIDAS is the best one. The authors include 12 months as lags and the exponential Almon lag function is restricted to ensure decaying weights.[29]

Ghysels and Valkanov (2006) systematically investigate the forecasting performance of the MIDAS model compared with single-frequency approaches. The authors simulate a bivariate infeasible high-frequency VAR(1) with different persistence and leading properties. The mixed-frequency sample is obtained by skipping the corresponding observations. They conduct an in-sample forecasting comparison of the MIDAS model (both with Almon and Beta weighting function), the infeasible high-frequency VAR (HF-VAR), and a low-frequency VAR. By definition, the infeasible HF-VAR is not outperformed by any model, but the MIDAS model performs relatively well. The RMSEs are only between 2 and 9 percent larger than the high-frequency VAR. Compared with the low-frequency VAR, the RMSE of the MIDAS model are on average between 17 and 40 per cent better. Ghysels and Valkanov (2006) find only small differences between the exponential Almon lag and the Beta weighting function. But the Almon lag performs constantly better across all models. The results hold for a variety of frequency mixtures

[28] This term was introduced by Mankiw and Shapiro (1986). Data rationality implies that there is no possibility for the prediction of data revisions.

[29] But this is not stated in the paper.

($m = \{5, 10, 20, 60, 120, 250\}$). However, the authors do not state how many lags are included and whether or not the weighting functions are restricted.

Ghysels and Valkanov (2006) also augment the information set by including information that becomes available in the m periods between t and $t+1$. As a further benchmark the low-frequency VAR is adjusted by "filling in" observations that might not be available between t and $t+1$. The authors employ the Kalman filter to estimate the missing observations. The approach builds upon the idea by Harvey and Pierse (1984). As *a priori* expected, the RMSEs decline further after adjusting the information set. The Kalman Filter approach works considerably good but is not able to outperform the MIDAS models.

Tab. 3.1: Characteristics of empirical MIDAS applications

	Application	Frequency mixture	Weighting Function	Restrictions	Included lags
Ghysels, Santa-Clara, and Valkanov (2005)	risk-return trade-off	daily-monthly	Almon	yes	252 days
Ghysels, Santa-Clara, and Valkanov (2006)	volatility forecasts	5-minutes - monthly	Beta	yes	50 days
Ghysels and Valkanov (2006)	Monte Carlo study	varies	Almon / Beta		
Ghysels, Sinko, and Valkanov (2007) a	risk-return trade of	daily-monthly	Almon / Beta	no	252 days
Ghysels, Sinko, and Valkanov (2007) b	volatility forecasts correcting for noise	5/30 minutes - monthly	Almon / Beta	no	varies
León, Nave, and Rubio (2007)	risk-return trade-off European equity markets	daily-monthly	Almon	no	248-253 days
Forsberg and Ghysels (2007)	volatility forecasts	daily-monthly	Beta	yes	50 days
Marcellino and Schumacher (2007)	now and forecasting German GDP	monthly-quarterly	Almon	yes	12 months
Kong, Liu, and Wang (2008)	risk-return trade-off Chinese stock markets	daily-daily	Beta	yes	88 days
Ghysels and Wright (2008)	forecasting of forecasts	daily-quarterly (monthly)	Beta	yes/no	varies
Clements and Galvao (2008)	forecasting US GDP	monthly-quarterly	Almon	yes	24 months
Hogrefe (2008)	forecasting US GDP revisions	monthly-quarterly	Almon	yes	12 months

3.6 Summary and Discussion

In this chapter we reviewed the literature on forecasting with mixed-frequency data. Currently, the standard approach is still to apply a two-step procedure. First, ensure that all data are sampled at the same frequency and then apply the time series models. Temporal aggregation plays the dominant role compared with the interpolation method. More recently state-space factor approaches where forecasts are generated at the higher-frequency have been applied more often.

In general we can state: mixed-frequency data matter, that is transformation of data leads, on average, to less accurate forecasts. Improvements were found, especially in the short-run.

The first approach to deal with mixed-frequency data was the linkage approach, where forecasts from different frequencies are combined in a formal way. The majority of applications generate forecasts from quarterly structural models which are combined with monthly time series forecasts. The combination also improved forecasts for variables that are only available at the lower frequency. Although the results were promising it seems that currently there is no more research in this area.[30] This is rather surprising, as currently forecasting combinations in general play a prominent role in the forecasting literature (see Timmermann (2006) for a recent survey). A possible reason is that large structural models play a less central role in the current forecasting literature.

By contrast, bridge equations are still widely used, especially at central banks. The survey demonstrated that bridge equations are a useful tool for forecasting, especially in the short run (nowcasting). As more information becomes available, the more accurate are the forecasts of the current quarter. Most of our reviewed articles found decreasing RMSEs within the quarter. As more information becomes available, the more accurate are the forecasts for the target variable. There is no predominant view on how to forecast the higher-frequency variables. Due to low computational cost, re-

[30] The last article on this topic we found was the one by Shen (1996).

cently more than one time series model is used (mostly AR, VAR, or BVAR). One central problem is that errors from the auxiliary model transmit to the bridge equation.

Pure mixed-frequency time series models have come into focus in recent years. There are two completely different approaches to handle mixed-frequency data. First we have state-space approaches, both in connection with dynamic factors as well as a pure VAR. State-space approaches are quite general and can handle any kind of data issues (mixed-frequency, reporting lags, measurement error, data revisions). This generality comes at a cost: with more data and higher-frequency mixtures the computational complexity rises dramatically. Therefore, only small-scale models with few variables are estimated. This is one reason why factor models are currently so dominant in the forecasting literature. Nevertheless the forecasting accuracy is better when compared with single-frequency models.

The second mixed-frequency approach is distributed lag models, most prominently the MIDAS approach which is advocated and promoted by Eric Ghysels and co-authors.[31] MIDAS models are parsimonious but can handle only equidistant data. MIDAS models can easily be estimated via NLS. Which weighting function should be employed remains undecided. Both are used and seem to produce similar results. Additionally, the weighting function is often restricted to ensure declining weights. But these restrictions are not based on theoretical reasoning. There are no comparisons between an unrestricted and a restricted version so far. Most applications of MIDAS can be found in the finance literature with very promising forecasting results. The three macroeconomic applications state that MIDAS is useful for short-term forecasting.

[31] The basic MATLAB code can be downloaded at www.unc.edu/ sinko/midas.zip (assessed August 2008).

4. EVALUATION CONSIDERATIONS

4.1 Preliminaries

In the following chapters we want to compare the mixed-frequency VAR model and the MIDAS approach in more detail. We have already noted, that these approaches have not been compared so far. We start by investigating model specification issues. Then we compare the forecast performance of the two techniques in a Monte Carlo study. In chapter 7 we extend the analysis to real data. In this chapter we outline the data generating processes (DGP) that will be used in the next chapters for the model comparisons.

To draw general conclusions from Monte Carlo studies the data generation process should cover as many data structures as possible. Different data structures are defined by different autocorrelation structures of the lower-frequency variable. As the mixed-frequency time series models are quite different in their approach to handling mixed-frequency data, it is not obvious from the theoretical point of view which model yields a more accurate forecasting performance, given a specific data pattern and (or) loss function.

In macroeconomic forecasting GDP is one of the variables forecast most often, as it considered as the most general representation of an economy. Figure 4.1 displays the autocorrelation structure of US (1954-2006) and German GDP (1991-2006) growth (quarterly and yearly). For the yearly growth rates one can see an oscillating autocorrelation structure for both countries. The oscillating pattern, with significant lags, can also be detected for US quarterly GDP growth. For German quarterly GDP growth we find no significant autocorrelations at any lag. We want to replicate these kinds of patterns and want to add further structures in our Monte Carlo study.

We define four different processes (labelled as Process I, II, III and IV respectively). With each process we generate mixed-frequency data. Within the different processes we can vary the relationship between the time series and the structure of the target (low-frequency) variable. All simulated processes are stationary.[1] As we focus on forecasting we will always assume that the higher-frequency variable can be regarded as an indicator for the lower-frequency variable. This is also the standard case in the literature.

In addition to the autocorrelation patterns of the lower-frequency variable we also define different 'strengths' of leading characteristics. We distinguish between cases with absolutely no relationships and cases with medium and strong predictive power between higher- and lower-frequency variables.

Processes I and II are generated from a bivariate VAR model and Process IV from a trivariate high-frequency VAR model. In Process IV we generate two predictors sampled at different frequencies. These processes assume that all variables are generated at the highest frequency. The mixed-frequency VAR model also assumes that all variables of the model are generated at the highest frequency. In later chapters we will see whether this is an advantage in comparison to the MIDAS model. In contrast, in Process III the data will be generated with a MIDAS model.

The higher frequency variable is sampled m-times higher than the lower frequency. In the generated high-frequency sample we observe for the lower frequency only every m-th observation; the other ones are skipped. To give an example, given $T = 100$ (length of the lower- frequency time series) and $m = 3$ (for instance monthly-quarterly data) we generate a bivariate data set with length $mT = 300$. The principle of skip-sampling is illustrated for monthly-quarterly ($m = 3$) in Table 4.1 . We simulate two time series (x_1 and x_2) with a VAR model, where x_2 is the lower frequency. We keep every third observation and the others are skipped.

As outlined in chapter 2, we focus on plausible macroeconomic mixtures, ad-

[1] This assumption is necessary as, first, the state-space framework of Zadrozny assumes stationary data. And second, the use of non-stationary data within the MIDAS framework is an unsolved issue.

ditionally to $m = 3$ we employ quarterly-yearly (or weekly-monthly) data ($m = 4$) and weekly-quarterly data ($m = 12$). This is in contrast to Ghysels and Valkanov (2006) who consider the following mixtures: $m = \{5, 10, 20, 60, 120, 250\}$.

For ease of computation all generated data are standardized.

Tab. 4.1: Example of skip-sampling for $m = 3$

	simulated data		skip-sampled data	
t	x_1	x_2	x_1	x_2
1	0.320	0.424	0.320	
2	0.056	0.164	0.056	
3	0.604	0.684	0.604	0.684
4	0.209	0.411	0.209	
5	0.915	0.672	0.915	
6	0.392	0.749	0.392	0.749

Fig. 4.1: Autocorrelations for US and German GDP growth rates

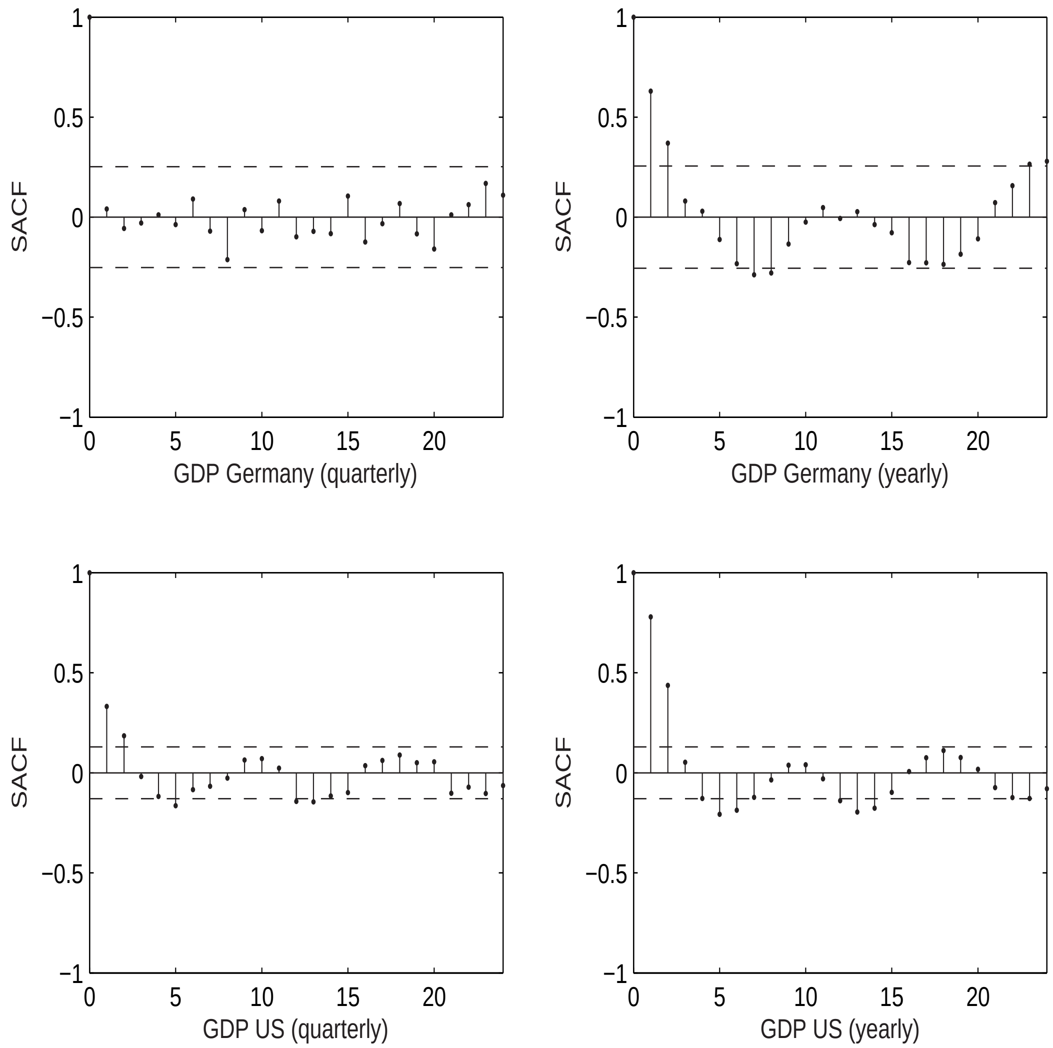

4.2 Process I

For the first process we follow Ghysels and Valkanov (2006). Use of this process allows us to compare our results with those of Ghysels and Valkanov (2006). We generate high-frequency data via a bivariate VAR(p) with $p = 1$.

$$X_\tau = \Gamma X_{\tau-1} + u_\tau \tag{4.1}$$

The matrix Γ is specified as:

$$\Gamma = \rho \times \begin{bmatrix} 1 & \delta_l \\ \delta_h & 1 \end{bmatrix}$$

such that a single parameter ρ determines the persistence of both series, whereas δ_l and δ_h capture the dependence between the two series. Generally, we will set $\delta_l = 0$. The bivariate random vector u_t is drawn from $N(0, I)$, where 0 is a bivariate zero-vector, and I is the identity matrix of dimension two. The data are generated for $\rho = \{0.10, 0.50, 0.90, 0.95\}$ and $\delta_h = \{0, -0.5, -1.5, -3.5\}$. The parameter ρ determines the persistence of time series and δ_h determines the 'predictive power' of the higher-frequency time series. The series are simulated for $m \times T$ observations, where $m = \{3, 4, 12\}$.

Tab. 4.2: Eigenvalues for Process I

		λ_1	λ_2
$\rho = 0.1$	$\delta_h = i$	0.10	0.10
$\rho = 0.5$	$\delta_h = i$	0.50	0.50
$\rho = 0.9$	$\delta_h = i$	0.90	0.90
$\rho = 0.95$	$\delta_h = i$	0.95	0.95

Notes: Table reports the eigenvalues of the coefficient matrix Γ in Equation (4.1). The parameter δ_h takes the values: $\delta_h = \{0, -0.5, -1.5, -3.5\}$.

In Table 4.2 we tabulate the eigenvalues of the coefficient matrix Γ for the simulated high-frequency process. All eigenvalues are smaller than 1 indicating stationary processes.

Figures 4.3 to 4.5 in the appendix to this chapter display the different autocorrelation structures (up to 24 lags) for $m = \{3, 4, 12\}$ and the different values of ρ and δ_h. The time series length is $T = 100$. Each graph in each figure plots the autocorrelation of the lower frequency variable. The dashed lines in the plots of the autocorrelations are the approximate two standard error bounds computed as $\pm 2/\sqrt{(}T)$. If the autocorrelation is within these bounds, it is not significantly different from zero at (approximately) the 5% significance level. Each row corresponds to a value of ρ with the corresponding values of δ_h. Note that the autocorrelation structure for a specific simulated *high-frequency* data set is almost the same as for the different frequency mixtures.[2] The different structures are due to the skipping of the unobserved high-frequency observations to obtain the low-frequency data. Therefore the resulting processes exhibit different dynamics between the variables.

First, we can report that at the first lag the correlation is positive for almost all examples. Furthermore, for low persistent time series ($\rho = 0.1$ and $\rho = 0.5$) there are almost no significant autocorrelations at any lags (independently of the parameter values of δ_h). The picture is different for highly persistent time series (lower two rows in each graph). One can see slow or fast decay and oscillating autocorrelation patterns. Comparing three graphs one can see that the autocorrelations at each lag 'die out' the higher the frequency mixture is, that is the structure is shifted from higher to lower lags.[3] For instance, take the last picture in each graph ($\rho = 0.95$ and $\delta_h = -3.5$). The number of positive significant figures gets lower as more observations are skipped, the higher is m.

4.3 Process II

The second process is an extension of the first process. We add an additional lag to the system. We consider the following VAR(2) data generating process

[2] However, the errors and length of the time series are different.

[3] This can be confirmed by looking at autocorrelations at even higher lags up to 72.

$$X_\tau = \Gamma_1 X_{\tau-1} + \Gamma_2 X_{\tau-2} + u_\tau \tag{4.2}$$

The matrix Γ_1 is specified as:

$$\Gamma_1 = \rho \times \begin{bmatrix} 1 & 0 \\ \delta_h & 1 \end{bmatrix}$$

where $\rho = \{0.10, 0.50, 0.90\}$ and $\delta_h = \{0, -0.5, -1.5\}$.[4] The matrix Γ_2 is specified as:

$$\Gamma_2 = \begin{bmatrix} -0.1 & 0.1 \\ -0.15 & -0.2 \end{bmatrix}$$

In Table 4.3 we present the absolute eigenvalues from Equation (4.2).[5] Again, we generate stationary time series as all eigenvalues are less than 1. Figures 4.6 to 4.8 in the appendix plot the corresponding sample autocorrelation functions. Compared with Process I, we have some similarities and differences. For low persistent time series we find again almost no significant autocorrelation at any lag. In contrast to Process I, the first lag is in some cases negative. For highly persistent series we can detect an oscillating pattern but there is often a change between positive and negative values after two consecutive lags. The figures demonstrate that Process II generates a completely different data structure in comparison to Process I.

4.4 Error Structure With GARCH Components

In addition to the homoscedastic errors used in Equation (4.1) and (4.2) we want to allow for heteroscedastic errors in the higher-frequency variable in Process I and II. It is well known in applied econometrics that both high- and low-frequency data display volatility clustering (see Bera and Higgins (1993)). The higher the frequency, the more likely is the aspect of time-

[4] We do not consider the values $\rho = 0.95$ and $\delta_h = -3.5$, as the process would not be stable.

[5] All eigenvalues are complex in this example.

Tab. 4.3: Eigenvalues for Process II

		λ_1	λ_2	λ_3	λ_4
$\rho = 0.1$	$\delta_h = 0$	0.417	0.417	0.448	0.448
	$\delta_h = -0.5$	0.459	0.459	0.406	0.406
	$\delta_h = -1.5$	0.481	0.481	0.388	0.388
$\rho = 0.5$	$\delta_h = 0$	0.349	0.349	0.535	0.535
	$\delta_h = -0.5$	0.598	0.598	0.312	0.312
	$\delta_h = -1.5$	0.688	0.688	0.271	0.271
$\rho = 0.9$	$\delta_h = 0$	0.770	0.770	0.242	0.242
	$\delta_h = -0.5$	0.860	0.860	0.217	0.217
	$\delta_h = -1.5$	0.971	0.971	0.192	0.192

Notes: Table reports the absolute eigenvalues of the coefficient matrices Γ_1 and Γ_2 from equation (4.2).

varying volatility. Consider again the data generation process for Process I

$$X_\tau = \Gamma X_{\tau-1} + u_\tau$$

and Process II

$$X_\tau = \Gamma_1 X_{\tau-1} + \Gamma_2 X_{\tau-2} + u_\tau$$

Let $u_{1\tau}$ be the higher-frequency variable distributed as $u_{1\tau}|\Omega_{\tau-1} \sim N(0, \sigma^2_{1\tau})$, where $\Omega_{\tau-1}$ denotes the history of the process. We use the standard GARCH(1,1) model to simulate the error variance

$$\sigma^2_{1\tau} = \alpha_0 + \alpha_1 u^2_{\tau 1} + \beta_1 \sigma^2_{1\tau-1}. \tag{4.3}$$

For $m = 3$ and $m = 4$ we set $\alpha_0 = 0.01$, $\alpha_1 = 0.13$, and $\beta_1 = 0.82$.[6] This set-up displays moderate shocks that persist over time. In contrast, for $m = 12$ (weekly-quarterly data) we set $\alpha_1 = 0.4$ and $\beta_1 = 0.5$ to model larger shocks which are less persistent.

[6] Our choice is inspired by Bollerslev (1986). These parameters were estimated for quarterly UK inflation.

4.5 Process III

The first two processes are generated with a bivariate VAR. It is expected that on average the mixed-frequency VAR performs better than the MIDAS approach (though, not theoretically derived), as the former is modelled at the same frequency as the data generating process. Therefore we generate the data for the third process with a MIDAS model. The data-generating process was suggested by Andreou, Ghysels, and Kourtellos (2007). We consider the following data generation process

$$y_t = \beta_0 + \beta_1 B\left(L^{1/m};\theta\right) x_{t-1}^{(m)} + \epsilon_t^{(m)}, \tag{4.4}$$

where the constant and the slope coefficients are given by: $\beta_0 = 0$ and $\beta_1 = \{0.6, 3, 6\}$. The error term is normally distributed with expected mean zero and variance 0.125. The different values of β_1 yield models of small, medium and large Signal to Noise Ratios (SNR), respectively. The dependent variable is measured at low frequency whereas the covariates are defined at high sampling frequency m, based on a K-dimensional vector of weights. The weighting function we parameterize as the exponential Almon lag with the unknown coefficients $\theta = (\theta_1, \theta_2)$ given by (2.56). We want to generate two different weighting schemes. We choose θ that yield both fast and slow decay of the weights. Thus we have $\theta = (7 \times 10^{-4}, -5 \times 10^{-2})$ and $\theta = (7 \times 10^{-4}, -6 \times 10^{-3})$ for the fast and slow decay of weights respectively. Figure 4.2 plots the two weighting functions with $K = 40$ lags.

For the data generation process (4.4) we assume for the high-frequency covariates $x_t^{(m)}$

$$x_{t/m}^{(m)} = c_0 + c_1 x_{t-1/m}^{(m)} + u_{t/m}. \tag{4.5}$$

For the parameters we assume $c_0 = 0$, $c_1 = 0.9$ and $u_{t/m} \sim N(0, 1)$. Figures 4.9 to 4.11 plot the corresponding sample autocorrelation patterns for Process III. The dominant structure is the oscillating pattern that we also observe in real data examples, as plotted in Figure 4.1. For $m = 3$ and $m = 4$, at the first three lags, all autocorrelations are positive and significant. For

$m = 12$ only the first and some higher lags are significant. Furthermore the oscillating pattern is less well pronounced.

Fig. 4.2: Shape of the weighting functions for Process III

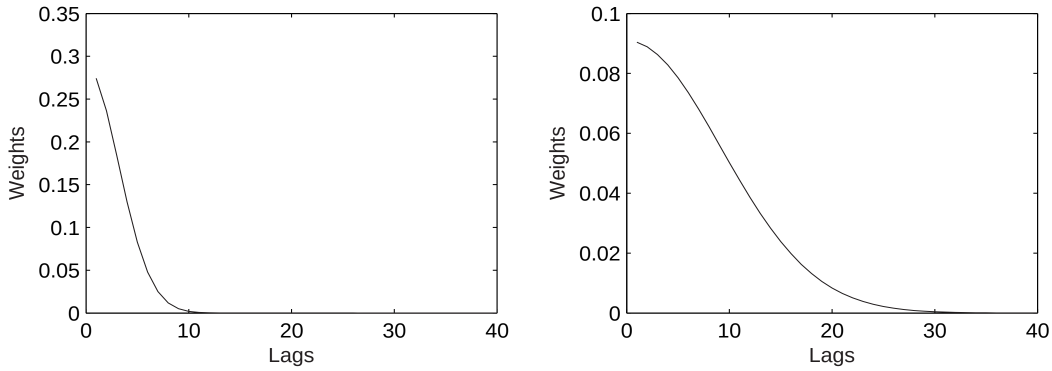

4.6 Process IV: Three Different Frequencies in One Model

The simulated processes so far were bivariate processes. The two time series approaches in focus are able to handle any number of variables sampled at any frequency. In empirical applications it is a natural step to consider indicators sampled at different frequencies, that is to include weekly indicators in addition to monthly ones. This aspect has not been investigated in the literature so far. Therefore, we extend our Monte Carlo study to three variables all sampled at different frequencies. Process IV can be considered as a starting point for future research on how many different frequencies can be handled in practice in one model.

We consider again the VAR(1) process

$$X_\tau = \Gamma X_{\tau-1} + u_\tau \tag{4.6}$$

The coefficient matrix Γ is now specified as

$$\Gamma = \rho \times \begin{bmatrix} 1 & 0 & 0 \\ -0.2 & 1 & 0.2 \\ -0.1 & -0.15 & 1 \end{bmatrix}$$

The parameter $\rho = \{0.10, 0.50, 0.90, 0.95\}$ again determines the persistence of the time series. The target variable is sampled at quarterly intervals simulated with homoscedastic errors. The second variable is sampled at monthly intervals with GARCH errors $\alpha_1 = 0.13$ and $\beta_1 = 0.82$ in Equation (4.3). Finally, the third variable is sampled weekly with GARCH error specification $\alpha_1 = 0.4$ and $\beta_1 = 0.5$ in Equation (4.3). In Table 4.6 the absolute eigenvalues are reported, which are again less than 1. The simulated time series are stationary as required for the mixed-frequency time series models. Figure 4.12 plots the autocorrelation function. Strong significant lags can only be detected for $\rho = 0.95$.

Tab. 4.4: Eigenvalues for Process IV

	λ_1	λ_2	λ_3
$\rho = 0.1$	0.101	0.101	0.100
$\rho = 0.5$	0.507	0.507	0.500
$\rho = 0.9$	0.913	0.913	0.900
$\rho = 0.95$	0.964	0.964	0.950

Notes: Table reports the absolute eigenvalues of the coefficient matrices Γ from Equation (4.6).

4.7 Appendix: Additional Figures

Fig. 4.3: Autocorrelations for Process I for $m = 3$

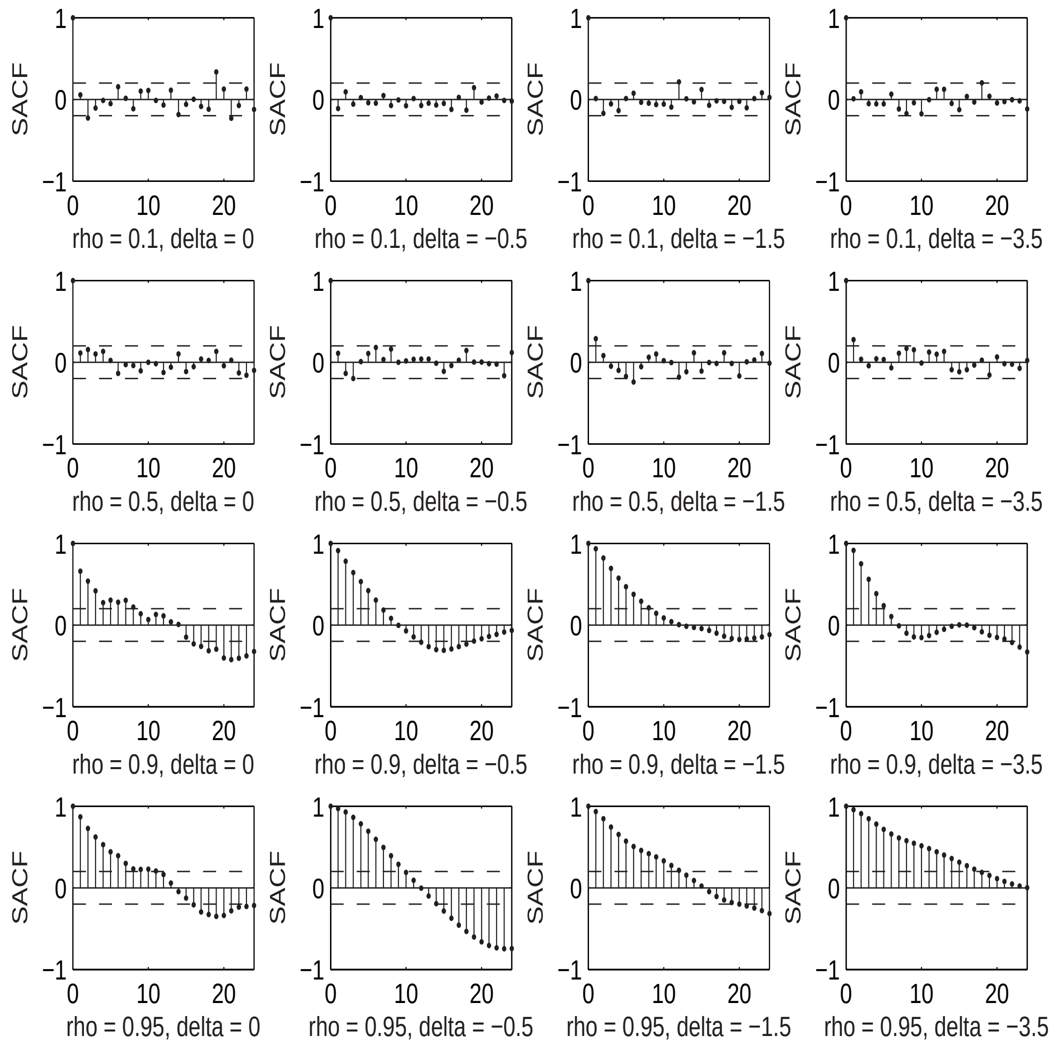

Fig. 4.4: Autocorrelations for Process I for $m = 4$

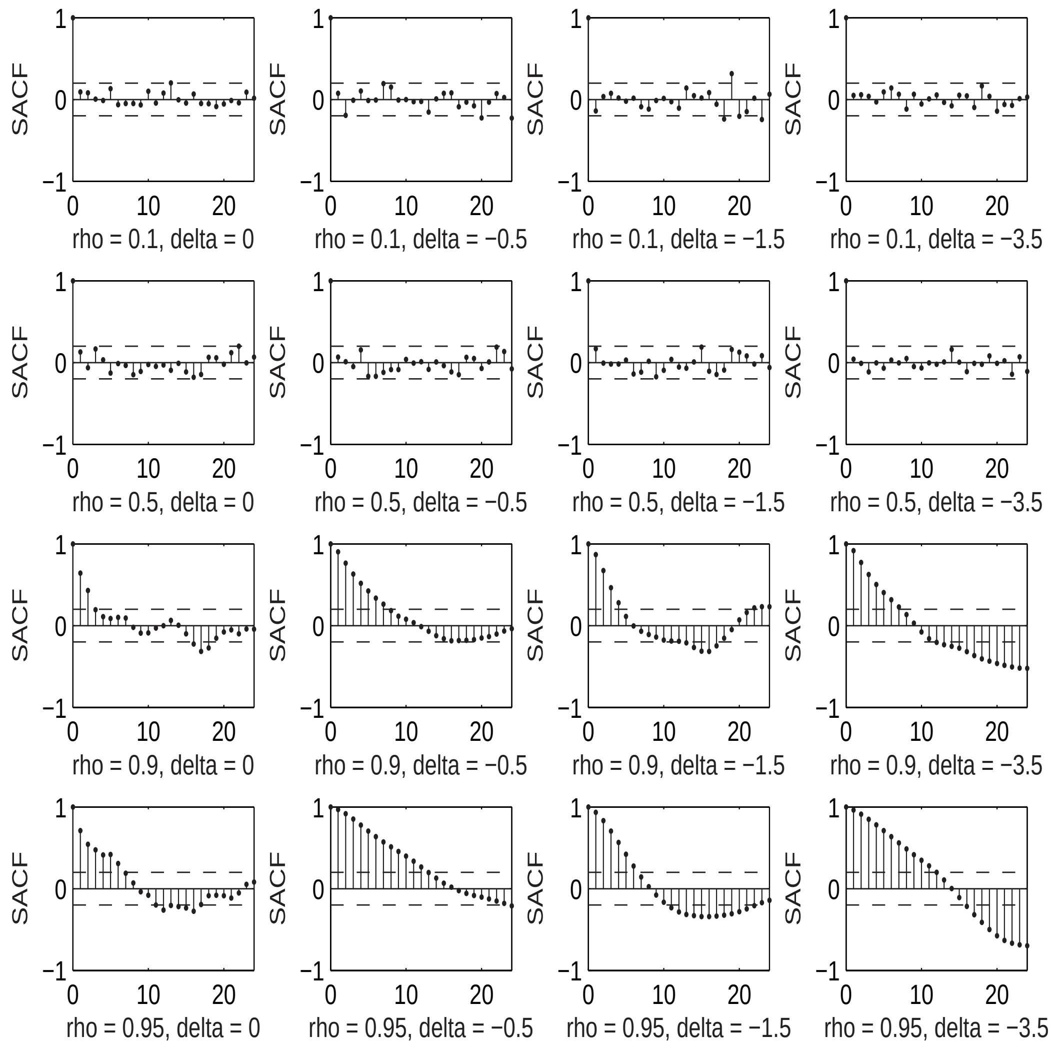

Fig. 4.5: Autocorrelations for Process I for $m = 12$

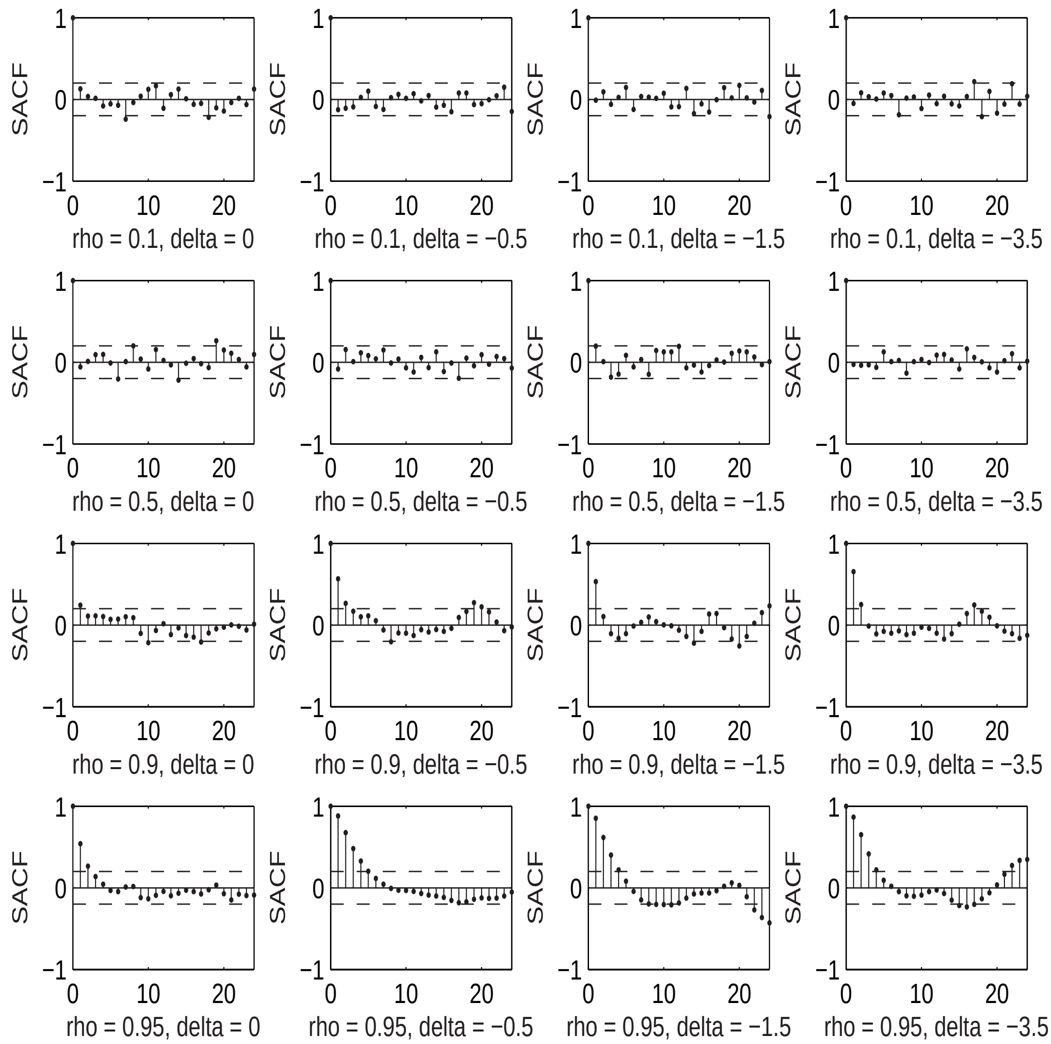

Fig. 4.6: Autocorrelations for Process II for $m = 3$

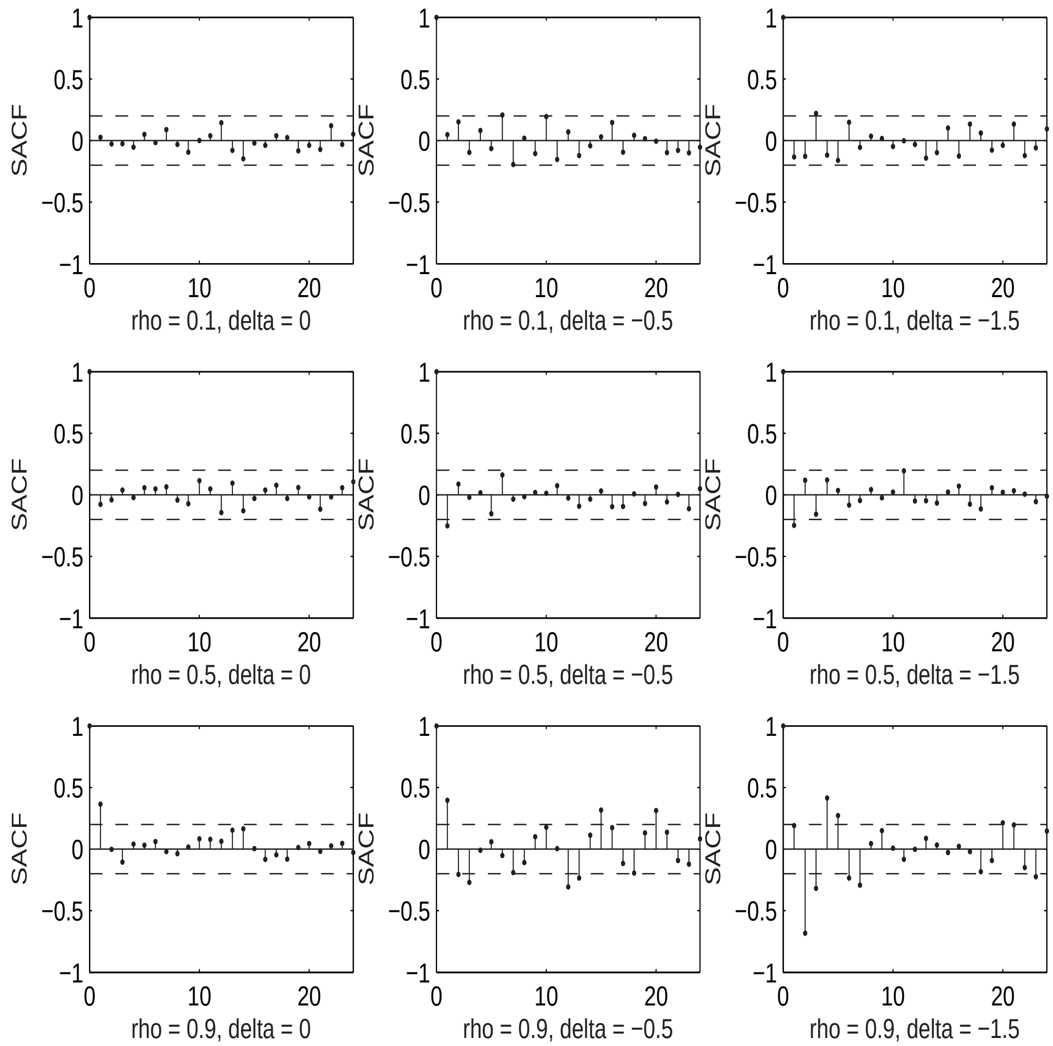

Fig. 4.7: Autocorrelations for Process II for $m = 4$

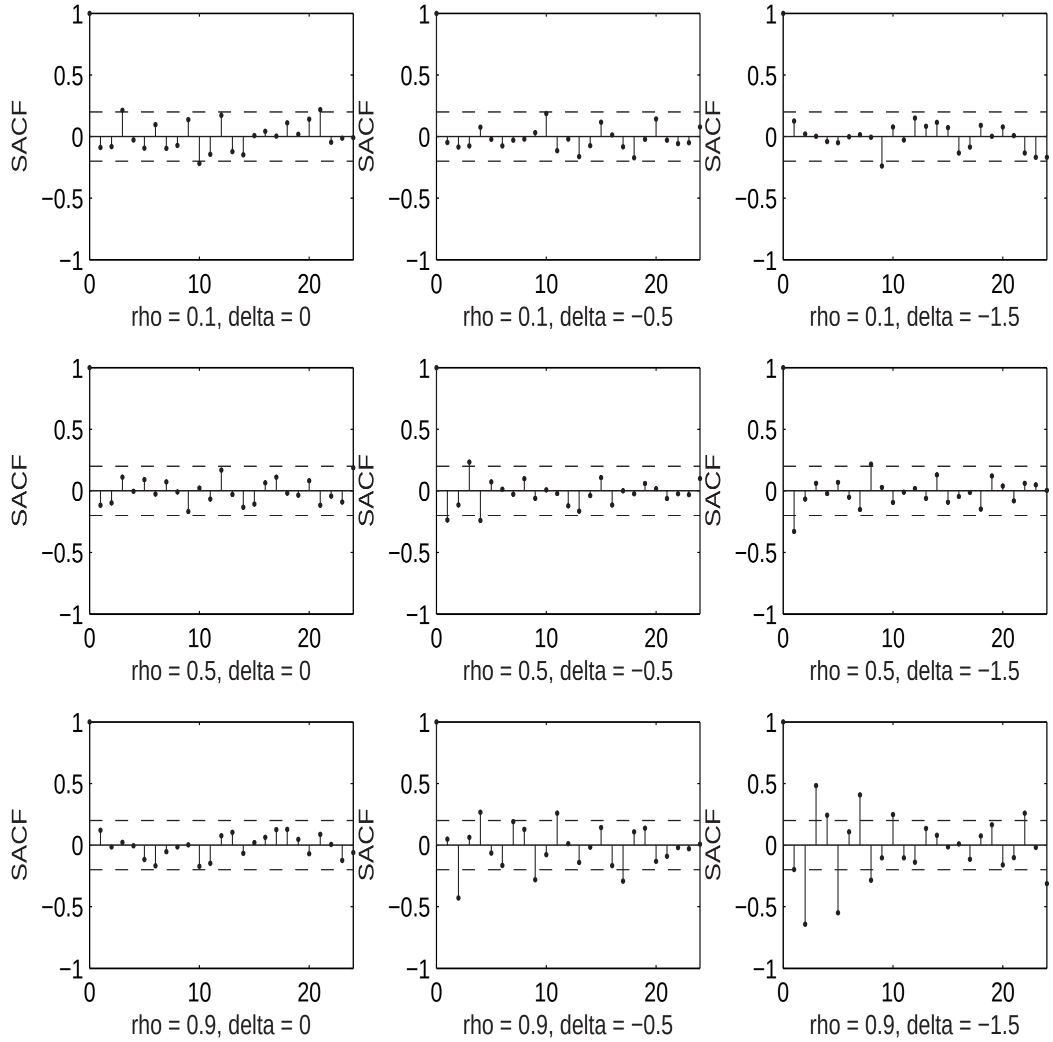

Fig. 4.8: Autocorrelations for Process II for $m = 12$

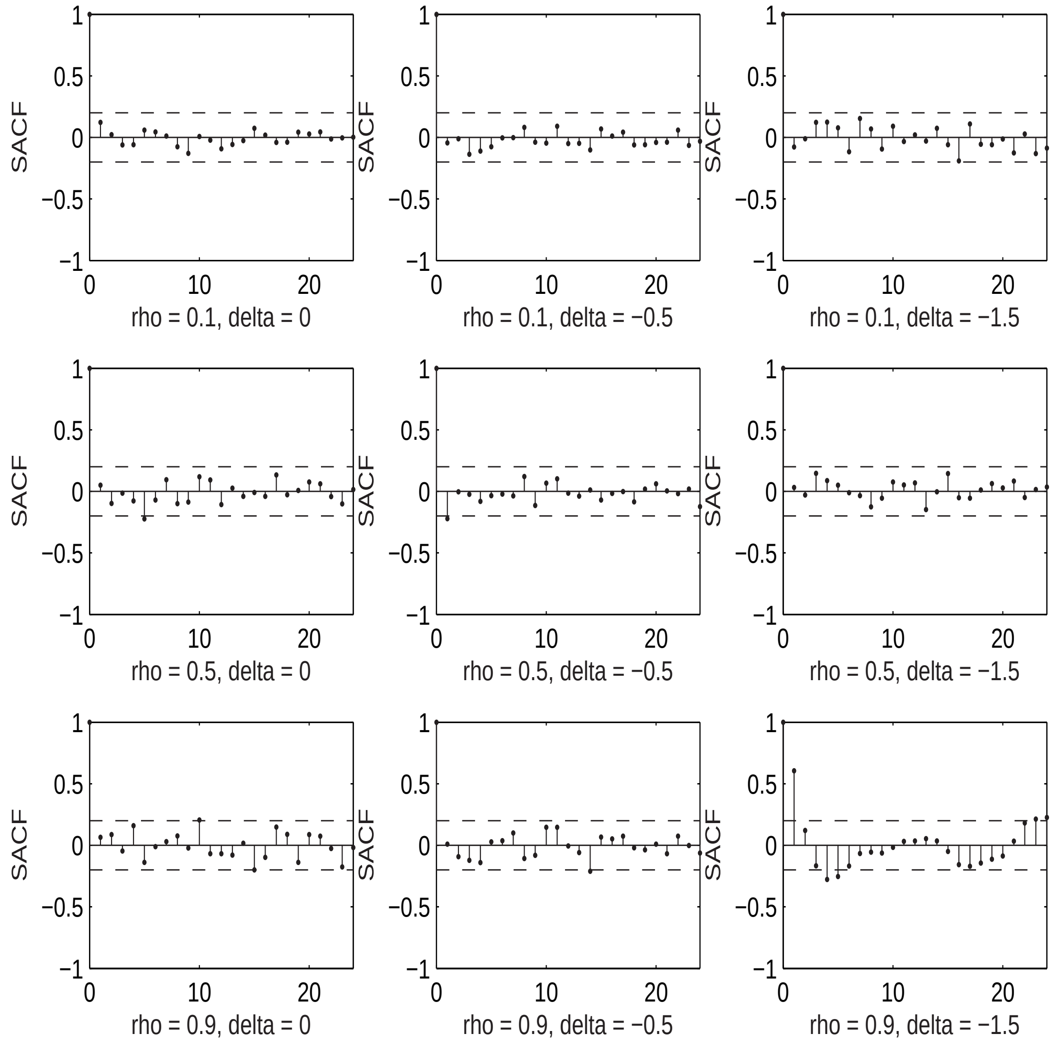

Fig. 4.9: Autocorrelations for Process III for $m = 3$

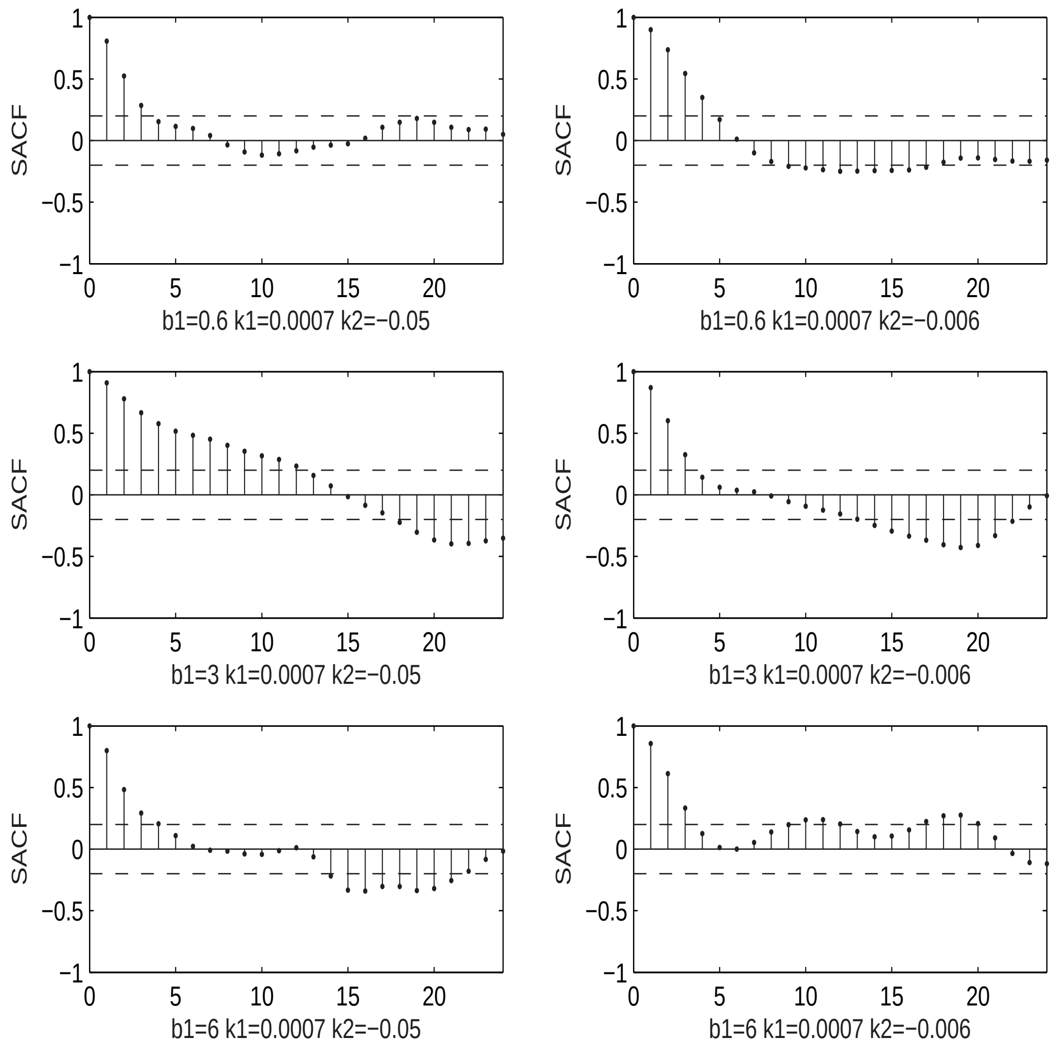

Fig. 4.10: Autocorrelations for Process III for $m = 4$

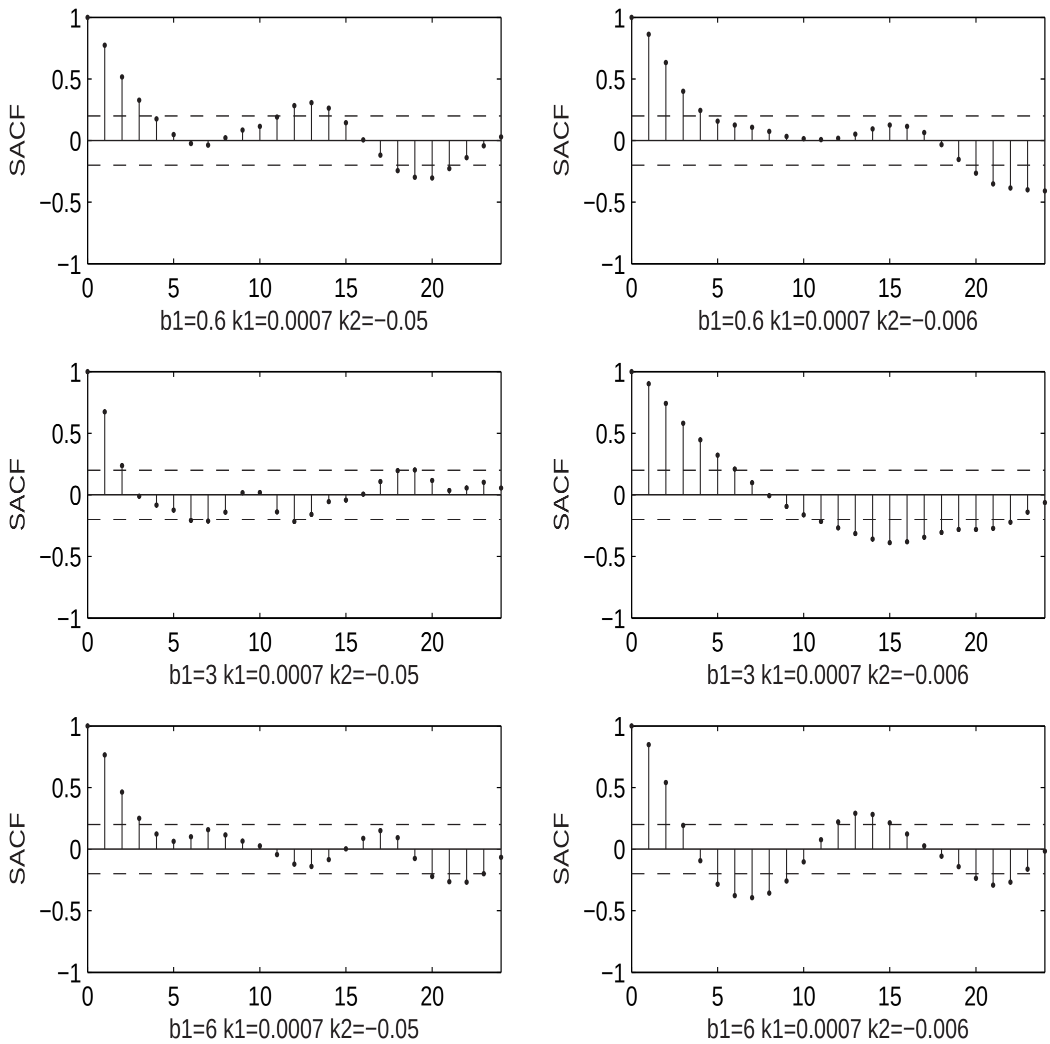

Fig. 4.11: Autocorrelations for Process III for $m = 12$

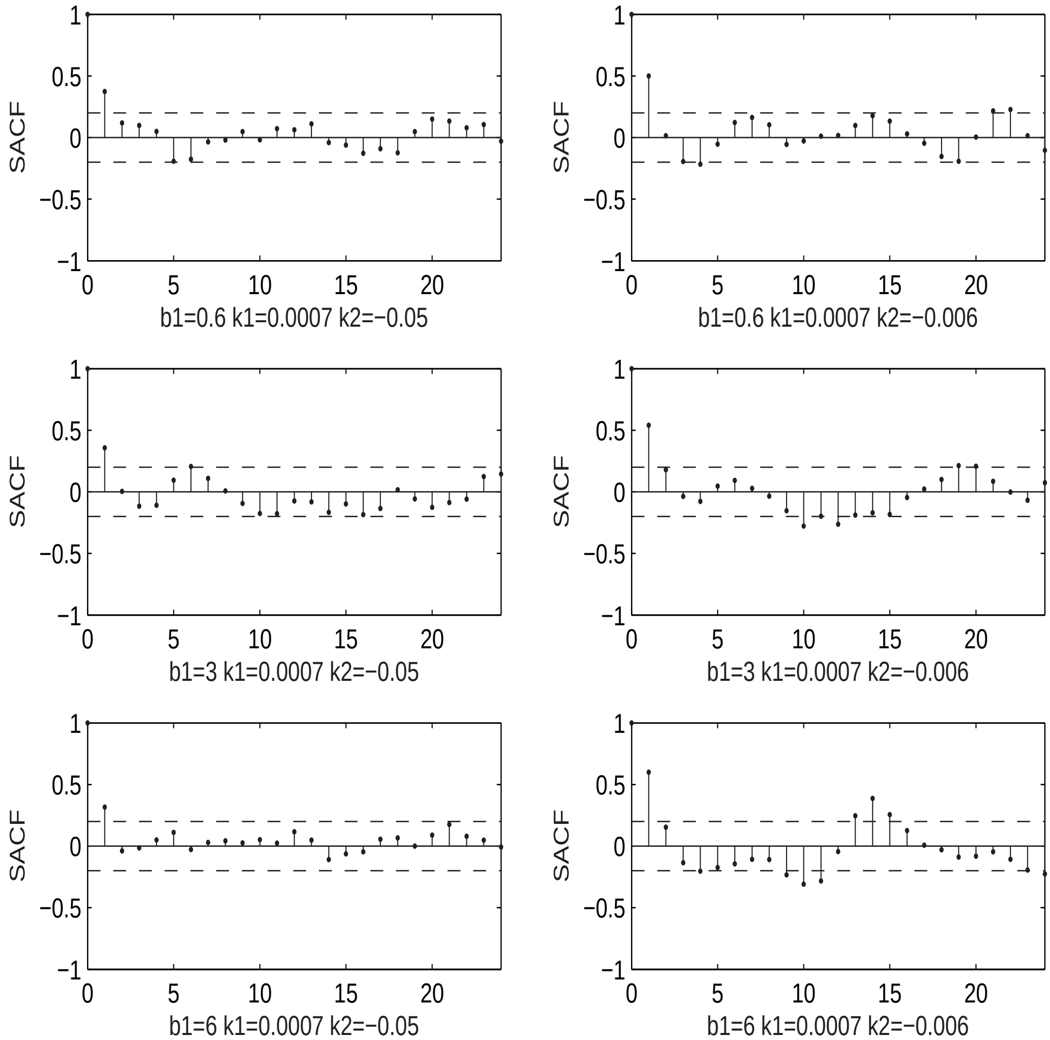

Fig. 4.12: Autocorrelations for Process IV for $m = 12$ and $m = 3$

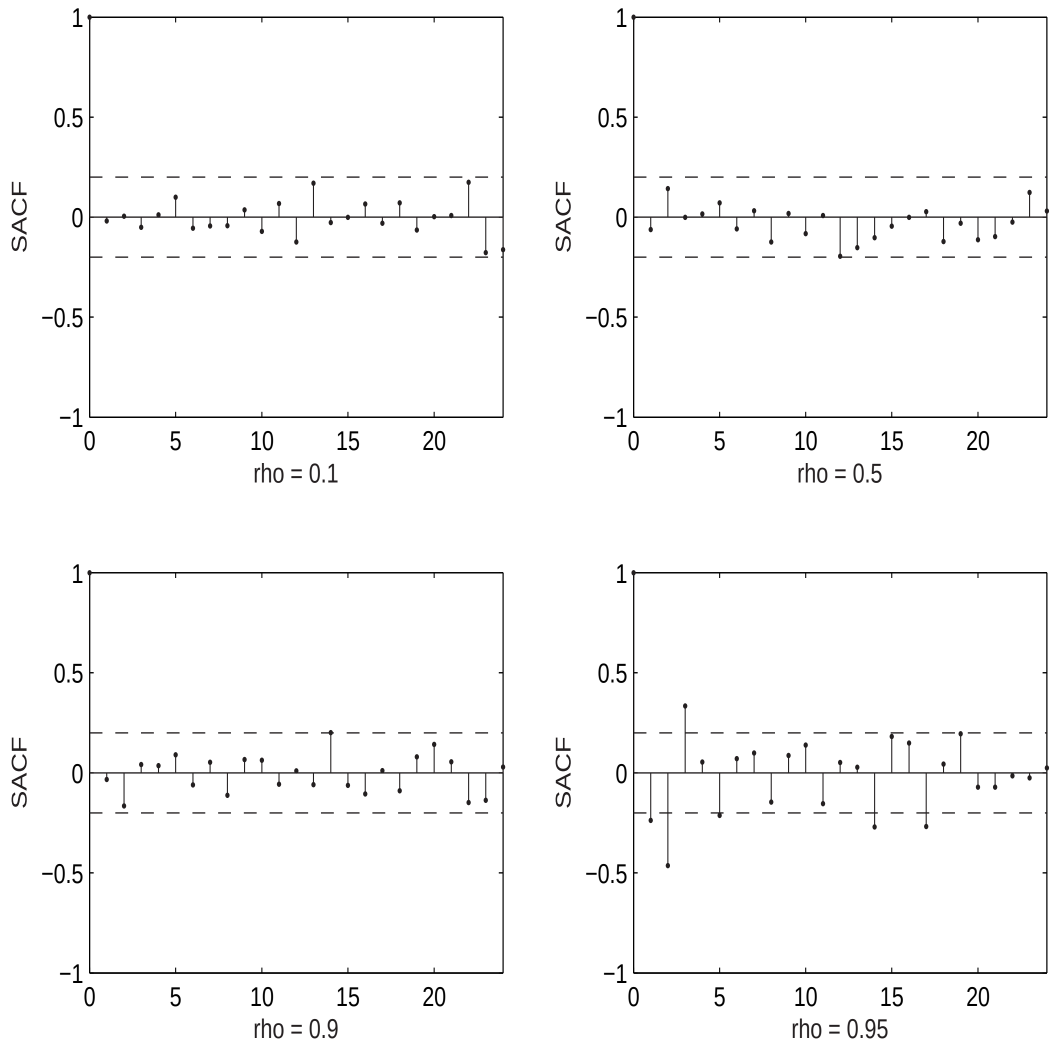

5. SPECIFICATION ISSUES

As we outlined in chapters 2 and 3, the model specification aspect has been neglected so far in the literature. In the empirical applications the model was chosen rather on an *ad hoc* basis and not on standard model selection criteria. Furthermore the weighting function in the MIDAS approaches is in some applications restricted, some not restricted, but without any explicit explanation given. In this chapter we want to take a closer look at model specification pertaining to forecasting economic time series. In particular we focus on lag selection in both approaches and whether or not the weighting function should be restricted . We will show that standard selection criteria can be computationally intensive and time consuming. Thus, results of the analysis can be interpreted as a first point of reference for applied forecasting.

5.1 Model Specification in Mixed-Frequency VAR Models

In the basic specification of the mixed-frequency VARMA with two variables, the MF-VAR(1), seven parameters have to be estimated.[1] The inclusion of more time series and higher-order lags leads to a large number of parameters to be estimated ($4p+3$ parameters in a bivariate VAR(p)). As the likelihood is calculated recursively, the computation time increases significantly with the number of parameters. Mittnik and Zadrozny (2005) state that forecasts are only feasible if variables are in compatible, cyclical form and not too many parameters have to be estimated.

We want to answer the following two questions. First, as the model operates

[1] We focus on model selection for the mixed-frequency VAR model based on the framework by Zadrozny (2008). The analysis can easily be applied to the framework by Hyung and Granger (2008).

at the highest frequency in the data, which lag length does describe the data well? Second, does the frequency mixture influence the lag selection? The importance of the second question can be illustrated with the following example: Suppose we find an AR(4) process (by some selection criterion) for a quarterly time series (for example GDP). Now we include a weekly time series and build a mixed-frequency VAR model operating at the weekly frequency. Allowing for the same past information set as in the univariate case, do we have to estimate an MF-VAR(48) model ($4m = 12 \cdot 4 = 48$, 48 weeks constitute 4 quarters)? The estimation of a model with 48 lags seems infeasible and contradicts the view of parsimony in econometrics.

To investigate these issues we conduct a small Monte Carlo study for Processes I and II illustrated in the previous chapter. We focus on these two processes, as the DGP is a (high-frequency) VAR process and the MF-VAR operates at the highest frequency. Thus we can draw clear cut conclusions about the order of the model. From the theoretical point of view we cannot decide which specific model to choose. Consider Process I which generates data from a VAR(1). As we skip $m-1$ observations to generate the mixed-frequency data set (every $m-th$ observation of the lower-frequency variable is observable), it is possible that the generated missing observations introduce new dynamics into the system. The figures in the appendices of the previous chapters demonstrate, that the autocorrelations patterns differ between the frequency mixtures m, although the data are generated with the same high-frequency VAR process.

We fit a mixed-frequency VAR(p) to each times series. We choose the optimal lag length via a corrected AIC criterion proposed by Hurvich and Tsay (1989) and used by Zadrozny (2008)

$$CAIC = L + \frac{2M}{1 - \frac{M+1}{mT}} \tag{5.1}$$

where L is -2 times the log-likelihood, M denotes the number of estimated parameters, and mT is the number of observations (higher frequency). We allow for a maximum of $2m$ lags for $m = 3$ and $m = 4$ and m lags for

$m = 12$. The time series length is $T = 100$. We simulate each process 100 times. This number is motivated by exponentially increasing computation time for estimating higher- order mixed-frequency VAR(p).[2]

Tables 5.1 and 5.2 report the average chosen lag lengths for each parameter combination and frequency mixture. For instance, consider Process I (Table 5.1). For $\rho = 0.1$, $\delta = 0$ and $m = 3$, on average one lag (1.000) was chosen by the corrected AIC criterion. For $\rho = 0.95$, $\delta = -3.5$, and $m = 12$ the majority of selected models were an VAR(1), as the average lag selection is 1.2. In general we can note, that the corrected AIC chooses low dimensional processes independently of the frequency mixture. For Process I there is a tendency for one lag, and for Process II (Table 5.2) for two lags for persistent series ($\rho > 0.5$). These results correspond to the fact that the true data-generating process is a VAR(1) and VAR(2), respectively. These findings suggest that there is no danger of parameter proliferation (at least in the bivariate case). Furthermore, it is demonstrated that small scale MF-VAR can capture rich data patterns, as graphed in the autocorrelation figures in the appendix to chapter 4.

[2] The estimation of all models up to a VAR(12) takes about 12h on a PC with an Intel Core 2 Quad processor.

Tab. 5.1: Lag selection in mixed-frequency VAR(p) models: Process I

		$m = 3$	$m = 4$	$m = 12$
$\rho = 0.1$	$\delta = 0$	1.000	1.000	1.000
	$\delta = -0.5$	1.000	1.000	1.000
	$\delta = -1.5$	1.000	1.000	1.000
	$\delta = -3.5$	1.000	1.000	1.000
$\rho = 0.5$	$\delta = 0$	1.000	1.000	1.000
	$\delta = -0.5$	1.000	1.000	1.000
	$\delta = -1.5$	1.060	1.180	1.140
	$\delta = -3.5$	1.040	1.180	1.120
$\rho = 0.9$	$\delta = 0$	1.400	1.040	1.160
	$\delta = -0.5$	1.200	1.220	1.320
	$\delta = -1.5$	1.060	1.180	1.280
	$\delta = -3.5$	1.060	1.020	1.220
$\rho = 0.95$	$\delta = 0$	1.400	1.040	1.000
	$\delta = -0.5$	1.140	1.200	1.280
	$\delta = -1.5$	1.060	1.120	1.220
	$\delta = -3.5$	1.160	1.040	1.200

Notes: Data are simulated from a high-frequency VAR as given in equation (4.1) with homoscedastic errors. For each parameter combination we estimate a mixed-frequency VAR(p) with $p_{max} = 2m$ ($p_{max} = m$ for $m = 12$). The table reports the average chosen lag length due to the corrected AIC criterion (5.1).

Tab. 5.2: Lag selection in mixed-frequency VAR(p) models: Process II

		$m = 3$	$m = 4$	$m = 12$
$\rho = 0.1$	$\delta = 0$	1.010	1.000	1.000
	$\delta = -0.5$	1.000	1.000	1.000
	$\delta = -1.5$	1.000	1.010	1.000
$\rho = 0.5$	$\delta = 0$	1.040	1.000	1.000
	$\delta = -0.5$	1.360	1.220	1.000
	$\delta = -1.5$	2.140	1.920	1.170
$\rho = 0.9$	$\delta = 0$	1.850	1.690	1.730
	$\delta = -0.5$	2.450	2.780	2.890
	$\delta = -1.5$	1.320	1.270	2.720

Notes: Data are simulated from a high-frequency VAR as given in equation (4.2) with homoscedastic errors. For each parameter combination we estimate a mixed-frequency VAR(p) with $p_{max} = 2m$ ($p_{max} = m$ for $m = 12$). The table reports the average chosen lag length due to the corrected AIC criterion (5.1).

5.2 MIDAS: Weighting Functions, Restrictions, and Lag Lengths

There are three important aspects in specifying a MIDAS model. First, one has to choose the weighting function itself. Second, should any restriction be imposed on the chosen weighting function? And third, how many lags should be included in the estimation? The first issue will be investigated in the Monte Carlo forecasting exercise in the next chapter. As we are interested in forecasting, the judgement will be based on forecast accuracy measures. Concerning the second and third issue, Ghysels, Sinko, and Valkanov (2007) suggest that, while estimating the weighting parameters, the lag selection is purely data driven. This implies that including another lag in the estimation which does not improve the estimation fit would be assigned a weight of zero. Furthermore it implies that no restrictions are imposed on the weighting function.

Loosely speaking, one should include as many lags as possible and let the data speak for themselves. This approach can possibly be recommended for financial applications where data are often sampled at 5-minute or daily frequency but not for macroeconomic data, which are often sampled at monthly and quarterly frequencies. The reason is that including more higher-frequency lags reduces the estimation sample in general, as less observations of the lower-frequency variable y_t can be used for estimation. In chapter 3 we reviewed the existing literature on MIDAS applications. In almost all (financial) applications the weighting function is restricted to obtain declining weights. This stands in contrast to the statement of Ghysels, Sinko, and Valkanov (2007) to let the data speak for themselves. We investigate the restriction and the issue of the number of lags subsequently, as we cannot rule these out as interrelated.

5.2.1 Restriction of the Weighting Function

We start with the issue of whether the weighting function should be restricted or not. We employ the three processes from the previous chapters for our analysis. To keep estimation issues simple, we allow only for two parameters in the exponential Almon lag specification (2.56).[3] We start at the point where we allow the data to speak for themselves as suggest by Ghysels, Sinko, and Valkanov (2007). Thus for the exponential Almon lag we have $-100 < \theta_1, \theta_2 < 100$ and $0 < \theta_1, \theta_2 < 300$ in our Beta weighting function.[4] We include $3m$ number of lags of the higher-frequency variable. For each simulated process set we estimate the two weighting parameters for each weighting function. We conduct 500 replications.

[3] All empirical applications so far used only two parameters in the weighting function.
[4] We are aware of the fact that these are restrictions, but we have to use them to avoid numbers which cannot be handled by the computer as they become too large (due to the exponential function).

Tab. 5.3: Average Parameter values of the exponential Almon and Beta weighting function: Process I

		$m=3$				$m=4$				$m=12$			
		Almon		Beta		Almon		Beta		Almon		Beta	
		θ_1	θ_2	θ_1	θ_2	θ_1	θ_2	θ_1	θ_2	θ_1	θ_2	θ_1	θ_2
$\rho=0.1$	$\delta=0$	7.48	-1.32	14.38	64.91	7.80	-1.09	14.92	71.40	4.97	-0.76	11.77	66.23
	$\delta=-0.5$	7.42	-1.28	13.95	60.74	6.79	-0.95	13.95	68.64	4.50	-0.69	11.39	70.82
	$\delta=-1.5$	8.10	-1.32	15.21	61.21	6.82	-1.01	16.82	72.99	4.37	-0.70	10.29	62.00
	$\delta=-3.5$	7.60	-1.44	14.40	58.46	6.79	-0.99	14.17	74.30	4.19	-0.63	12.17	65.24
$\rho=0.5$	$\delta=0$	10.84	-1.56	29.86	71.33	9.99	-1.44	26.39	82.73	5.75	-0.90	16.86	79.21
	$\delta=-0.5$	10.97	-1.93	27.90	75.50	10.66	-1.45	31.83	90.88	6.51	-0.94	18.78	84.36
	$\delta=-1.5$	11.54	-1.82	34.22	79.17	12.07	-1.70	32.58	91.14	7.36	-1.12	18.59	85.12
	$\delta=-3.5$	11.50	-2.16	42.07	85.45	11.29	-1.58	39.99	101.33	7.10	-1.11	22.55	86.25
$\rho=0.9$	$\delta=0$	10.45	-1.74	52.07	63.69	10.89	-2.02	58.68	81.60	10.97	-1.41	37.26	98.08
	$\delta=-0.5$	-19.40	-15.93	0.00	5.01	-17.52	-14.04	0.00	5.01	-4.83	-7.24	12.27	18.80
	$\delta=-1.5$	-21.09	-16.94	0.00	5.01	-19.23	-14.75	0.00	5.01	-5.24	-6.08	10.43	18.93
	$\delta=-3.5$	-14.17	-9.98	0.00	5.01	-13.98	-9.09	0.00	5.01	-3.65	-4.10	8.78	15.79
$\rho=0.95$	$\delta=0$	5.79	-1.79	49.58	48.77	7.58	-1.41	57.40	60.59	9.70	-1.76	47.19	79.18
	$\delta=-0.5$	-11.35	-8.91	0.00	5.01	-14.02	-10.22	0.00	5.02	-9.55	-8.40	0.60	5.30
	$\delta=-1.5$	-10.79	-8.54	0.00	5.01	-14.70	-11.24	0.00	5.02	-7.92	-6.84	0.54	5.50
	$\delta=-3.5$	-11.67	-8.96	0.00	5.01	-13.43	-9.42	0.00	5.02	-7.43	-6.66	0.55	5.28

Notes: Data are simulated from a high-frequency VAR as in equation (4.1). The parameter values are as follows: $\rho=\{0.10,0.50,0.90,0.95\}$ and $\delta_l=\{0,-0.5,-1.5,-3.5\}$. The series are simulated for $m\times T$ observations, where $m=\{3,4,12\}$ and $T=100$ (lower frequency). The low-frequency data are skip-sampled once every m observations from the simulated series. Each process is simulated 500 times. This table reports the average parameters within the exponential Almon lag (2.56) and Beta (2.57) weighting function. The parameters are not restricted. The number of included lags is fixed to $3m$.

Tab. 5.4: Average Parameter values of the exponential Almon and Beta weighting function: Process II

		$m=3$				$m=4$				$m=12$			
		Almon		Beta		Almon		Beta		Almon		Beta	
		θ_1	θ_2	θ_1	θ_2	θ_1	θ_2	θ_1	θ_2	θ_1	θ_2	θ_1	θ_2
$\rho=0.1$	$\delta=0$	7.61	-1.11	6.65	28.59	6.84	-1.23	6.16	29.74	3.77	-0.64	6.75	34.31
	$\delta=-0.5$	6.19	-1.12	5.57	25.80	5.85	-0.95	4.70	23.75	3.69	-0.57	5.97	33.38
	$\delta=-1.5$	7.10	-1.26	5.51	24.62	5.80	-1.16	4.05	23.50	3.24	-0.57	7.65	31.52
$\rho=0.5$	$\delta=0$	5.58	-1.84	10.73	29.59	9.47	-1.52	12.66	38.61	4.97	-0.87	8.55	37.20
	$\delta=-0.5$	2.86	-3.22	14.50	32.51	7.99	-1.24	11.59	39.20	4.82	-0.90	8.15	36.66
	$\delta=-1.5$	3.68	-2.85	15.56	35.94	9.75	-1.76	7.33	39.00	4.78	-0.85	7.71	34.21
$\rho=0.9$	$\delta=0$	0.88	-1.24	1.56	11.11	0.49	-3.02	1.12	19.00	7.94	-1.06	12.81	41.93
	$\delta=-0.5$	-0.63	-1.39	0.94	15.61	-5.91	-6.62	0.34	16.18	3.46	-1.98	1.89	19.99
	$\delta=-1.5$	-2.44	-3.72	0.21	15.73	5.78	-0.48	1.76	10.71	4.15	-0.45	6.10	49.90

Notes: Data are simulated from a high-frequency VAR as in equation (4.2). The parameter values are as follows: $\rho = \{0.10, 0.50, 0.90\}$ and $\delta_l = \{0, -0.5, -1.5\}$. The series are simulated for $m \times T$ observations, where $m = \{3, 4, 12\}$ and T = 100 (lower frequency). The low-frequency data are skip-sampled once every m observations from the simulated series. Each process is simulated 500 times. This table reports the average parameters within the exponential Almon lag (2.56) and Beta (2.57) weighting function. The parameters are not restricted. The number of included lags is fixed to $3m$.

In Table 5.3 we display the average two parameters (θ_1 and θ_2) for both the Almon and the Beta weighting functions for Process I. Each row represents a combination of ρ (persistence) and δ ('lead'). On average, relatively large values are chosen from the given data. Let us first consider the exponential Almon lag weighting function. For moderate persistent simulated time series ($\rho \leq 0.5$), a positive θ_1 and a negative θ_2 is estimated. This corresponds to the sharp, peaked weights where the peak is located between the second and fourth observation of the higher-frequency variable. Independently of the included lags, up to five positive weights (> 0.05) are assigned. For higher simulated persistent times series ($\rho \geq 0.9$) both weighting parameters are negative. This corresponds to a sharp declining weighting function with assigning a weight of almost 1 to the first lag of the higher- frequency variable. These results hold across different frequency mixtures. We plot the two typical shapes of the weighting function in Figure 5.1. The results obtained for the Beta weighting function can be interpreted in exactly the same way. The estimated parameters produce the same shape for the weighting function as the exponential Almon lag weighting function. This points to the fact that both weighting functions interpret the data in the same way.

Fig. 5.1: Examples of estimated weighting functions for Process I

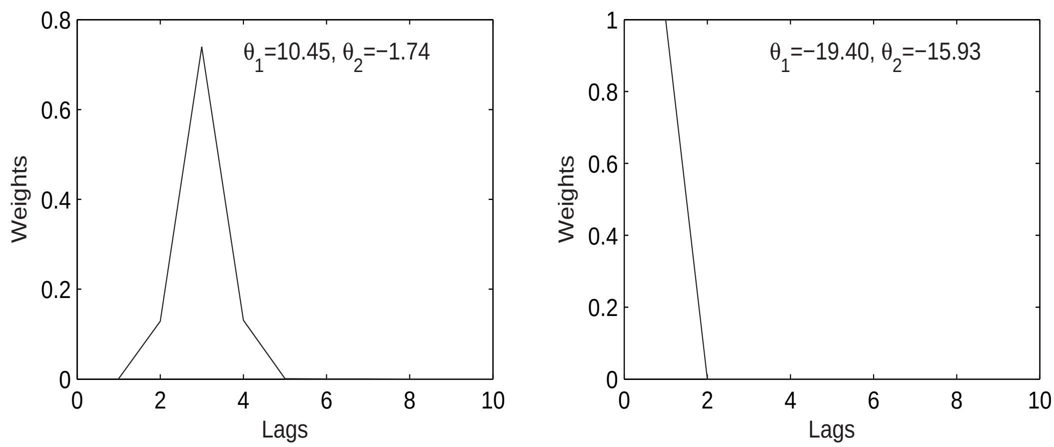

We find similar results for Process II. The average weighting parameters are reported in Table 5.4. In general we find either a sharp peak around lag 5 or

Tab. 5.5: Average Parameter values of the exponential Almon and Beta weighting function: Process III

		Almon		Beta	
		θ_1	θ_2	θ_1	θ_2
			$m = 3$		
WF 1	$\beta_1 = 0.6$	-1.394	0.066	0.966	4.348
	$\beta_1 = 3$	-1.357	0.054	0.965	4.275
	$\beta_1 = 6$	-1.388	0.066	0.961	4.335
WF 2	$\beta_1 = 0.6$	-0.376	0.009	0.975	1.628
	$\beta_1 = 3$	-0.367	0.008	0.976	1.643
	$\beta_1 = 6$	-0.370	0.008	0.976	1.638
			$m = 4$		
WF 1	$\beta_1 = 0.6$	-0.988	-0.531	0.955	11.986
	$\beta_1 = 3$	-0.961	-0.538	0.948	11.865
	$\beta_1 = 6$	-0.399	-0.784	0.863	12.493
WF 2	$\beta_1 = 0.6$	-0.435	0.014	0.968	1.957
	$\beta_1 = 3$	-0.429	0.013	0.968	1.979
	$\beta_1 = 6$	-0.430	0.013	0.968	1.979
			$m = 12$		
WF 1	$\beta_1 = 0.6$	-2.169	-5.028	0.643	16.770
	$\beta_1 = 3$	-0.528	-4.869	0.609	15.414
	$\beta_1 = 6$	-1.712	-4.685	0.449	11.008
WF 2	$\beta_1 = 0.6$	-0.727	-1.667	0.907	15.783
	$\beta_1 = 3$	-0.545	-1.375	0.820	13.421
	$\beta_1 = 6$	-0.282	-1.370	0.563	11.463

Notes: Data are simulated from a MIDAS regression as in equation (4.4). Each process is simulated 500 times. WF1 corresponds to an exponential Almon lag function with $\theta = (7 \times 10^{-4}, -5 \times 10^{-2})$ and WF2 with parameters $\theta = (7 \times 10^{-4}, -6 \times 10^{-3})$. This table reports the average parameters within the exponential Almon lag (2.56) and Beta (2.57). The parameters are not restricted. The number of included lags is fixed to $3m$.

a concentration of weights on the first two lags. The shape is in general the same as plotted in Figure 5.1.

For Process III we have tabulated the results in Table 5.5. For this process the *a priori* parameters are known from the data generation process. First we can conclude that for this process, the average parameters do not converge to their true values. Nevertheless, the obtained values for both weighting functions generate declining weights. But this decline is rather sharp in all cases, that is the second weighting function (WF 2) cannot be reproduced.

The results for Processes I and II may not be intuitive for the forecasting process. Why should the value of the leading indicator obtained three or four periods before be more important than the most recent one. Furthermore, why should only the most recent one be used in forecasting, while discarding the other lags? A declining weight function is more intuitive for the forecaster. Does the forecast accuracy improve as we restrict the weighting function to ensure declining weights, in comparison to unrestricted weighting functions? As we outlined in the literature section, in most applications the weighting function is restricted. Again, we employ our Monte Carlo study to answer this question. In the case of the exponential Almon lag both θ_1 and θ_2 have to be negative to ensure declining weights.[5] As $|\theta_i|$ increases the more concentrated are the weights on the first lags. The Beta weighting function is declining if $\theta_1 = 1$ and $\theta_2 > 1$. As θ_2 increases, the more rapidly the weights decline.

For each simulated time series we conduct a one-step-ahead forecast of the lower-frequency variable. We compare the squared error obtained with the unrestricted weighting functions and the restricted weighting functions. For the exponential Almon lag we impose the restriction $-1 < \theta_1 < 0$ and $-0.1 < \theta_2 < 0$ which ensures declining weights. The higher the absolute values of θ_i, the more weight is assigned to the more recent observations of the leading time series. For the Beta weighting function we impose $\theta_1 = 1$

[5] Ghysels, Sinko, and Valkanov (2007) state that $\theta_2 < 0$ guarantees declining weights (p.57). This is wrong. For some positive values of θ_1 the weighting function can be hump shaped or increasing in weights. Ghysels, Sinko, and Valkanov (2007) even show this in a graph in their paper (p. 58).

and $1 < \theta_2 < 15$. We restrict the sharpness of decline in both cases, as from the previous results we know that there is a tendency to assign the whole weight to the first lag.

Table 5.6 displays the results for Process I. Each entry represents the ratio of the mean-squared error (MSE) of a one-step-ahead forecast of the MIDAS model of unrestricted and restricted weighting functions. A ratio smaller than 1 indicates that the unrestricted weighting function produces smaller MSEs. We find that there is no difference between the restricted and unrestricted cases. Only for the case $\rho = 0.95$ and $\delta_h < 0$ is there a marginal advantage of the unrestricted weighting function. Given the results there is no need to put any restrictions on the weighting function.

Similar results can be reported for Process II tabulated in Table 5.7. In case of $\rho = 0.9$ and $\delta = -1.5$ the restricted weighting function results in higher forecasting errors. For other parameter combinations there seems to be no difference in forecasting accuracy between restricted and unrestricted weighting functions.

Conducting the same exercise for Process III seems not to make any sense. This is first, because the data-generating process assumes declining weights. And second, we confirmed that for this data structure, decreasing weights are estimated. Table 5.8 reports the ratio of the unrestricted and the restricted weighting functions. The results provide an interesting interpretation. In most cases the restricted weighting function results in higher MSEs; in some cases, (for example WF 1: $\beta_1 = 6$ and $m = 4$), the restriction yields lower MSEs. Thus, we have found another situation where restriction makes sense. This points to the fact that there can be an 'optimal restriction' for the weighting function. But could be difficult to find in applied work.

This section demonstrates that there is no clear answer to whether the weighting function should be restricted or not. At least there seems to be no disadvantage if we prefer one of the options. There some exceptions. In some cases (highly persistent time series) there is a tendency to restrict the weighting function which leads to lower forecasting errors.

In general we cannot determine whether our results hold in general, as the number of included lags was fixed. Furthermore our imposed restrictions are only an example out of many possible other restrictions. There may be room for improvement in forecasting accuracy with restricted weighting functions when the number of lags is increased or decreased. We will investigate this issue in the next subsection.

Tab. 5.6: MSE ratios of restricted and unrestricted weighting functions: Process I

		$m = 3$		$m = 4$		$m = 12$	
		Almon	Beta	Almon	Beta	Almon	Beta
$\rho = 0.1$	$\delta = 0$	1.017	1.005	1.011	1.009	1.002	1.008
	$\delta = -0.5$	1.016	1.006	1.009	1.008	1.012	1.011
	$\delta = -1.5$	1.018	1.003	1.003	1.003	1.018	1.011
	$\delta = -3.5$	1.012	1.012	1.018	1.003	1.012	1.016
$\rho = 0.5$	$\delta = 0$	1.016	1.005	1.013	0.999	1.009	1.006
	$\delta = -0.5$	1.010	1.004	1.009	1.004	0.998	1.001
	$\delta = -1.5$	1.011	1.007	1.003	1.004	1.015	1.019
	$\delta = -3.5$	1.001	1.002	1.007	1.005	1.010	1.014
$\rho = 0.9$	$\delta = 0$	1.002	1.001	1.006	1.002	1.004	1.004
	$\delta = -0.5$	0.983	0.991	0.986	0.986	0.996	0.990
	$\delta = -1.5$	0.983	0.990	0.987	0.987	1.004	0.991
	$\delta = -3.5$	0.984	0.991	0.985	0.985	0.996	0.990
$\rho = 0.95$	$\delta = 0$	1.000	1.000	1.007	1.006	1.006	1.000
	$\delta = -0.5$	0.985	0.993	0.986	0.986	0.996	0.977
	$\delta = -1.5$	0.984	0.992	0.985	0.985	0.996	0.977
	$\delta = -3.5$	0.986	0.992	0.987	0.987	0.997	0.982

Notes: Data are simulated from a high-frequency VAR as in equation (4.1) as in the previous table. We estimate a MIDAS model with restricted and unrestricted weighting functions and calculate one-step ahead forecasts of the lower frequency. We used the following restrictions: exponential Almon lag weighting function (2.56): $-1 < \theta_1 < 0$ and $-0.1 < \theta_2 < 0$, Beta weighting function (2.57): $\theta_1 = 1$ and $1 < \theta_2 < 15$. This table displays the ratio between MSE from MIDAS regressions with restricted and unrestricted weighting functions. The number of included lags is fixed to $3m$.

Tab. 5.7: MSE ratios of restricted and unrestricted weighting functions: Process II

		$m = 3$		$m = 4$		$m = 12$	
		Almon	Beta	Almon	Beta	Almon	Beta
$\rho = 0.1$	$\delta = 0$	1.012	0.980	1.011	1.020	1.000	1.017
	$\delta = -0.5$	0.987	1.000	0.995	1.004	1.021	1.002
	$\delta = -1.5$	1.017	1.023	1.018	1.000	1.019	1.024
$\rho = 0.5$	$\delta = 0$	1.024	0.984	1.003	1.018	1.036	1.020
	$\delta = -0.5$	1.005	1.029	1.003	0.995	1.001	0.988
	$\delta = -1.5$	0.999	1.046	1.013	1.012	1.025	1.004
$\rho = 0.9$	$\delta = 0$	1.007	1.003	1.006	0.995	1.006	1.002
	$\delta = -0.5$	0.992	1.009	0.956	0.954	1.000	0.982
	$\delta = -1.5$	0.638	0.831	0.425	0.812	0.858	0.745

Notes: Data are simulated from a high-frequency VAR as in equation (4.2) as in the previous table. Each process is simulated 500 times. We estimate a MIDAS model with restricted and unrestricted weighting functions and calculate one-step ahead forecasts of the lower frequency. We used the following restrictions: exponential Almon lag weighting function (2.56): $-1 < \theta_1 < 0$ and $-0.1 < \theta_2 < 0$, Beta weighting function (2.57): $\theta_1 = 1$ and $1 < \theta_2 < 15$. This table displays the ratio between MSE from MIDAS regressions with restricted and unrestricted weighting functions. The number of included lags is fixed to $3m$.

5.2.2 Choosing the Number of Lags

Finally, we want to investigate how many lags one should include in the estimation of MIDAS models. Surprisingly the current literature is relatively uninformative on how many lags are included and on which criterion the number of lags is chosen. There are no formal selection criteria and the lengths seem to be chosen rather *ad hoc*.

In the simplest bivariate case, in MIDAS regressions, only four parameters have to be estimated independently of the number of included lags. So there is no threat of proliferation of parameters and (or) overparametrization as in standard linear models by inclusion of more lags. Loosely speaking, we may include as many lags as possible, and positive weight assignments determine the lag lengths implicitly. In practice this approach is limited. The more lags

Tab. 5.8: MSE ratios of restricted and unrestricted weighting functions: Process III

		$m = 3$		$m = 4$		$m = 12$	
		Almon	Beta	Almon	Beta	Almon	Beta
WF 1	$\beta_1 = 0.6$	1.281	0.934	1.551	0.790	1.004	0.668
	$\beta_1 = 3$	1.299	0.816	1.771	0.796	0.974	0.662
	$\beta_1 = 6$	1.415	0.962	2.079	0.918	1.008	0.704
WF 2	$\beta_1 = 0.6$	0.995	0.841	0.960	0.767	1.168	0.520
	$\beta_1 = 3$	0.984	0.812	0.952	0.794	1.055	0.477
	$\beta_1 = 6$	0.975	0.826	0.946	0.777	1.146	0.538

Notes: Data are simulated from a MIDAS regression as in equation (4.4). Each process is simulated 500 times. WF1 corresponds to an exponential Almon lag function with $\theta = (7\times10^{-4}, -5\times10^{-2})$ and WF2 with parameters $\theta = (7\times10^{-4}, -6\times10^{-3})$. We estimate a MIDAS model with restricted and unrestricted weighting functions and calculate one-step ahead forecasts of the lower frequency. We used the following restrictions: exponential Almon lag weighting function (2.56): $-1 < \theta_1 < 0$ and $-0.1 < \theta_2 < 0$, Beta weighting function (2.57): $\theta_1 = 1$ and $1 < \theta_2 < 15$. This table displays the ratio between MSE from MIDAS regressions with restricted and unrestricted weighting functions. The number of included lags is fixed to $3m$.

that are included, the smaller the estimation sample. In financial applications with minute and daily data available, the forecaster can be more generous with the number of included lags. But in macroeconomic forecasting this problem can be severe.[6] Ghysels, Santa-Clara, and Valkanov (2004) state that standard selection procedures such as the Akaike or Schwarz criterion (which are often used for model selection in forecasting) can be applied in the MIDAS context.[7] But no article has applied such a criterion. We consider one of the standard selection criteria, the Bayesian Information criterion (BIC)

$$BIC = \ln(\hat{\sigma}^2) + \frac{M}{T} \ln T$$

[6] To give an example: for the purpose of forecasting German GDP growth the data set often starts in 1991 to avoid structural breaks due to reunification. Using a monthly indicator with 12 included lags reduces the estimation sample by one year. The inclusion of 24 lags reduces the sample by two years, and so forth.

[7] The authors do not investigate this issue in their paper.

where $\hat{\sigma}^2$ is the estimated variance of the error term and M is the number of estimated parameters. The classical model selection criteria impose a penalty on additional estimated parameters. This does not apply in the MIDAS context, as the number of estimated parameters is constant. The fit of the regression is only determined by the error variance and the time series length T (included in the estimation). The BIC is decreasing with increasing T. This demonstrates a trade-off between included lags and the number of observations (and information) included for estimation.

In our Monte Carlo experiment we use two criteria to assess how many lags should be included for the estimation and forecasting process. For each simulated time series model we allow for m up to $8m$ lags. On the one hand, we choose the optimal lag due to the BIC criterion (in-sample criterion). On the other hand, we assess the optimality due to forecast accuracy by forecasting the lower frequency time series (OSC criterion) one-step-ahead. Granger (1993) pointed out that in-sample selection measures (as BIC) frequently fail to provide strong implications for the out-of-sample performance. We employ both the restricted and unrestricted weighting functions to disentangle the possible existing relationship between forecast performance, lag length and weighting function restrictions.[8]

Tables 5.9 tabulates the results for the Almon weighting function for Process I. The third and fourth columns report the average chosen lag length due to the BIC criterion (restricted versus unrestricted weighting function). The fifth and seventh columns report the corresponding lags chosen due to the OSC criterion. Columns six and eight exhibit the average RMSE at the chosen lag due to the OSC criterion.

First we can note that on the one hand there are some differences in the number of lags between the BIC and the OSC. For $m = 3$ and $m = 4$ there is a tendency of the OSC to choose a fewer lags in comparison to the BIC criterion. For $m = 12$ both criteria deliver similar results. On the other hand restricted and unrestricted weighting functions deliver the same results. One notable exception is for $m = 12$ where in the restricted BIC case and $\rho = 0.95$

[8] We employ the same restriction as in the previous subsection.

the number of chosen lags are considerably higher. For lower persistent time series the average number of lags is about $4m$ or $5m$. For higher persistent series there is a tendency of the BIC criterion to include the maximum of allowed lags (24 for $m = 3$ and 32 for $m = 4$).

Comparing the RMSEs for OSC criterion we find interesting results. In contrast to the previous results (Table 5.6) there is a difference in forecasting performance between restricted and unrestricted weighting functions. In case of persistent time series ($\rho \geq 0.9$) the RMSE is lower (comparing the fourth and sixth columns) for restricted weighting functions. The opposite is true for $\rho \leq 0.5$ where the RMSE is higher for restricted weighting functions.

The interpretations for the exponential Almon lag function can also be drawn for the Beta weighting function. The results are displayed in Table 5.10.

For Process II we find similar results concerning the lag length. The lag lengths to be chosen are around $4m$ independently of the selection criteria. The results are reported in Tables 5.11 and 5.12. There is one notable exception to Process I. Restrictions on the weighting function do not improve the forecasting accuracy. The RMSE of restricted weighting functions are in any case higher than the unrestricted counterpart. Again, the interpretations for the Beta weighting functions are the same as for the exponential Almon lag weighting function.

Taking these results and the results from the previous subsection, we can conclude that forecast performance, the number of included lags and weighting function restrictions are interrelated. Our results stand in contrast to the statement of Ghysels, Sinko, and Valkanov (2007) to let the data speak for themselves when forecasting is concerned. Lag length selection and the restriction of the weighting function should explicitly investigated to optimise forecasting results.

Tab. 5.9: Exponential Almon lag weighting function: Lag length selection for Process I

		BIC		Out-of-sample			
				unrestricted		restricted	
		$m = 3$					
		unrestricted	restricted	lag	RMSE	lag	RMSE
$\rho = 0.1$	$\delta = 0$	12	11	12	0.768	11	0.861
	$\delta = -0.5$	13	12	12	0.841	11	0.945
	$\delta = -1.5$	13	11	12	0.867	12	0.944
	$\delta = -3.5$	12	11	12	0.909	12	0.994
$\rho = 0.5$	$\delta = 0$	13	12	12	0.884	12	0.949
	$\delta = -0.5$	12	11	13	1.204	13	1.259
	$\delta = -1.5$	11	10	13	1.852	12	1.769
	$\delta = -3.5$	10	10	13	1.927	12	1.852
$\rho = 0.9$	$\delta = 0$	15	13	12	0.773	12	0.843
	$\delta = -0.5$	21	23	15	0.292	15	0.243
	$\delta = -1.5$	23	24	17	0.266	16	0.194
	$\delta = -3.5$	23	24	16	0.186	15	0.135
$\rho = 0.95$	$\delta = 0$	16	15	13	0.883	12	0.960
	$\delta = -0.5$	24	24	16	0.147	17	0.118
	$\delta = -1.5$	24	24	17	0.143	17	0.113
	$\delta = -3.5$	24	24	17	0.136	18	0.111
		$m = 4$					
		unrestricted	restricted	lag	RMSE	lag	RMSE
$\rho = 0.1$	$\delta = 0$	16	15	17	0.897	15	1.000
	$\delta = -0.5$	16	15	16	0.841	15	0.941
	$\delta = -1.5$	17	16	17	0.783	15	0.871
	$\delta = -3.5$	16	14	16	0.815	16	0.898
$\rho = 0.5$	$\delta = 0$	16	15	16	0.839	15	0.929
	$\delta = -0.5$	16	15	17	0.940	17	1.001
	$\delta = -1.5$	14	13	18	1.207	15	1.197
	$\delta = -3.5$	14	13	18	1.522	16	1.492
$\rho = 0.9$	$\delta = 0$	19	18	16	0.893	17	0.959
	$\delta = -0.5$	23	28	21	0.339	18	0.307
	$\delta = -1.5$	23	29	22	0.348	19	0.291
	$\delta = -3.5$	23	30	22	0.273	20	0.227
$\rho = 0.95$	$\delta = 0$	20	19	17	1.096	17	1.153
	$\delta = -0.5$	32	32	23	0.152	23	0.123
	$\delta = -1.5$	32	32	22	0.102	22	0.077
	$\delta = -3.5$	32	32	24	0.093	22	0.073
		$m = 12$					
		unrestricted	restricted	lag	RMSE	lag	RMSE
$\rho = 0.1$	$\delta = 0$	40	43	49	0.884	46	0.962
	$\delta = -0.5$	41	43	51	0.943	47	1.026
	$\delta = -1.5$	42	42	49	0.827	48	0.919
	$\delta = -3.5$	44	44	46	0.849	44	0.919
$\rho = 0.5$	$\delta = 0$	40	40	49	0.825	47	0.897
	$\delta = -0.5$	45	44	47	0.877	48	0.940
	$\delta = -1.5$	42	44	47	0.817	43	0.889
	$\delta = -3.5$	41	42	48	0.893	48	0.976
$\rho = 0.9$	$\delta = 0$	48	48	48	0.874	48	0.946
	$\delta = -0.5$	37	43	50	1.236	51	1.282
	$\delta = -1.5$	36	44	55	1.243	54	1.269
	$\delta = -3.5$	32	41	55	1.286	52	1.304
$\rho = 0.95$	$\delta = 0$	51	51	49	1.066	48	1.105
	$\delta = -0.5$	37	64	59	0.252	49	0.287
	$\delta = -1.5$	36	64	59	0.272	47	0.285
	$\delta = -3.5$	36	64	61	0.271	46	0.289

Notes: Data are simulated from a high-frequency VAR as in equation (4.1). Each process is simulated 500 times. The maximum lag length is $8m$. This tables reports the average lag length due to the BIC criterion and the out-of-sample criterion both for restricted and unrestricted weighting functions for the exponential Almon lag weighting function. For the OSC criterion the average RMSE is reported for the chosen lag.

Tab. 5.10: Beta lag weighting function: Lag length selection for Process I

		BIC		Out-of-sample			
				unrestricted		restricted	
		$m = 3$					
		unrestricted	restricted	lag	RMSE	lag	RMSE
$\rho = 0.1$	$\delta = 0$	12	11	13	0.848	12	0.876
	$\delta = -0.5$	13	11	12	0.913	12	0.969
	$\delta = -1.5$	12	11	13	0.900	12	0.955
	$\delta = -3.5$	12	11	12	0.986	12	1.028
$\rho = 0.5$	$\delta = 0$	13	11	12	0.929	12	0.961
	$\delta = -0.5$	12	10	12	1.173	15	1.251
	$\delta = -1.5$	11	8	10	1.571	17	1.545
	$\delta = -3.5$	11	8	8	1.561	18	1.525
$\rho = 0.9$	$\delta = 0$	14	13	12	0.807	12	0.834
	$\delta = -0.5$	23	24	12	0.402	14	0.230
	$\delta = -1.5$	24	24	12	0.414	15	0.182
	$\delta = -3.5$	24	24	12	0.259	14	0.130
$\rho = 0.95$	$\delta = 0$	15	15	12	0.886	12	0.943
	$\delta = -0.5$	23	24	15	0.219	17	0.113
	$\delta = -1.5$	23	24	15	0.286	17	0.107
	$\delta = -3.5$	23	24	15	0.239	17	0.108
		$m = 4$					
$\rho = 0.1$	$\delta = 0$	16	14	18	0.944	16	1.002
	$\delta = -0.5$	16	14	17	0.924	15	0.953
	$\delta = -1.5$	16	14	17	0.844	16	0.886
	$\delta = -3.5$	15	14	17	0.847	16	0.901
$\rho = 0.5$	$\delta = 0$	16	14	17	0.880	15	0.934
	$\delta = -0.5$	15	13	16	0.954	19	0.999
	$\delta = -1.5$	13	11	16	1.073	22	1.108
	$\delta = -3.5$	13	10	15	1.271	24	1.296
$\rho = 0.9$	$\delta = 0$	18	18	17	0.908	16	0.951
	$\delta = -0.5$	28	31	13	0.462	17	0.306
	$\delta = -1.5$	31	32	13	0.570	18	0.282
	$\delta = -3.5$	31	32	13	0.401	19	0.222
$\rho = 0.95$	$\delta = 0$	19	19	17	1.078	16	1.117
	$\delta = -0.5$	32	32	18	0.314	22	0.120
	$\delta = -1.5$	32	32	17	0.196	21	0.078
	$\delta = -3.5$	32	32	17	0.215	22	0.075
		$m = 12$					
$\rho = 0.1$	$\delta = 0$	43	41	52	0.885	45	0.940
	$\delta = -0.5$	46	41	55	0.913	46	1.000
	$\delta = -1.5$	43	41	49	0.808	47	0.894
	$\delta = -3.5$	44	42	49	0.853	44	0.911
$\rho = 0.5$	$\delta = 0$	41	40	50	0.809	48	0.876
	$\delta = -0.5$	44	42	51	0.865	49	0.915
	$\delta = -1.5$	45	42	50	0.802	44	0.864
	$\delta = -3.5$	44	40	52	0.887	48	0.953
$\rho = 0.9$	$\delta = 0$	50	47	51	0.851	48	0.911
	$\delta = -0.5$	41	27	48	1.631	65	0.709
	$\delta = -1.5$	42	26	43	1.693	66	0.649
	$\delta = -3.5$	39	26	45	1.731	64	0.674
$\rho = 0.95$	$\delta = 0$	52	50	47	1.041	47	1.094
	$\delta = -0.5$	65	77	47	0.973	54	0.288
	$\delta = -1.5$	66	77	44	1.244	50	0.278
	$\delta = -3.5$	67	80	43	1.149	51	0.282

Notes: Data are simulated from a high-frequency VAR as in equation (4.1). Each process is simulated 500 times. The maximum lag length is $8m$. This tables reports the average lag length due to the BIC criterion and the out-of-sample criterion both for restricted and unrestricted weighting functions for the Beta weighting function. For the OSC criterion the average RMSE is reported for the chosen lag.

Tab. 5.11: Exponential Almon lag weighting function: Lag length selection for Process II

		BIC		Out-of-sample			
		unrestricted	restricted	unrestricted		restricted	
		$m = 3$					
		lag	lag	lag	RMSE	lag	RMSE
$\rho = 0.1$	$\delta = 0$	13	11	13	0.875	11	0.974
	$\delta = -0.5$	12	10	13	0.672	12	0.764
	$\delta = -1.5$	13	12	12	0.677	11	0.742
$\rho = 0.5$	$\delta = 0$	12	10	14	0.955	12	1.033
	$\delta = -0.5$	10	9	14	1.044	13	1.043
	$\delta = -1.5$	10	10	13	0.982	12	1.007
$\rho = 0.9$	$\delta = 0$	12	11	12	1.127	11	1.175
	$\delta = -0.5$	11	10	14	1.135	12	1.841
	$\delta = -1.5$	10	7	13	2.003	14	2.462
		$m = 4$					
$\rho = 0.1$	$\delta = 0$	16	15	15	0.695	15	0.783
	$\delta = -0.5$	17	16	17	0.672	15	0.765
	$\delta = -1.5$	16	15	15	0.780	15	0.874
$\rho = 0.5$	$\delta = 0$	17	15	16	0.908	16	0.975
	$\delta = -0.5$	16	15	17	0.817	15	0.919
	$\delta = -1.5$	15	14	17	0.683	14	0.697
$\rho = 0.9$	$\delta = 0$	15	15	18	1.218	17	1.227
	$\delta = -0.5$	14	13	18	2.191	17	2.104
	$\delta = -1.5$	16	17	17	0.207	16	0.422
		$m = 12$					
$\rho = 0.1$	$\delta = 0$	43	44	49	0.935	48	1.014
	$\delta = -0.5$	44	44	48	0.889	46	0.968
	$\delta = -1.5$	43	44	50	0.788	46	0.863
$\rho = 0.5$	$\delta = 0$	43	45	49	0.933	47	0.993
	$\delta = -0.5$	43	44	48	0.753	46	0.829
	$\delta = -1.5$	43	42	50	0.830	47	0.927
$\rho = 0.9$	$\delta = 0$	42	44	49	0.882	48	0.956
	$\delta = -0.5$	44	44	50	0.571	50	0.634
	$\delta = -1.5$	40	39	46	0.593	39	0.931

Notes: Data are simulated from a high-frequency VAR as in equation (4.1). Each process is simulated 500 times. The maximum lag length is $8m$. This tables reports the average lag length due to the BIC criterion and the out-of-sample criterion both for restricted and unrestricted weighting functions for the exponential Almon lag weighting function. For the OSC criterion the average RMSE is reported for the chosen lag.

Tab. 5.12: Beta lag weighting function: Lag length selection for Process II

		BIC		Out-of-sample			
		unrestricted	restricted	unrestricted		restricted	
			$m = 3$				
		lag	lag	lag	RMSE	lag	RMSE
$\rho = 0.1$	$\delta = 0$	12	11	13	0.960	1	1.070
	$\delta = -0.5$	12	10	13	0.769	1	0.901
	$\delta = -1.5$	12	11	13	0.754	1	0.871
$\rho = 0.5$	$\delta = 0$	12	10	12	0.968	1	0.829
	$\delta = -0.5$	11	8	14	0.940	1	0.942
	$\delta = -1.5$	9	6	12	0.754	1	0.783
$\rho = 0.9$	$\delta = 0$	11	11	9	0.973	1	0.810
	$\delta = -0.5$	12	6	7	1.652	1	1.505
	$\delta = -1.5$	10	3	12	1.442	2	1.528
			$m = 4$				
		lag	lag	lag	RMSE	lag	RMSE
$\rho = 0.1$	$\delta = 0$	15	14	17	0.727	14	0.779
	$\delta = -0.5$	16	16	17	0.729	15	0.760
	$\delta = -1.5$	15	15	16	0.861	14	0.867
$\rho = 0.5$	$\delta = 0$	16	14	17	0.929	18	0.977
	$\delta = -0.5$	17	14	17	0.882	17	0.936
	$\delta = -1.5$	14	13	17	0.788	15	0.707
$\rho = 0.9$	$\delta = 0$	15	12	13	1.097	22	1.063
	$\delta = -0.5$	13	5	19	1.400	26	1.315
	$\delta = -1.5$	22	4	19	0.396	23	0.493
			$m = 12$				
$\rho = 0.1$	$\delta = 0$	44	42	51	0.928	47	0.999
	$\delta = -0.5$	46	43	51	0.859	47	0.936
	$\delta = -1.5$	45	41	53	0.779	47	0.839
$\rho = 0.5$	$\delta = 0$	45	42	52	0.891	47	0.974
	$\delta = -0.5$	45	44	51	0.730	45	0.798
	$\delta = -1.5$	45	41	51	0.862	48	0.913
$\rho = 0.9$	$\delta = 0$	46	43	49	0.877	50	0.936
	$\delta = -0.5$	45	31	53	0.525	33	0.621
	$\delta = -1.5$	39	16	44	0.521	35	0.878

Notes: Data are simulated from a high-frequency VAR as in equation (4.1). Each process is simulated 500 times. The maximum lag length is $8m$. This tables reports the average lag length due to the BIC criterion and the out-of-sample criterion both for restricted and unrestricted weighting functions for the Beta weighting function. For the OSC criterion the average RMSE is reported for the chosen lag.

5.3 Suggestions for Applied Forecasting

Given our specification results we want to draw some conclusions and to give the following recommendations for specifying a mixed-frequency time series model:

1. Consider only small orders (up to 4) of mixed-frequency VAR models, independently of the frequency mixture and the data structure. Low scale MF-VAR can handle many data structures even for high-frequency mixtures. Larger models would increase the computational costs.

2. Forecast performance, lag length selection and restricting the weighting functions in the MIDAS framework are interrelated.

3. The number of included lags in MIDAS models is approximately $4m$ or $5m$.

4. Within the MIDAS framework the BIC criterion delivers often the same results concerning lag selection as an out-of-sample criterion across different data structures. For persistent time series the BIC suggests to include too many lags for estimation.

5. Whether the weighting functions should be restricted or not depends on the data structure. In the case of strongly persistent time series the weighting function can lead to greater forecasting accuracy. On the other side there are data structure where the restricted weighting functions deliver higher forecasting errors.

6. There is no difference in specification between the exponential Almon lag and the Beta weighting function.

We suggest the following procedure for MIDAS model specification :

1. The number of included lags should be chosen via the BIC criterion. Start with $4m$ and vary the number of included lags by $\pm 2m$.

2. Check the autocorrelation function of your target variable. For strongly persistent series restrict your weighting function otherwise not.

3. In case of a pseudo-real-time out-of-sample forecasting exercise try both the restricted and the unrestricted scheme.

6. MONTE CARLO FORECASTING STUDY

In this chapter we systematically compare the forecasting performance of the two mixed-frequency time series models. We do not compare them only against each other, but also with single-frequency time series models. We analyze whether there is a systematic improvement in forecasting accuracy by employing more advanced mixed-frequency models. We focus on short-run forecasting: one-step ahead and intra-period forecasts. Given our rich data structure from the four processes outlined in chapter 4, we expect to draw clear-cut conclusions. Before we study each process separately we outline some theoretical results concerning the use of mixed-frequency data for forecasting.

6.1 Forecasting with Mixed-frequency Data: Some Theoretical Reasoning

Ghysels and Valkanov (2006) are the first to investigate theoretically the gains in forecasting from using mixed-data sampling. These authors consider three different information sets. If all data are available at the high frequency, the largest information set is denoted by $\mathcal{I}_t$. This is the best but in practice often infeasible. The second-best solution is to use mixed high- and low-frequency variables ($\mathcal{I}_t^M$). The third information set ($\mathcal{I}_t^A$) is obtained from temporal aggregation, where all data are aggregated to the least frequency. To appraise the forecasting performance, we define the mean squared error of three linear predictors:

$$\begin{aligned} P[y_{t+h}|\mathcal{I}_t]: \quad MSE(h,\mathcal{I}_t) &\equiv E\left(y_{t+h} - P[y_{t+h}|\mathcal{I}_t]\right) \\ P[y_{t+h}|\mathcal{I}_t^M]: \quad MSE(h,\mathcal{I}_t^M) &\equiv E\left(y_{t+h} - P[y_{t+h}|\mathcal{I}_t^M]\right) \\ P[y_{t+h}|\mathcal{I}_t^A]: \quad MSE(h,\mathcal{I}_t^A) &\equiv E\left(y_{t+h} - P[y_{t+h}|\mathcal{I}_t^A]\right) \end{aligned} \tag{6.1}$$

for $h \in N$. In general we expect the following ranking

$$MSE(h,\mathcal{I}_t) \leq MSE(h,\mathcal{I}_t^M) \leq MSE(h,\mathcal{I}_t^A) \quad \forall h \in \mathbb{N}. \tag{6.2}$$

Ghysels and Valkanov (2006) prove within a VAR framework that under certain conditions the following ranking holds

$$MSE(h,\mathcal{I}_t) = MSE(h,\mathcal{I}_t^M) < MSE(h,\mathcal{I}_t^A) \quad \forall h \in \mathbb{N}. \tag{6.3}$$

In other words, there are circumstances where mixed-data sampling achieves the same predictive accuracy as we would get if we had all the disaggregated information available; we would be better off using aggregate data for comparison. The proof is based on Granger causality properties in a framework provided by Dufour and Renault (1998). However, these conditions cannot be tested empirically, as they are built upon the availability of the entire high-frequency process.

Hyung and Granger (2008) prove that a mixed-frequency VAR model produces the lowest MSE when compared with temporally aggregated single-frequency models and models which use a within-quarter variable (as in Rathjens and Robins (1993)). This result holds only when all parameters are known and does not necessarily carry over to estimated processes. The proof relies on state-space representations.

The theorems in Ghysels and Valkanov (2006) and Hyung and Granger (2008) provide a theoretical basis for the empirical findings, that using data sampled at different frequencies improves forecasting accuracy in comparison with temporally aggregated data.

6.2 The General Set-up

Our Monte Carlo study extends and modifies that of Ghysels and Valkanov (2006) (GV henceforth) in several ways. First, GV consider only one data-generating process (Process I in this thesis) for bivariate time series. We investigate the forecasting performance for four processes, where one of them is a trivariate process. GV generate only in-sample forecasts, whereas we focus on the out-of-sample performance. GV include an autoregressive lag in the MIDAS regressions. We focus on the basic specification. In addition to the (infeasible) high-frequency VAR, the low-frequency VAR and MIDAS regressions, we include the mixed-frequency VAR in the forecast comparisons. GV employ as time series lengths for lower frequency $T = 500$ and $T = 1000$. We use $T = 100$ and $T = 500$. The first choice represents typical time series lengths in applied forecasting. For instance, $T = 100$ represents 100 quarters (25 years). The longer time series length ($T = 500$) is chosen to investigate by how far the forecast performance improves. GV simulate only homoscedastic errors, whereas we also allow for GARCH effects in the errors. GV consider the following frequency mixtures: $m = \{5, 10, 20, 60, 120, 250\}$. For reasons of brevity we only consider $m = \{3, 4, 12\}$. The focus of this thesis is also on typical macroeconomic mixtures with yearly, quarterly, and monthly data. Furthermore feasible estimates for higher-frequency mixtures (such as daily-quarterly) are difficult to obtain for the mixed-frequency VAR. In line with GV, we simulate each process 1,000 times. Furthermore we also augment the information set to allow for intra-period information.

In sum, we compare the forecasting performance of five time series models. The infeasible high-frequency VAR (HF-VAR), the low-frequency VAR (LF-VAR), the MIDAS model with both exponential Almon and Beta weighting function, and the mixed-frequency VAR (MF-VAR) within the framework provided by Zadrozny (1990). For process III we substitute the LF-VAR with an AR benchmark model. We forecast the lower-frequency variable one-step-ahead. Thus we focus on short-term forecasting. In our Monte Carlo study it is the final observation T. We use all information available up to $T-1$ to generate the forecast, that is $mT-m$ high-frequency observations. This notation is valid for the MIDAS models and the LF-VAR where no auxiliary forecasts are necessary to obtain the desired forecast (direct forecasting approach). For the HF-VAR and the MF-VAR the final forecast value is denoted by mT and we use information up to $mT - m$, but the forecasting horizon is different. For MIDAS and LF-VAR the horizon is $h = 1$; for the HF-VAR and MF-VAR the horizon is $h = m$, as we have to iterate up to the desired forecast (iterated forecasting approach). To avoid confusion, all forecasts use the same information set. As the forecasting evaluation criterion, we employ the Root Mean Squared Error (RMSE)

$$RMSE_{iT} = \frac{1}{1000} \sum \sqrt{(y_T - \hat{y}_T)^2} \tag{6.4}$$

$$RMSE_{iT} = \frac{1}{1000} \sum \sqrt{(y_T - \hat{y}_{mT})^2}, \tag{6.5}$$

where i denotes the model. The first definition is the notation for LF-VAR and the MIDAS approach, whereas the second is the RMSE notation for the other models.

We note that in the case of the HF-VAR and MF-VAR our notation deviates from the standard definition of the RMSE. To obtain the desired forecast for the actual value in T we have to iterate the forecast up to mT starting in $mT - m$. Thus, in the case of monthly and quarterly data, we generate two auxiliary forecasts to get the desired forecast. The standard RMSE in this case is the average of the forecasting errors up to the forecast horizon. Applying the standard RMSE criterion would introduce bias into our comparisons for two reasons. First, the number of generated forecasts is different; second, the longer the forecasting horizon, the higher is the average forecasting error. This should be borne in mind for interpretation of the forthcoming results.

In a second step we augment the information set to allow for intra-period information. Thus the high-frequency information set contains observations up to $mT - 1$. *A priori* this should improve the forecasting accuracy, as shown by Ghysels and Valkanov (2006). We also outlined in chapter 3 many examples where accounting for intra-period information reduces forecasting errors. The augmentation of the information set also mimics real-time situations, where actual information is processed to update the current estimate of GDP (nowcasting). This situation is often found at central banks where monetary policy is conducted in real time. But as we have outlined in chapter 3, bridge equations are still the dominant approach for nowcasting at central banks.

6.3 Process I

Given our findings in chapter 5 we can specify how to estimate the mixed-frequency time series models. In the MIDAS framework we select the number of included lags with the BIC. The range of tested lags is $4m \pm 2m$. We restrict the weighting function only for $\rho \geq 0.9$, as given in chapter 4. In the case of the mixed-frequency VAR, we allow for a maximum of three lags. For lag selection we use the corrected AIC criterion as stated in equation (5.1). In the case of the HF-VAR we generate the forecasts with a VAR(1) process, as the true data generating process is a VAR(1).[1] For the LF-VAR we allow for a maximum of $2m$ lags. We take the BIC to choose the optimal lag.[2]

The results for Process I are tabulated in ten tables which are all contained in an appendix to this chapter. We will focus on cases where $\rho > 0.5$ as these exhibit significant autocorrelations as outlined in chapter 4. We start with the comparison to the infeasible high-frequency VAR benchmark in Table 6.1. A ratio smaller than 1 indicates that the HF-VAR exibits a higher RMSE compared to its competitor. In general, for $m = 3$ and $m = 4$ no model outperforms the benchmark. For $\rho = 0.95$ and $m = \{3, 4\}$ the HF-VAR clearly dominates its competitors. This result is in line with Ghysels and Valkanov (2006) who also find ratios larger than 1 although for in-sample comparisons. The superiority of the HF-VAR holds for the LF-VAR and the MF-VAR for $m = 12$ but not for the MIDAS models. For larger persistent

[1] We experimented with lag selection criteria but in almost all cases the VAR(1) was chosen.

[2] The results indicate that in most cases a VAR(1) is selected.

series ($\rho \geq 0.9$) both MIDAS models outperform the HF-VAR by 10 to 20 %. This finding is due to the calculation of forecasts. For $m = 12$ the HF-VAR generates 11 auxiliary forecasts to obtain the desired forecasts, which increases the forecast uncertainty. In contrast, in case of MIDAS the forecast is calculated directly. Therefore the accumulated forecast error of the HF-VAR is higher than the direct forecast error of the MIDAS models.

The more interesting part is the comparison of the feasible models in reality. Therefore we report the ratios of the mixed-frequency models relative to the low-frequency VAR. The results are depicted in Table 6.2. Note, that in this table the same results are used as in Table 6.1. A ratio below 1 indicates that model i obtains on average lower RMSEs than the LF-VAR. The results are heterogenous. For $m = 3$ only the MIDAS models outperform the benchmark, but not for strongly persistent series ($\rho = 0.95$), whereas the MF-VAR exhibits larger RMSEs in each case for this frequency mixture. In this scenario, the temporal-aggregated data seems to entail enough information for forecasting. This argument does not hold for $m = 4$ where the MIDAS models outperform the LF-VAR also for $\rho = 0.95$ (by 5 to 15%) but not as clearly as for $\rho = 0.9$ (up to 55 %). The MF-VAR is only better for $\rho = 0.9$. The most striking results we find for the largest frequency mixture $m = 12$. All mixed-frequency models outperform the LF-VAR up to 40%. Thus, the use of untransformed mixed-frequency data clearly increases forecasting accuracy. Comparing MIDAS and MF-VAR the former one is better choice. Concerning the two weighting function, the exponential Almon lag provides constantly lower RMSEs but on a small scale level (between 1 and 5 percentage points).

It is well known that increasing the estimation sample reduces estimation uncertainty and therefore fosters forecasting accuracy. We increase the lower-frequency time series length from 100 to 500. Table 6.3 displays the ratio of the RMSEs obtained from $T = 500$ relative to $T = 100$ for each model separately. A ratio below 1 indicates that the forecast errors are reduced by an increased sample size. As expected *a priori*, we find that in almost all cases forecasts improve. For higher frequency mixtures the highest gain in forecasting accuracy is detected for the LF-VAR. The MF-VAR reduces the average RMSE relatively more than the MIDAS competitors.

Now we go a step further and model GARCH errors in the high-frequency DGP. Table 6.4 reports the ratio relative to the HF-VAR. In contrast to the homoscedastic case, the HF-VAR is almost never outperformed by the mixed-frequency models. Comparing the ratios relative to the LF-VAR (Table 6.5), for $\rho = 0.9$, the benchmark is always outperformed. For $m = 12$ we find similar results as in Table 6.1. Both mixed-frequency approaches outperform the LF-VAR for $\rho \geq 0.9$. Again we find that the exponential Almon lag performs better than the Beta weighting function. Increasing the sample size to 500 does not always improve the forecasting accuracy in any case as reported in Table 6.6. For $m = 12$ we find almost no improvement for the mixed-frequency models. In general we can summarize that there seems to be a sensitivity to heteroscedastic data which may reduce forecast accuracy.

In a final step we augment the information set and include information up to $mT - 1$.[3] Table 6.7 reports the ratio relative to the LF-VAR, which includes information up to $mT - m$. The results are striking. The average forecast errors are reduced up to 90% compared to the LF-VAR. For $m = 3$ and $m = 4$ the MF-VAR clearly outperforms the MIDAS models. This is due to the true DGP. Our MIDAS approach does not contain an autoregressive lag which seems to be the striking advantage for this data structure. This may also explain the puzzling results for $m = 3$ and $\rho = 0.95$ where even the LF-VAR outperforms the MIDAS models.

Increasing the sample size reduces again the forecast errors. All ratios are lower than 1 in Table 6.8.

We also augment the information set for the GARCH case. The results are essentially the same as for the homoscedastic case. Again, with the MF-VAR we obtained the lowest RMSEs as reported in Table 6.9. And, for $\rho = 0.95$ and lower frequency mixtures the MIDAS models do not outperform the LF-VAR benchmark. Increasing the sample size increases again the forecasting

[3] To give an example, $mT - 1$ corresponds to the February value to forecast the March value of the quarterly target variable.

accuracy but not as much as in the homoscedastic case. Table 6.10 reports the corresponding results.

6.4 Process II

For Process II we apply the same model selection criteria as for Process I with two exceptions. First, we do not restrict the weighting functions in any case. Second, the forecasts for the HF-VAR are generated with a VAR(2) process. The tables can be found in the appendix to this chapter. We proceed in the same way as for Process I.

We start by comparing the different approaches relative to the infeasible HF-VAR. Table 6.11 reports the results. The high-frequency benchmark model is not outperformed considerably in any case . This stands in contrast to our findings for Process I (Table 6.1), where for $m = 12$ the MIDAS models obtained lower RMSEs in comparison to the HF-VAR.

Similar results are found in Table 6.12, where we report the ratios relative to the LF-VAR benchmark. Only for $m = 12$ and $\rho \geq 0.5$ the mixed-frequency approaches outperform the single-frequency benchmark. In Table 6.12 the Beta weighting function exhibits lower RMSEs in comparison to the exponential Almon lag.

In Table 6.13 we increase the sample size of the lower frequency from $T = 100$ to $T = 500$. Compared to Process I we find only small improvements up to 10%. In some cases there is a slight deterioration of the forecasting accuracy.

As in case of Process I we simulate GARCH in the errors of the DGP. Table 6.14 displays the corresponding ratios relative to the infeasible HF-VAR. As expected, the benchmark is not outperformed, but the ratios are relative close to 1. Using the LF-VAR as the benchmark (Table 6.15) the mixed-frequency approaches are only better for $m = 12$. As in the homoscedastic case, the increase of the sample size leads to small improvements of the forecasting accuracy (Table 6.16).

Finally we augment the information set. Table 6.17 reports the results for homoscedastic errors. Generally we can state, that the higher m and the higher ρ the lower is the average RMSE of the mixed-frequency models. In contrast to Process I the autoregressive part does not drive the results as we find some cases where the MIDAS is better than the MF-VAR. Similar conclusions can be drawn by allowing for GARCH errors. In this case the Beta weighting function delivers no ratios smaller than 1 for $m = 3, 4$ but is better than the exponential Almon lag for $m = 12$. The results are reported in Table 6.19.

Increasing the sample size reduces the forecasting errors in both cases with a lower extend in the GARCH case. The results are depicted in Tables 6.18 and 6.20, respectively.

6.5 Process III

For Process III we cannot estimate a VAR(p) benchmark model, as the data generating process is a MIDAS model. In this case we estimate an AR(p) for the lower-frequency time series as the benchmark model. The number of included lags is $3m \pm 2m$ and we do not restrict the weighting functions. In case of the MV-VAR model we allow again for a maximum of three lags.

Table 6.21 (Panel A) reports the ratio of the mixed-frequency models compared to the AR benchmark. A ratio below one indicates that the AR model exhibits a higher RMSE than its competitor. We obtain clear cut results. All mixed-frequency models outperform the AR benchmark clearly. For larger mixtures the AR gets better but the RMSE a still higher by at least 20% compared to the mixed-frequency models. For lower-frequency mixtures both weighting functions exhibit similar average RMSEs, but for $m = 12$ the Beta weighting function clearly outperforms the exponential Almon lag. Comparing MIDAS and MF-VAR we find that MIDAS is on average better than the MF-VAR. This result is not surprising as the DGP is a MIDAS model. Panel B reports the results for an increase of the estimation sample. We find rather small improvements (up to 12%). This result is important for applied forecasting. It states that we obtain clear cut forecasting results with standard time series lengths. Or to put it differently, increasing the sample size would not result in significant improvements. This conclusion draws on the similarity of the autocorrelation structures of US and German GDP (Figure 4.1) and Process III (Figure 4.9).

The augmentation of the information set improves the forecasting accuracy markedly. In case of the MIDAS models we find a reduction of the RMSE in comparison to the AR benchmark up to 95% (Table 6.22, Panel A). The Beta weighting function constantly outperforms the exponential Almon lag. The reductions for the MF-VAR are rather moderate. Panel B in Table 6.22 reports that the increase of the sample size leads to comparably better reduction in the RMSEs than the standard information set.

6.6 Process IV: Two Indicators at Different Frequencies

For Process IV with two leading indicators we did not conduct a specification investigation as we did for the other processes in chapter 5. Given the results for Processes I and II, we allow for a maximum of $4m \pm 2m$ lags; we do not restrict the weighting function, as the autocorrelation function shows a low persistent pattern. In case of the MV-VAR model, we now allow for a maximum lag of five. The HF-VAR forecasts are generated with an VAR(1) model.

For Process IV we find clear-cut results. The HF-VAR remains infeasible (Table 6.23, Panel A), but compared with the LF-VAR, the mixed-frequency approaches are better than the LF-VAR (Panel B). In contrast to *a priori* expectations the increase of the estimation sample lowers the forecasting accuracy for the mixed-frequency approaches (Panel C). Augmenting the information set, the RMSEs decrease substantially. The MIDAS models are up to 70% and the MF-VAR up to 35% better than the LF-VAR (Panel D). There are no improvement by increasing the estimation sample (Panel E).

6.7 Summary

We summarize the results of our Monte Carlo study in the following key points:

- Mixed-frequency time series models are at least as good as single-frequency time series models.
- In case of data with significant autocorrelation structures, one of the two mixed-frequency approaches outperforms the temporal aggregated benchmark model (in some cases up to 50%).
- For time series with a strong autoregressive component, the MF-VAR outperforms MIDAS.
- For strong persistent time series and low-frequency mixtures, temporal-aggregation is an option for forecasting.
- Augmenting the information set by allowing for intra-period information increases forecasting accuracy in almost all cases. The improvement in comparison to temporal-aggregated and single-frequency models is substantial. A reduction of forecast errors up to 95% is feasible.
- The choice for the exponential Almon lag or Beta weighting function remains undecided. In many cases both deliver the same results; thus they can be used interchangeably for forecasting purposes. For short-term forecasting (allowing for intra-period information) the Beta functions obtains on average lower forecasting errors.
- For persistent series and low-frequency mixtures the augmentation of the information set does not necessarily improve forecasting accuracy in the MIDAS case.
- Using more than two different frequencies for forecasting the mixed-frequency time series models is feasible and demonstrates remarkable forecasting improvements in comparison with temporal aggregated benchmark models.

- Increasing the sample size reduces the average forecast errors in many cases. But we have also find structures where no or only small improvements occur.

- Our results indicate that for standard macroeconomic time series the generated forecasts would not be more accurate, on average, than very long macroeconomic time series.

- GARCH effects in macroeconomic time series do affect the forecasting performance of the mixed-frequency approaches. The relative gain is smaller in comparison to the homoscedastic case. Nevertheless, there may be room for improvement by adjusting the models for GARCH dynamics.

6.8 Appendix: Tables

Tab. 6.1: RMSE Ratios vs. HF-VAR: Process I with homoscedastic errors ($T = 100$)

		$m = 3$				$m = 4$				$m = 12$			
		LF-VAR	Almon	Beta	MF-VAR	LF-VAR	Almon	Beta	MF-VAR	LF-VAR	Almon	Beta	MF-VAR
$\rho = 0.1$	$\delta = 0$	1.052	0.967	0.956	1.055	0.965	1.463	1.451	1.596	1.093	1.231	1.301	1.437
	$\delta = -0.5$	1.034	1.005	0.997	1.098	1.306	1.357	1.321	1.452	1.261	1.286	1.348	1.485
	$\delta = -1.5$	1.033	1.009	1.025	1.130	1.354	1.340	1.372	1.509	1.145	1.205	1.255	1.384
	$\delta = -3.5$	1.073	1.064	1.119	1.222	1.379	1.431	1.415	1.556	1.176	1.182	1.230	1.361
$\rho = 0.5$	$\delta = 0$	1.027	1.055	1.052	1.151	1.352	1.355	1.333	1.472	1.461	1.297	1.300	1.438
	$\delta = -0.5$	1.170	0.992	1.003	1.103	1.267	1.455	1.487	1.631	1.317	1.194	1.273	1.404
	$\delta = -1.5$	1.061	0.982	0.944	1.049	1.369	1.402	1.379	1.515	1.048	1.080	1.105	1.215
	$\delta = -3.5$	1.009	0.919	0.852	0.950	1.143	1.272	1.241	1.369	0.903	0.985	1.025	1.120
$\rho = 0.9$	$\delta = 0$	1.206	1.166	1.144	1.264	1.336	1.437	1.425	1.568	1.060	0.992	1.025	1.131
	$\delta = -0.5$	1.410	1.188	1.200	1.474	2.001	1.149	1.213	1.479	1.220	0.912	0.960	1.087
	$\delta = -1.5$	1.612	1.283	1.326	1.708	2.258	1.116	1.195	1.536	1.061	0.740	0.775	1.038
	$\delta = -3.5$	1.595	1.158	1.235	1.626	1.851	0.837	0.943	1.232	1.281	0.868	0.858	1.070
$\rho = 0.95$	$\delta = 0$	1.162	1.635	1.643	1.796	1.412	1.690	1.671	1.847	0.990	1.053	1.051	1.163
	$\delta = -0.5$	1.626	3.256	3.266	3.838	2.276	1.911	1.907	2.395	1.331	0.871	0.877	1.055
	$\delta = -1.5$	2.120	4.168	4.199	5.002	2.607	2.257	2.379	2.972	1.307	0.858	0.875	1.066
	$\delta = -3.5$	1.626	3.611	3.626	4.338	2.483	2.334	2.333	2.953	1.568	0.893	0.934	1.138

Notes: Data is simulated from a high-frequency VAR as given in equation (4.1) with homoscedastic errors. The time series length is $T = 100$ (lower-frequency). Using the high frequency data, we estimate a VAR(p) and conduct an out-of-sample one-step ahead forecast. The root mean squared error of this forecast, $RMSE(1, \mathcal{I}_t)$, is used as a benchmark. The competitors are the low-frequency VAR(p), where the higher-frequency variable is aggregated to the lower frequency by taking averages over the corresponding low-frequency period ($RMSE(1, \mathcal{I}_t^A)$), the MIDAS model with exponential Almon lag and Beta lag weighting function and the mixed-frequency VAR ($RMSE_{\{i\}}(1, \mathcal{I}_t^M), i = \{Almon, Beta, MF-VAR\}$). The table reports the ratio $RMSE(1, \mathcal{I}_t^A)/RMSE(1, \mathcal{I}_t)$ and $RMSE_{\{i\}}(1, \mathcal{I}_t^M)/RMSE(1, \mathcal{I}_t)$ for frequency mixtures $m = 3, 4, 12$.

Tab. 6.2: RMSE Ratios vs. LF-VAR: Process I with homoscedastic errors ($T = 100$)

		$m = 3$			$m = 4$			$m = 12$		
		Almon	Beta	MF-VAR	Almon	Beta	MF-VAR	Almon	Beta	MF-VAR
$\rho = 0.1$	$\delta = 0$	0.920	0.909	1.003	1.515	1.503	1.653	1.127	1.191	1.315
	$\delta = -0.5$	0.973	0.964	1.062	1.039	1.012	1.112	1.020	1.069	1.178
	$\delta = -1.5$	0.976	0.993	1.093	0.990	1.013	1.114	1.053	1.096	1.209
	$\delta = -3.5$	0.991	1.042	1.138	1.038	1.026	1.128	1.005	1.046	1.158
$\rho = 0.5$	$\delta = 0$	1.027	1.024	1.121	1.002	0.986	1.089	0.887	0.890	0.984
	$\delta = -0.5$	0.848	0.858	0.943	1.148	1.174	1.287	0.906	0.966	1.066
	$\delta = -1.5$	0.926	0.890	0.988	1.024	1.007	1.107	1.031	1.054	1.160
	$\delta = -3.5$	0.911	0.845	0.942	1.113	1.085	1.198	1.090	1.135	1.239
$\rho = 0.9$	$\delta = 0$	0.967	0.949	1.048	1.075	1.067	1.173	0.936	0.967	1.067
	$\delta = -0.5$	0.842	0.851	1.045	0.574	0.606	0.739	0.748	0.787	0.891
	$\delta = -1.5$	0.796	0.822	1.059	0.494	0.529	0.680	0.698	0.731	0.829
	$\delta = -3.5$	0.726	0.774	1.019	0.452	0.510	0.666	0.678	0.669	0.767
$\rho = 0.95$	$\delta = 0$	1.406	1.413	1.545	1.197	1.184	1.308	1.064	1.061	1.174
	$\delta = -0.5$	2.003	2.009	2.361	0.840	0.838	1.052	0.654	0.659	0.793
	$\delta = -1.5$	1.966	1.980	2.359	0.866	0.913	1.140	0.656	0.669	0.816
	$\delta = -3.5$	2.221	2.231	2.668	0.940	0.940	1.189	0.570	0.595	0.726

Notes: This table builds upon the same data as in Table 6.1 but the benchmark is the LF-VAR. The table reports the ratio $RMSE_{\{i\}}(1, \mathcal{I}_t^M)/RMSE(1, \mathcal{I}_t^A)$ for frequency mixtures $m = 3, 4, 12$.

Tab. 6.3: RMSE Ratios $T = 100$ vs. $T = 500$: Process I with homoscedastic errors

		$m = 3$				$m = 4$				$m = 12$			
		LF-VAR	Almon	Beta	MF-VAR	LF-VAR	Almon	Beta	MF-VAR	LF-VAR	Almon	Beta	MF-VAR
$\rho = 0.1$	$\delta = 0$	0.881	1.000	0.995	0.933	1.283	0.894	0.906	0.854	0.679	1.019	1.023	0.958
	$\delta = -0.5$	0.906	0.915	0.913	0.860	0.948	0.949	0.980	0.922	0.582	0.998	0.972	0.913
	$\delta = -1.5$	0.965	0.977	0.964	0.906	0.960	1.033	0.990	0.931	0.625	1.079	1.061	0.993
	$\delta = -3.5$	0.908	0.909	0.842	0.799	0.923	0.967	0.934	0.879	0.587	1.089	1.106	1.032
$\rho = 0.5$	$\delta = 0$	0.921	0.911	0.930	0.879	0.970	1.024	1.022	0.956	0.490	1.021	1.029	0.962
	$\delta = -0.5$	0.838	0.965	0.967	0.910	1.012	0.932	0.929	0.876	0.538	1.068	0.987	0.926
	$\delta = -1.5$	1.024	1.021	1.001	0.934	0.930	0.965	0.907	0.855	0.606	1.030	1.036	0.972
	$\delta = -3.5$	0.959	0.972	0.974	0.909	1.053	0.996	0.972	0.912	0.547	0.899	0.862	0.817
$\rho = 0.9$	$\delta = 0$	0.836	1.009	1.008	0.944	0.858	1.046	0.979	0.920	0.545	1.013	1.033	0.967
	$\delta = -0.5$	0.916	0.903	0.964	0.847	0.731	0.958	0.927	0.820	0.601	1.002	0.984	0.908
	$\delta = -1.5$	0.934	0.766	0.854	0.738	0.767	0.812	0.855	0.739	0.513	0.960	0.941	0.869
	$\delta = -3.5$	0.883	0.767	0.780	0.672	0.770	0.895	0.884	0.755	0.549	0.954	0.985	0.902
$\rho = 0.95$	$\delta = 0$	0.901	0.899	0.899	0.851	0.885	1.009	1.023	0.957	0.622	0.968	0.970	0.909
	$\delta = -0.5$	0.810	0.635	0.632	0.588	0.626	0.809	0.828	0.728	0.628	0.867	0.874	0.782
	$\delta = -1.5$	0.738	0.653	0.648	0.598	0.617	0.769	0.759	0.674	0.699	0.887	0.886	0.787
	$\delta = -3.5$	0.849	0.690	0.686	0.628	0.620	0.731	0.771	0.679	0.581	0.967	0.940	0.831

Notes: Data is simulated from a high-frequency VAR as given in equation (4.1) with homoscedastic errors. The time series length is $T = 500$ (lower-frequency). The table reports the ratio $RMSE_{\{i\}}(500)/RMSE_{\{i\}}(100)$ for each model i for frequency mixtures $m = 3, 4, 12$.

Tab. 6.4: RMSE Ratios vs. HF-VAR: Process I with GARCH errors ($T = 100$)

		$m=3$				$m=4$				$m=12$			
		LF-VAR	Almon	Beta	MF-VAR	LF-VAR	Almon	Beta	MF-VAR	LF-VAR	Almon	Beta	MF-VAR
$\rho = 0.1$	$\delta = 0$	1.041	0.990	0.992	1.046	1.090	1.352	1.351	1.072	1.120	1.121	1.189	1.046
	$\delta = -0.5$	1.017	1.022	0.987	1.035	1.098	1.345	1.316	1.025	1.139	1.148	1.217	1.136
	$\delta = -1.5$	1.014	1.039	1.044	1.013	1.097	1.353	1.376	1.001	1.079	1.127	1.200	1.052
	$\delta = -3.5$	1.025	1.058	1.016	1.152	1.104	1.417	1.337	1.066	1.079	1.105	1.140	1.097
$\rho = 0.5$	$\delta = 0$	0.974	1.021	0.981	1.167	1.088	1.268	1.252	1.047	1.132	1.157	1.211	1.123
	$\delta = -0.5$	1.032	1.025	1.014	1.109	1.102	1.346	1.293	1.068	1.103	1.160	1.178	1.035
	$\delta = -1.5$	1.091	1.033	1.046	1.461	1.119	1.294	1.271	1.034	1.079	1.110	1.215	1.085
	$\delta = -3.5$	1.062	0.955	0.885	1.157	1.072	1.295	1.220	1.068	1.077	1.036	1.107	1.389
$\rho = 0.9$	$\delta = 0$	1.134	1.275	1.257	1.361	1.208	1.476	1.447	1.028	1.072	1.005	1.100	1.362
	$\delta = -0.5$	1.471	1.215	1.293	1.555	1.619	1.317	1.332	1.410	1.409	1.006	1.065	1.322
	$\delta = -1.5$	1.852	1.335	1.494	1.498	1.949	1.117	1.187	1.535	1.229	0.922	1.063	0.961
	$\delta = -3.5$	1.689	1.190	1.328	1.400	1.680	0.988	1.079	1.614	1.364	0.974	1.017	1.223
$\rho = 0.95$	$\delta = 0$	1.049	1.578	1.546	1.369	1.311	1.636	1.651	1.136	1.146	1.059	1.084	1.176
	$\delta = -0.5$	1.623	2.682	2.698	2.227	1.770	2.009	1.994	2.061	1.721	1.002	1.108	1.410
	$\delta = -1.5$	2.192	3.627	3.629	2.925	1.974	2.216	2.217	2.419	1.745	1.055	1.215	1.596
	$\delta = -3.5$	1.992	3.060	3.055	2.122	1.951	2.231	2.305	2.392	1.728	1.063	1.230	1.296

Notes: Data is simulated from a high-frequency VAR as given in equation (4.1) with GARCH errors (see equation (4.3)). The time series length is $T = 100$ (lower-frequency). Using the high frequency data, we estimate a VAR(p) and conduct an out-of-sample one-step ahead forecast. The root mean squared error of this forecast, $RMSE(1, \mathcal{I}_t)$, is used as a benchmark. The competitors are the low-frequency VAR(p), where the higher-frequency variable is aggregated to the lower frequency by taking averages over the corresponding low-frequency period ($RMSE(1, \mathcal{I}_t^A)$), the MIDAS model with exponential Almon lag and Beta lag weighting function and the mixed-frequency VAR ($RMSE_{\{i\}}(1, \mathcal{I}_t^M), i = \{Almon, Beta, MF-VAR\}$). The table reports the ratio $RMSE(1, \mathcal{I}_t^A)/RMSE(1, \mathcal{I}_t)$ and $RMSE_{\{i\}}(1, \mathcal{I}_t^M)/RMSE(1, \mathcal{I}_t)$ for frequency mixtures $m = 3, 4, 12$.

Tab. 6.5: RMSE Ratios vs. LF-VAR: Process I with GARCH errors ($T = 100$)

		$m=3$			$m=4$			$m=12$		
		Almon	Beta	MF-VAR	Almon	Beta	MF-VAR	Almon	Beta	MF-VAR
$\rho = 0.1$	$\delta = 0$	0.952	0.953	1.005	1.240	1.240	0.984	1.001	1.062	0.934
	$\delta = -0.5$	1.004	0.971	0.906	1.225	1.198	0.933	1.008	1.069	0.997
	$\delta = -1.5$	1.024	1.029	0.998	1.234	1.255	0.912	1.045	1.112	0.975
	$\delta = -3.5$	1.031	0.990	1.123	1.284	1.211	0.966	1.024	1.057	1.017
$\rho = 0.5$	$\delta = 0$	1.048	1.006	1.197	1.165	1.150	0.962	1.022	1.070	0.992
	$\delta = -0.5$	0.992	0.982	1.074	1.221	1.174	0.899	1.052	1.068	0.835
	$\delta = -1.5$	0.947	0.959	1.339	1.156	1.136	0.924	1.029	1.126	0.910
	$\delta = -3.5$	0.899	0.833	1.089	1.208	1.139	0.843	0.961	1.027	1.289
$\rho = 0.9$	$\delta = 0$	1.124	1.108	1.201	1.222	1.198	0.851	0.937	1.026	1.270
	$\delta = -0.5$	0.826	0.879	1.057	0.814	0.823	0.871	0.714	0.756	0.938
	$\delta = -1.5$	0.721	0.807	0.809	0.573	0.609	0.788	0.750	0.864	0.782
	$\delta = -3.5$	0.705	0.786	0.829	0.588	0.642	0.961	0.714	0.745	0.897
$\rho = 0.95$	$\delta = 0$	1.504	1.473	1.305	1.248	1.259	0.866	0.924	0.946	1.027
	$\delta = -0.5$	1.652	1.662	1.372	1.135	1.127	1.165	0.582	0.644	0.819
	$\delta = -1.5$	1.655	1.656	1.335	1.123	1.123	1.226	0.604	0.696	0.914
	$\delta = -3.5$	1.536	1.534	1.065	1.144	1.182	1.226	0.615	0.712	0.750

Notes: This table builds upon the same data as in Table 6.4 but the benchmark is the LF-VAR. The table reports the ratio $RMSE_{\{i\}}(1, \mathcal{I}_t^M)/RMSE(1, \mathcal{I}_t^A)$ for frequency mixtures $m = 3, 4, 12$.

Tab. 6.6: RMSE Ratios $T = 100$ vs. $T = 500$: Process I with GARCH errors

		$m=3$				$m=4$				$m=12$			
		LF-VAR	Almon	Beta	MF-VAR	LF-VAR	Almon	Beta	MF-VAR	LF-VAR	Almon	Beta	MF-VAR
$\rho = 0.1$	$\delta = 0$	0.939	0.986	1.006	0.944	1.307	0.925	0.956	0.899	1.032	0.941	0.929	0.872
	$\delta = -0.5$	0.930	0.953	0.933	0.878	0.823	0.856	0.902	0.851	0.908	0.950	0.922	0.868
	$\delta = -1.5$	1.036	1.037	1.034	0.969	0.864	0.935	0.909	0.857	1.030	0.998	0.977	0.917
	$\delta = -3.5$	0.921	0.916	0.866	0.821	0.912	0.913	0.924	0.870	0.943	0.993	0.953	0.894
$\rho = 0.5$	$\delta = 0$	0.934	0.913	0.919	0.899	0.902	0.923	0.952	0.893	0.820	0.985	0.963	0.902
	$\delta = -0.5$	0.826	0.983	0.978	0.884	1.022	0.908	0.891	0.842	0.819	0.960	0.912	0.858
	$\delta = -1.5$	0.942	0.940	0.951	1.260	0.876	0.897	0.857	0.810	0.967	0.968	0.974	0.915
	$\delta = -3.5$	1.090	0.976	1.002	1.081	1.023	0.902	0.899	0.846	1.152	1.201	1.135	1.067
$\rho = 0.9$	$\delta = 0$	0.914	0.989	0.987	0.951	0.880	0.885	0.919	0.866	1.013	1.161	1.092	1.021
	$\delta = -0.5$	0.923	1.055	1.043	0.807	0.758	0.959	0.927	0.820	1.093	1.127	1.086	1.061
	$\delta = -1.5$	1.016	0.937	0.913	0.579	0.744	0.930	0.922	0.791	1.006	1.053	1.040	0.936
	$\delta = -3.5$	1.021	0.953	0.930	0.540	0.833	0.997	0.919	0.782	1.177	1.160	1.200	1.075
$\rho = 0.95$	$\delta = 0$	0.825	0.904	0.906	0.649	0.966	0.995	1.024	0.636	1.172	0.922	0.943	0.884
	$\delta = -0.5$	0.891	0.736	0.720	0.662	0.692	0.906	0.896	0.713	1.110	1.052	1.053	1.011
	$\delta = -1.5$	0.829	0.740	0.719	0.657	0.661	0.880	0.827	0.662	1.157	1.068	1.113	1.025
	$\delta = -3.5$	0.881	0.802	0.783	0.654	0.717	0.918	0.887	0.700	0.929	1.058	1.057	0.900

Notes: Data is simulated from a high-frequency VAR as given in equation (4.1) with GARCH errors (see equation (4.3)). The time series length is $T = 500$ (lower-frequency). The table reports the ratio $RMSE_{\{i\}}(500)/RMSE_{\{i\}}(100)$ for each model i for frequency mixtures $m = 3, 4, 12$.

Tab. 6.7: RMSE Ratios vs. LF-VAR: Process I with homoscedastic errors - Augmented information set ($T = 100$)

		$m = 3$			$m = 4$			$m = 12$		
		Almon	Beta	MF-VAR	Almon	Beta	MF-VAR	Almon	Beta	MF-VAR
$\rho = 0.1$	$\delta = 0$	1.145	1.132	1.014	1.366	1.411	0.631	1.155	1.195	0.696
	$\delta = -0.5$	0.973	0.961	0.921	0.944	0.981	0.727	1.026	1.119	0.569
	$\delta = -1.5$	0.991	0.981	1.013	0.953	0.969	0.772	1.044	1.051	0.604
	$\delta = -3.5$	0.975	0.955	0.970	0.956	0.963	0.713	0.997	1.082	0.630
$\rho = 0.5$	$\delta = 0$	1.035	1.024	0.941	0.924	0.963	0.732	0.864	0.900	0.744
	$\delta = -0.5$	0.844	0.833	0.930	1.047	1.010	0.797	0.882	0.883	0.601
	$\delta = -1.5$	0.677	0.639	0.740	0.860	0.669	0.528	0.933	0.737	0.555
	$\delta = -3.5$	0.430	0.372	0.517	0.860	0.447	0.459	0.788	0.383	0.361
$\rho = 0.9$	$\delta = 0$	0.845	0.910	0.760	0.991	1.024	0.631	0.892	0.873	0.706
	$\delta = -0.5$	0.792	0.809	0.473	0.598	0.508	0.442	0.427	0.345	0.315
	$\delta = -1.5$	0.720	0.721	0.287	0.507	0.386	0.326	0.243	0.137	0.281
	$\delta = -3.5$	0.706	0.707	0.102	0.473	0.341	0.206	0.197	0.087	0.128
$\rho = 0.95$	$\delta = 0$	1.284	1.365	0.806	1.128	1.154	0.430	0.950	0.939	0.524
	$\delta = -0.5$	1.911	1.959	0.380	1.015	0.862	0.404	0.406	0.280	0.383
	$\delta = -1.5$	1.933	1.971	0.192	1.086	0.918	0.279	0.330	0.216	0.196
	$\delta = -3.5$	2.270	2.318	0.219	1.160	0.976	0.300	0.286	0.177	0.087

Notes: Data is simulated from a high-frequency VAR as given in equation (4.1) with homoscedastic errors. The time series length is $T = 100$ (lower-frequency). Using the low-frequency data (temporally aggregated), we estimate a VAR(p) and conduct an out-of-sample one-step ahead forecast. The root mean squared error of this forecast, $RMSE(1, \mathcal{I}_t^A)$, is used as a benchmark. The competitors are the MIDAS model with exponential Almon lag and Beta lag weighting function and the mixed-frequency VAR ($RMSE_{\{i\}}(1, \mathcal{I}_t^M), i = \{Almon, Beta, MF-VAR\}$). The information set is augmented and includes information up to $mT - 1$. The table reports the ratio $RMSE_{\{i\}}(1, \mathcal{I}_t^M)/RMSE(1, \mathcal{I}_t^A)$ for frequency mixtures $m = 3, 4, 12$.

Tab. 6.8: RMSE Ratios $T = 100$ vs. $T = 500$: Process I with homoscedastic errors - Augmented information set

		$m = 3$				$m = 4$				$m = 12$			
		LF-VAR	Almon	Beta	MF-VAR	LF-VAR	Almon	Beta	MF-VAR	LF-VAR	Almon	Beta	MF-VAR
$\rho = 0.1$	$\delta = 0$	1.134	1.023	1.093	0.994	1.283	0.943	1.008	0.927	0.679	1.045	0.884	0.944
	$\delta = -0.5$	0.906	0.921	0.881	0.979	0.948	1.003	1.015	0.972	0.582	1.029	0.839	0.999
	$\delta = -1.5$	0.965	0.969	0.999	0.974	0.960	0.967	1.143	0.979	0.625	1.085	0.998	0.983
	$\delta = -3.5$	0.908	0.898	0.966	0.988	0.923	0.929	1.014	0.960	0.587	1.092	0.931	0.950
$\rho = 0.5$	$\delta = 0$	0.921	0.908	0.930	0.974	0.970	1.029	1.057	0.976	0.490	1.047	0.889	0.998
	$\delta = -0.5$	0.838	0.964	0.916	0.973	1.012	0.959	1.023	0.995	0.538	1.050	0.945	0.953
	$\delta = -1.5$	1.024	0.982	0.968	0.986	0.930	0.944	0.993	0.942	0.606	1.010	0.894	0.983
	$\delta = -3.5$	0.959	0.953	0.932	0.969	1.053	1.004	1.132	0.930	0.547	1.022	0.885	0.934
$\rho = 0.9$	$\delta = 0$	0.836	1.011	0.877	0.973	0.858	1.024	0.974	0.963	0.545	1.016	0.943	0.977
	$\delta = -0.5$	0.916	0.860	0.831	0.998	0.731	0.917	0.827	0.947	0.601	0.891	0.835	0.920
	$\delta = -1.5$	0.934	0.794	0.756	0.764	0.767	0.839	0.700	0.881	0.513	0.914	0.830	0.899
	$\delta = -3.5$	0.883	0.773	0.745	0.247	0.770	0.815	0.708	0.888	0.549	0.860	0.860	0.800
$\rho = 0.95$	$\delta = 0$	0.901	0.908	0.911	0.984	0.885	1.048	0.932	0.948	0.622	0.997	0.961	0.964
	$\delta = -0.5$	0.810	0.757	0.706	0.840	0.626	0.951	0.766	0.868	0.628	0.854	0.991	0.928
	$\delta = -1.5$	0.738	0.785	0.699	0.967	0.617	0.925	0.722	0.738	0.699	0.847	0.722	0.811
	$\delta = -3.5$	0.849	0.744	0.648	0.373	0.620	0.912	0.639	0.913	0.581	0.796	0.682	0.941

Notes: Data is simulated from a high-frequency VAR as given in equation (4.1). The time series length is $T = 500$ (lower-frequency). The information set is augmented and includes information up to $mT - 1$. The table reports the ratio $RMSE_{\{i\}}(500)/RMSE_{\{i\}}(100)$ for each model i for frequency mixtures $m = 3, 4, 12$.

Tab. 6.9: RMSE Ratios vs. LF-VAR: Process I with GARCH errors - Augmented information set ($T = 100$)

		$m=3$			$m=4$			$m=12$		
		Almon	Beta	MF-VAR	Almon	Beta	MF-VAR	Almon	Beta	MF-VAR
$\rho = 0.1$	$\delta = 0$	0.999	0.971	0.807	1.182	1.246	0.634	1.059	1.130	0.825
	$\delta = -0.5$	0.975	0.983	0.742	1.145	1.195	0.669	1.054	1.117	0.738
	$\delta = -1.5$	1.030	1.000	0.863	1.143	1.219	0.616	1.060	1.174	0.762
	$\delta = -3.5$	0.986	0.966	0.855	1.107	1.205	0.614	1.066	1.160	0.762
$\rho = 0.5$	$\delta = 0$	1.016	1.018	0.975	1.131	1.206	0.763	1.062	1.151	0.751
	$\delta = -0.5$	0.938	0.918	0.826	1.121	1.187	0.688	1.028	1.085	0.603
	$\delta = -1.5$	0.706	0.671	0.682	0.971	0.870	0.529	0.960	0.696	0.593
	$\delta = -3.5$	0.454	0.368	0.391	0.924	0.468	0.330	0.748	0.353	0.281
$\rho = 0.9$	$\delta = 0$	0.982	1.136	0.833	1.116	1.151	0.600	0.923	0.878	0.627
	$\delta = -0.5$	0.777	0.766	0.624	0.784	0.684	0.428	0.419	0.303	0.359
	$\delta = -1.5$	0.672	0.641	0.352	0.567	0.439	0.347	0.251	0.104	0.168
	$\delta = -3.5$	0.709	0.649	0.292	0.524	0.363	0.308	0.233	0.069	0.118
$\rho = 0.95$	$\delta = 0$	1.413	1.537	1.244	1.194	1.260	0.471	0.875	0.844	0.684
	$\delta = -0.5$	1.788	1.797	0.552	1.308	1.119	0.621	0.362	0.196	0.295
	$\delta = -1.5$	1.855	1.863	0.372	1.361	1.136	0.338	0.158	0.144	0.349
	$\delta = -3.5$	1.727	1.720	0.249	1.413	1.160	0.373	0.133	0.146	0.141

Notes: Data is simulated from a high-frequency VAR as given in equation (4.1) with GARCH errors (see equation (4.3)). The time series length is $T = 100$ (lower-frequency). Using the low-frequency data (temporally aggregated), we estimate a VAR(p) and conduct an out-of-sample one-step ahead forecast. The root mean squared error of this forecast, $RMSE(1, \mathcal{I}_t^A)$, is used as a benchmark. The competitors are the MIDAS model with exponential Almon lag and Beta lag weighting function and the mixed-frequency VAR ($RMSE_{\{i\}}(1, \mathcal{I}_t^M), i = \{Almon, Beta, MF-VAR\}$). The information set is augmented and includes information up to $mT - 1$. The table reports the ratio $RMSE_{\{i\}}(1, \mathcal{I}_t^M)/RMSE(1, \mathcal{I}_t^A)$ for frequency mixtures $m = 3, 4, 12$.

Tab. 6.10: RMSE Ratios $T = 100$ vs. $T = 500$: Process I with GARCH errors - Augmented information set

		$m = 3$				$m = 4$				$m = 12$			
		LF-VAR	Almon	Beta	MF-VAR	LF-VAR	Almon	Beta	MF-VAR	LF-VAR	Almon	Beta	MF-VAR
$\rho = 0.1$	$\delta = 0$	0.949	0.938	0.976	0.989	1.158	1.010	0.971	0.958	1.007	0.963	0.958	0.992
	$\delta = -0.5$	0.945	0.967	0.927	0.994	0.978	0.900	0.890	0.969	1.005	0.994	0.980	0.965
	$\delta = -1.5$	1.056	1.008	1.017	0.984	1.067	0.967	0.943	0.955	1.093	0.964	0.948	0.976
	$\delta = -3.5$	0.964	0.943	0.958	0.976	1.139	1.015	0.981	0.986	1.027	0.924	0.903	0.965
$\rho = 0.5$	$\delta = 0$	0.985	0.953	0.973	0.989	1.122	1.027	0.990	0.972	1.059	0.975	0.982	0.956
	$\delta = -0.5$	0.935	0.918	0.947	0.990	1.175	1.064	1.016	0.944	0.978	0.942	0.910	0.954
	$\delta = -1.5$	0.916	0.865	0.930	0.959	1.071	0.903	0.952	0.988	0.939	0.823	0.901	0.989
	$\delta = -3.5$	1.035	0.767	0.963	0.949	1.092	0.660	1.101	0.932	0.966	1.010	0.881	0.987
$\rho = 0.9$	$\delta = 0$	0.972	1.048	0.891	0.987	0.974	0.893	0.954	0.973	1.002	0.988	0.915	0.962
	$\delta = -0.5$	0.885	0.980	1.025	0.971	0.937	0.759	0.938	0.906	0.947	0.965	0.924	0.968
	$\delta = -1.5$	0.885	0.786	0.839	0.974	0.862	0.757	1.020	0.930	0.868	1.017	1.250	0.962
	$\delta = -3.5$	0.965	0.849	0.919	0.967	0.918	0.680	0.917	0.917	1.105	0.948	0.945	0.945
$\rho = 0.95$	$\delta = 0$	0.914	0.981	0.911	0.971	1.040	1.044	1.026	0.970	1.012	0.961	0.936	0.956
	$\delta = -0.5$	0.893	0.882	0.888	0.839	0.890	0.847	0.930	0.942	0.859	0.935	0.904	0.961
	$\delta = -1.5$	0.802	0.809	0.814	0.790	0.873	0.895	1.044	0.972	0.866	0.850	0.970	0.990
	$\delta = -3.5$	0.719	0.868	0.881	0.733	0.912	0.886	1.046	0.778	0.843	0.923	0.916	0.876

Notes: Data is simulated from a high-frequency VAR as given in equation (4.1) with GARCH errors (see equation (4.3)).The time series length is $T = 500$ (lower-frequency). The information set is augmented and includes information up to $mT - 1$. The table reports the ratio $RMSE_{\{i\}}(500)/RMSE_{\{i\}}(100)$ for each model i for frequency mixtures $m = 3, 4, 12$.

Tab. 6.11: RMSE Ratios vs. HF-VAR: Process II with homoscedastic errors ($T = 100$)

		$m = 3$				$m = 4$				$m = 12$			
		LF-VAR	Almon	Beta	MF-VAR	LF-VAR	Almon	Beta	MF-VAR	LF-VAR	Almon	Beta	MF-VAR
$\rho = 0.1$	$\delta = 0$	1.012	0.978	1.000	1.032	1.016	0.985	1.005	1.040	1.061	1.023	1.012	1.046
	$\delta = -0.5$	1.036	1.018	1.024	1.059	1.013	1.054	1.006	1.041	0.999	0.977	1.015	1.047
	$\delta = -1.5$	1.010	1.007	1.002	1.037	1.015	1.019	1.017	1.052	0.964	1.009	0.959	1.023
$\rho = 0.5$	$\delta = 0$	1.016	0.985	1.005	1.040	1.002	1.010	1.011	1.045	1.094	1.005	0.953	1.022
	$\delta = -0.5$	1.013	1.054	1.006	1.041	1.018	1.025	1.009	1.045	1.337	1.040	0.833	1.087
	$\delta = -1.5$	1.015	1.019	1.017	1.052	1.065	0.977	0.942	1.010	2.105	1.100	0.919	1.061
$\rho = 0.9$	$\delta = 0$	1.061	1.023	1.012	1.046	1.060	1.022	0.989	1.023	1.043	1.032	1.019	1.054
	$\delta = -0.5$	0.999	0.977	1.015	1.047	1.030	1.001	1.011	1.043	0.998	0.970	0.954	1.034
	$\delta = -1.5$	0.964	1.009	0.959	1.023	1.023	0.984	1.010	1.041	1.302	1.093	0.942	1.003

Notes: Data are simulated from a high-frequency VAR as given in equation (4.2) with homoscedastic errors. The time series length is $T = 100$ (lower-frequency). Using the high frequency data, we estimate a VAR(p) and conduct a out-of-sample one-step ahead forecast. The root mean squared error of this forecast, $RMSE(1, \mathcal{I}_t)$, is used as a benchmark. The competitors are the low-frequency VAR(p), where the higher-frequency variable is aggregated to the lower frequency by taking averages over the corresponding low-frequency period ($RMSE(1, \mathcal{I}_t^A)$), the MIDAS model with exponential Almon lag and Beta lag weighting function and the mixed-frequency VAR ($RMSE_{\{i\}}(1, \mathcal{I}_t^M), i = \{Almon, Beta, MF-VAR\}$). The table reports the ratio $RMSE(1, \mathcal{I}_t^A)/RMSE(1, \mathcal{I}_t)$ and $RMSE\{i\}(1, \mathcal{I}_t^M)/RMSE_(1, \mathcal{I}_t)$ for frequency mixtures $m = 3, 4, 12$.

Tab. 6.12: RMSE Ratios vs. LF-VAR: Process II with homoscedastic errors ($T = 100$)

		$m = 3$			$m = 4$			$m = 12$		
		Almon	Beta	MF-VAR	Almon	Beta	MF-VAR	Almon	Beta	MF-VAR
$\rho = 0.1$	$\delta = 0$	0.966	0.988	1.020	0.969	0.988	1.023	0.964	0.954	0.986
	$\delta = -0.5$	0.983	0.989	1.022	1.041	0.993	1.028	0.978	1.016	1.049
	$\delta = -1.5$	0.997	0.992	1.026	1.004	1.001	1.036	1.047	0.994	1.061
$\rho = 0.5$	$\delta = 0$	0.969	0.988	1.023	1.008	1.009	1.043	0.918	0.871	0.934
	$\delta = -0.5$	1.041	0.993	1.028	1.006	0.991	1.026	0.778	0.623	0.813
	$\delta = -1.5$	1.004	1.001	1.036	0.918	0.885	0.949	0.522	0.437	0.504
$\rho = 0.9$	$\delta = 0$	0.964	0.954	0.986	0.964	0.933	0.965	0.989	0.977	1.010
	$\delta = -0.5$	0.978	1.016	1.049	0.972	0.982	1.013	0.972	0.956	1.035
	$\delta = -1.5$	1.047	0.994	1.061	0.962	0.988	1.018	0.840	0.724	0.771

Notes: This table builds upon the same data as in Table 6.11 but the benchmark is the LF-VAR. The table reports the ratio $RMSE_{\{i\}}(1, \mathcal{I}_t^M)/RMSE(1, \mathcal{I}_t^A)$ for frequency mixtures $m = 3, 4, 12$.

Tab. 6.13: RMSE Ratios $T = 100$ vs. $T = 500$: Process II with homoscedastic errors

		$m=3$				$m=4$				$m=12$			
		LF-VAR	Almon	Beta	MF-VAR	LF-VAR	Almon	Beta	MF-VAR	LF-VAR	Almon	Beta	MF-VAR
$\rho = 0.1$	$\delta = 0$	0.965	0.944	0.986	0.885	0.947	0.957	0.933	0.885	0.983	0.977	0.997	0.899
	$\delta = -0.5$	0.953	0.893	0.923	0.913	1.002	1.028	0.986	0.913	1.013	1.000	0.979	1.063
	$\delta = -1.5$	0.921	0.977	0.958	0.907	0.996	0.964	0.927	0.907	1.047	1.052	1.021	1.106
$\rho = 0.5$	$\delta = 0$	0.947	0.957	0.933	0.858	1.025	0.996	0.980	0.858	0.990	0.887	0.932	1.070
	$\delta = -0.5$	1.002	1.028	0.986	0.926	0.898	0.891	0.867	0.926	0.984	0.948	0.919	0.821
	$\delta = -1.5$	0.996	0.964	0.927	0.981	1.054	0.952	0.961	0.981	0.931	0.778	0.933	0.886
$\rho = 0.9$	$\delta = 0$	0.983	0.977	0.997	0.899	0.997	1.078	1.032	0.899	0.950	0.918	0.958	0.834
	$\delta = -0.5$	1.013	1.000	0.979	1.063	1.009	1.001	0.975	1.063	0.916	0.867	0.847	0.820
	$\delta = -1.5$	1.047	1.052	1.021	1.006	0.858	0.842	0.810	1.010	0.869	0.992	0.898	0.837

Notes: Data is simulated from a high-frequency VAR as given in equation (4.2) with homoscedastic errors. The time series length is $T = 500$ (lower-frequency). The table reports the ratio $RMSE_{\{i\}}(500)/RMSE_{\{i\}}(100)$ for each model i for frequency mixtures $m = 3, 4, 12$.

Tab. 6.14: RMSE Ratios vs. HF-VAR: Process II with GARCH errors ($T = 100$)

		$m=3$				$m=4$				$m=12$			
		LF-VAR	Almon	Beta	MF-VAR	LF-VAR	Almon	Beta	MF-VAR	LF-VAR	Almon	Beta	MF-VAR
$\rho = 0.1$	$\delta = 0$	0.974	0.984	0.941	1.008	1.046	1.004	1.050	1.085	1.038	1.033	1.002	1.095
	$\delta = -0.5$	0.967	1.026	0.995	1.029	1.051	1.016	1.065	1.071	1.060	1.072	1.025	1.034
	$\delta = -1.5$	1.031	0.971	1.012	1.046	1.019	1.032	1.071	1.028	1.057	1.032	1.040	1.122
$\rho = 0.5$	$\delta = 0$	1.046	1.004	1.050	1.085	0.978	1.000	1.014	1.257	1.101	1.090	0.968	1.003
	$\delta = -0.5$	1.051	1.016	1.065	1.071	1.042	1.032	1.046	1.071	1.285	0.977	0.851	1.035
	$\delta = -1.5$	1.019	1.032	1.071	1.028	0.966	1.004	0.989	1.019	2.199	1.298	0.931	1.012
$\rho = 0.9$	$\delta = 0$	1.038	1.033	1.002	1.095	1.070	1.006	0.986	1.132	1.051	1.067	0.980	1.196
	$\delta = -0.5$	1.060	1.072	1.025	1.034	1.040	1.074	1.022	1.114	0.996	1.053	1.037	1.068
	$\delta = -1.5$	1.057	1.032	1.040	1.122	1.027	1.034	1.023	1.071	1.194	1.000	0.893	1.125

Notes: Data are simulated from a high-frequency VAR as given in equation (4.2) with GARCH errors (see equation (4.3)). The time series length is $T = 100$ (lower-frequency). Using the high frequency data, we estimate a VAR(p) and conduct a out-of-sample one-step ahead forecast. The root mean squared error of this forecast, $RMSE(1, \mathcal{I}_t)$, is used as a benchmark. The competitors are the low-frequency VAR(p), where the higher-frequency variable is aggregated to the lower frequency by taking averages over the corresponding low-frequency period ($RMSE(1, \mathcal{I}_t^A)$), the MIDAS model with exponential Almon lag and Beta lag weighting function and the mixed-frequency VAR ($RMSE_{\{i\}}(1, \mathcal{I}_t^M), i = \{Almon, Beta, MF-VAR\}$). The table reports the ratio $RMSE(1, \mathcal{I}_t^A)/RMSE(1, \mathcal{I}_t)$ and $RMSE_{\{i\}}(1, \mathcal{I}_t^M)/RMSE(1, \mathcal{I}_t)$ for frequency mixtures $m = 3, 4, 12$.

Tab. 6.15: RMSE Ratios vs. LF-VAR: Process II with GARCH errors ($T = 100$)

		$m=3$			$m=4$			$m=12$		
		Almon	Beta	MF-VAR	Almon	Beta	MF-VAR	Almon	Beta	MF-VAR
$\rho = 0.1$	$\delta = 0$	1.010	0.941	1.034	0.960	1.050	1.037	0.995	1.002	1.054
	$\delta = -0.5$	1.061	0.995	1.064	0.967	1.065	1.019	1.012	1.025	0.976
	$\delta = -1.5$	0.942	1.012	1.014	1.013	1.071	1.009	0.976	1.040	1.062
$\rho = 0.5$	$\delta = 0$	0.960	1.050	1.037	1.023	1.014	1.286	0.989	0.968	0.911
	$\delta = -0.5$	0.967	1.065	1.019	0.991	1.046	1.028	0.760	0.851	0.805
	$\delta = -1.5$	1.013	1.071	1.009	1.039	0.989	1.054	0.590	0.931	0.460
$\rho = 0.9$	$\delta = 0$	0.995	1.002	1.054	0.940	0.986	1.058	1.015	0.980	1.138
	$\delta = -0.5$	1.012	1.025	0.976	1.032	1.022	1.071	1.057	1.037	1.072
	$\delta = -1.5$	0.976	1.040	1.062	1.007	1.023	1.043	0.837	0.893	0.942

Notes: This table builds upon the same data as in Table 6.14 but the benchmark is the LF-VAR. The table reports the ratio $RMSE_{\{i\}}(1, \mathcal{I}_t^M)/RMSE(1, \mathcal{I}_t^A)$ for frequency mixtures $m = 3, 4, 12$.

Tab. 6.16: RMSE Ratios $T = 100$ vs. $T = 500$: Process II with GARCH errors

		$m = 3$				$m = 4$				$m = 12$			
		LF-VAR	Almon	Beta	MF-VAR	LF-VAR	Almon	Beta	MF-VAR	LF-VAR	Almon	Beta	MF-VAR
$\rho = 0.1$	$\delta = 0$	0.886	0.876	0.882	0.874	0.998	0.968	0.927	0.874	1.009	0.907	0.923	0.976
	$\delta = -0.5$	0.990	0.935	0.953	0.984	1.012	0.977	0.936	0.984	1.094	0.967	0.995	0.992
	$\delta = -1.5$	1.013	1.049	0.993	0.868	1.084	0.972	0.931	0.868	1.045	0.978	0.932	0.979
$\rho = 0.5$	$\delta = 0$	0.998	0.968	0.927	0.942	1.086	0.963	0.952	0.942	0.952	0.864	0.927	0.930
	$\delta = -0.5$	1.012	0.977	0.936	0.865	0.979	0.905	0.913	0.865	1.053	0.977	0.963	0.998
	$\delta = -1.5$	1.084	0.972	0.931	0.875	1.205	1.023	1.054	0.875	0.878	0.769	0.858	0.790
$\rho = 0.9$	$\delta = 0$	1.009	0.907	0.923	0.976	1.068	1.018	1.025	0.976	0.910	0.936	0.987	0.924
	$\delta = -0.5$	1.094	0.967	0.995	0.992	1.044	0.938	0.989	0.992	0.996	0.937	0.972	0.878
	$\delta = -1.5$	1.045	0.978	0.932	0.979	0.895	0.903	0.923	0.979	0.979	1.045	0.975	0.887

Notes: Data is simulated from a high-frequency VAR as given in equation (4.2) with GARCH errors (see equation (4.3)). The time series length is $T = 500$ (lower-frequency). The table reports the ratio $RMSE_{\{i\}}(500)/RMSE_{\{i\}}(100)$ for each model i for frequency mixtures $m = 3, 4, 12$.

Tab. 6.17: RMSE Ratios vs. LF-VAR: Process II with homoscedastic errors - Augmented information set ($T = 100$)

		$m = 3$			$m = 4$			$m = 12$		
		Almon	Beta	MF-VAR	Almon	Beta	MF-VAR	Almon	Beta	MF-VAR
$\rho = 0.1$	$\delta = 0$	0.974	0.951	0.924	0.911	0.912	0.692	0.788	1.008	0.782
	$\delta = -0.5$	0.996	1.040	0.980	0.989	1.038	1.057	0.794	0.972	0.768
	$\delta = -1.5$	0.966	1.010	0.904	0.984	0.977	1.011	0.842	0.992	0.833
$\rho = 0.5$	$\delta = 0$	0.911	0.912	0.692	0.972	1.001	0.881	0.881	0.957	1.095
	$\delta = -0.5$	0.989	1.038	1.057	0.951	0.953	1.077	0.805	0.656	0.751
	$\delta = -1.5$	0.984	0.977	1.011	0.871	0.666	0.637	0.166	0.249	0.177
$\rho = 0.9$	$\delta = 0$	0.788	1.008	0.782	0.799	1.035	0.809	0.784	0.953	0.753
	$\delta = -0.5$	0.794	0.972	0.768	0.748	0.957	0.767	0.724	0.523	0.700
	$\delta = -1.5$	0.842	0.992	0.833	0.713	0.760	0.583	0.652	0.247	0.546

Notes: Data is simulated from a high-frequency VAR as given in equation (4.2) with homoscedastic errors. The time series length is $T = 100$ (lower-frequency). Using the low-frequency data (temporally aggregated), we estimate a VAR(p) and conduct an out-of-sample one-step ahead forecast. The root mean squared error of this forecast, $RMSE(1, \mathcal{I}_t^A)$, is used as a benchmark. The competitors are the MIDAS model with exponential Almon lag and Beta lag weighting function and the mixed-frequency VAR ($RMSE_{\{i\}}(1, \mathcal{I}_t^M), i = \{Almon, Beta, MF-VAR\}$). The information set is augmented and includes information up to $mT - 1$. The table reports the ratio $RMSE_{\{i\}}(1, \mathcal{I}_t^M)/RMSE(1, \mathcal{I}_t^A)$ for frequency mixtures $m = 3, 4, 12$.

Tab. 6.18: RMSE Ratios $T = 100$ vs. $T = 500$: Process II with homoscedastic errors - Augmented information set

		$m=3$				$m=4$				$m=12$			
		LF-VAR	Almon	Beta	MF-VAR	LF-VAR	Almon	Beta	MF-VAR	LF-VAR	Almon	Beta	MF-VAR
$\rho = 0.1$	$\delta = 0$	0.965	0.942	0.918	1.074	0.947	0.927	0.924	1.074	0.983	0.842	0.895	0.726
	$\delta = -0.5$	0.953	0.940	0.953	1.208	1.002	1.003	1.016	1.208	1.013	0.942	0.973	0.836
	$\delta = -1.5$	0.921	0.954	0.944	1.009	0.996	0.964	1.002	1.009	1.047	0.948	1.005	0.865
$\rho = 0.5$	$\delta = 0$	0.947	0.927	0.924	0.624	1.025	1.040	0.993	0.624	0.990	0.959	0.948	0.953
	$\delta = -0.5$	1.002	1.003	1.016	0.982	0.898	0.951	0.944	0.982	0.984	1.011	0.968	0.680
	$\delta = -1.5$	0.996	0.964	1.002	1.036	1.054	0.921	0.931	1.036	0.931	1.166	0.978	0.963
$\rho = 0.9$	$\delta = 0$	0.983	0.842	0.895	0.726	0.997	0.924	1.009	0.726	0.950	0.840	0.954	0.774
	$\delta = -0.5$	1.013	0.942	0.973	0.836	1.009	0.919	0.916	0.836	0.916	0.867	0.912	0.827
	$\delta = -1.5$	1.047	0.948	1.005	0.865	0.858	0.777	0.769	0.865	0.869	0.832	0.642	1.214

Notes: Data is simulated from a high-frequency VAR as given in equation (4.2). The time series length is $T = 500$ (lower-frequency). The information set is augmented and includes information up to $mT - 1$. The table reports the ratio $RMSE_{\{i\}}(500)/RMSE_{\{i\}}(100)$ for each model i for frequency mixtures $m = 3, 4, 12$.

Tab. 6.19: RMSE Ratios vs. LF-VAR: Process II with GARCH errors - Augmented information set ($T = 100$)

		$m = 3$			$m = 4$			$m = 12$		
		Almon	Beta	MF-VAR	Almon	Beta	MF-VAR	Almon	Beta	MF-VAR
$\rho = 0.1$	$\delta = 0$	1.006	0.971	0.839	0.976	1.019	1.073	0.863	1.034	1.032
	$\delta = -0.5$	1.027	0.988	0.819	0.981	1.021	1.054	0.875	1.051	0.901
	$\delta = -1.5$	0.981	1.004	0.820	1.001	1.017	0.979	0.871	1.037	0.911
$\rho = 0.5$	$\delta = 0$	0.976	1.019	1.073	0.952	0.981	1.162	0.874	0.932	0.820
	$\delta = -0.5$	0.981	1.021	1.054	0.909	0.887	0.952	0.763	0.579	0.464
	$\delta = -1.5$	1.001	1.017	0.979	0.922	0.669	0.665	0.134	0.216	0.125
$\rho = 0.9$	$\delta = 0$	0.863	1.034	1.032	0.862	1.027	0.847	0.877	0.919	0.884
	$\delta = -0.5$	0.875	1.051	0.901	0.831	1.010	0.912	0.776	0.482	0.781
	$\delta = -1.5$	0.871	1.037	0.911	0.786	0.756	0.770	0.761	0.343	0.403

Notes: Data is simulated from a high-frequency VAR as given in equation (4.2) with GARCH errors (see equation (4.3)). The time series length is $T = 100$ (lower-frequency). Using the low-frequency data (temporally aggregated), we estimate a VAR(p) and conduct an out-of-sample one-step ahead forecast. The root mean squared error of this forecast, $RMSE(1, \mathcal{I}_t^A)$, is used as a benchmark. The competitors are the MIDAS model with exponential Almon lag and Beta lag weighting function and the mixed-frequency VAR ($RMSE_{\{i\}}(1, \mathcal{I}_t^M), i = \{Almon, Beta, MF-VAR\}$). The information set is augmented and includes information up to $mT-1$. The table reports the ratio $RMSE_{\{i\}}(1, \mathcal{I}_t^M)/RMSE(1, \mathcal{I}_t^A)$ for frequency mixtures $m = 3, 4, 12$.

Tab. 6.20: RMSE Ratios $T = 100$ vs. $T = 500$: Process I with GARCH errors - Augmented information set

		$m = 3$				$m = 4$				$m = 12$			
		LF-VAR	Almon	Beta	MF-VAR	LF-VAR	Almon	Beta	MF-VAR	LF-VAR	Almon	Beta	MF-VAR
$\rho = 0.1$	$\delta = 0$	0.886	0.875	0.891	0.965	0.998	0.935	0.952	0.965	1.009	0.943	0.879	0.970
	$\delta = -0.5$	0.990	0.959	0.912	0.974	1.012	0.959	1.005	0.974	1.094	1.042	0.988	0.985
	$\delta = -1.5$	1.013	1.007	0.979	0.961	1.084	0.947	0.989	0.961	1.045	0.988	0.976	0.966
$\rho = 0.5$	$\delta = 0$	0.998	0.935	0.952	0.994	1.086	1.032	1.064	0.994	0.952	0.925	0.971	0.979
	$\delta = -0.5$	1.012	0.959	1.005	1.000	0.979	0.989	1.056	1.000	1.053	1.049	0.964	0.978
	$\delta = -1.5$	1.084	0.947	0.989	0.991	1.205	0.926	0.983	0.991	0.878	1.166	0.966	0.954
$\rho = 0.9$	$\delta = 0$	1.009	0.943	0.879	0.970	1.068	1.059	0.991	0.970	0.910	1.054	0.977	0.986
	$\delta = -0.5$	1.094	1.042	0.988	0.985	1.044	1.091	0.909	0.985	0.996	1.027	1.004	0.964
	$\delta = -1.5$	1.045	0.988	0.976	0.966	0.895	1.035	0.833	0.966	0.979	0.952	0.522	0.975

Notes: Data is simulated from a high-frequency VAR as given in equation (4.2) with GARCH errors (see equation (4.3)).The time series length is $T = 500$ (lower-frequency). The information set is augmented and includes information up to $mT - 1$. The table reports the ratio $RMSE_{\{i\}}(500)/RMSE_{\{i\}}(100)$ for each model i for frequency mixtures $m = 3, 4, 12$.

Tab. 6.21: RMSE Ratios vs. LF-AR and ($T = 100$) vs. ($T = 500$) : Process III

		Panel A: Ratio vs. LF-AR								
		$m = 3$			$m = 4$			$m = 12$		
		Almon	Beta	MF-VAR	Almon	Beta	MF-VAR	Almon	Beta	MF-VAR
WF 1	$\beta_1 = 0.6$	0.392	0.390	0.462	0.473	0.461	0.463	0.901	0.921	0.921
	$\beta_1 = 3$	0.359	0.359	0.361	0.498	0.498	0.475	0.939	0.894	0.895
	$\beta_1 = 6$	0.366	0.378	0.365	0.458	0.464	0.430	0.941	0.882	0.892
WF 2	$\beta_1 = 0.6$	0.280	0.279	0.355	0.333	0.321	0.340	0.823	0.659	0.662
	$\beta_1 = 3$	0.290	0.277	0.392	0.331	0.320	0.417	0.829	0.649	0.655
	$\beta_1 = 6$	0.266	0.250	0.303	0.328	0.306	0.487	0.771	0.633	0.640
		Panel B: Ratio ($T = 100$) vs. ($T = 500$)								
WF 1	$\beta_1 = 0.6$	0.857	0.838	0.960	0.967	0.979	0.959	1.024	0.997	1.000
	$\beta_1 = 3$	0.881	0.885	0.947	0.942	0.941	0.981	1.008	0.991	0.992
	$\beta_1 = 6$	0.928	0.887	0.951	0.974	0.943	0.974	1.025	0.999	1.005
WF 2	$\beta_1 = 0.6$	0.939	0.877	0.986	0.917	0.938	0.988	1.015	1.064	1.064
	$\beta_1 = 3$	0.923	0.904	0.983	0.945	0.915	0.951	0.927	0.924	0.927
	$\beta_1 = 6$	0.901	0.901	0.858	0.932	0.985	0.988	0.927	0.953	0.958

Notes: Data are simulated from a MIDAS model (4.4). The time series length is $T = 100$. We use an AR(p) model and calculate an out-of-sample one-step ahead forecast. The root mean squared error of this forecast, $RMSE(1, \mathcal{I}_t^A)$, is used as a benchmark. The competitors are the MIDAS model with exponential Almon lag and Beta weighting function and the mixed-frequency VAR ($RMSE_{\{i\}}(1, \mathcal{I}_t^M), i = \{Almon, Beta, MF-VAR\}$). Panel A reports the ratio $RMSE_{\{i\}}(1, \mathcal{I}_t^M)/RMSE_{AR}(1, \mathcal{I}_t^A)$ for frequency mixtures $m = 3, 4, 12$. Panel B reports the ratio $RMSE_{\{i\}}(500)/RMSE_{\{i\}}(100)$ for each model i for frequency mixtures $m = 3, 4, 12$.

Tab. 6.22: RMSE Ratios vs. LF-AR and ($T = 100$) vs. ($T = 500$) : Process III - Augmented information set

					Panel A: Ratio vs. LF-AR					
			$m = 3$			$m = 4$			$m = 12$	
		Almon	Beta	MF-VAR	Almon	Beta	MF-VAR	Almon	Beta	MF-VAR
WF 1	$\beta_1 = 0.6$	0.126	0.128	0.405	0.185	0.115	0.205	0.581	0.106	0.407
	$\beta_1 = 3$	0.026	0.027	0.231	0.446	0.031	0.132	0.463	0.041	0.600
	$\beta_1 = 6$	0.014	0.016	0.213	0.232	0.019	0.129	0.463	0.032	0.561
WF 2	$\beta_1 = 0.6$	0.377	0.155	0.337	0.248	0.155	0.341	0.432	0.124	0.549
	$\beta_1 = 3$	0.125	0.055	0.365	0.183	0.065	0.308	0.215	0.039	0.503
	$\beta_1 = 6$	0.074	0.042	0.515	0.034	0.060	0.321	0.181	0.031	0.549
					Panel B: Ratio ($T = 100$) vs. ($T = 500$)					
WF 1	$\beta_1 = 0.6$	0.832	0.952	0.989	0.772	0.893	0.953	0.995	0.983	0.402
	$\beta_1 = 3$	0.669	0.950	0.978	0.988	0.792	0.892	0.945	0.729	0.599
	$\beta_1 = 6$	0.689	0.856	0.882	0.996	0.721	0.940	0.983	0.727	0.560
WF 2	$\beta_1 = 0.6$	0.913	0.873	0.978	0.930	0.880	0.890	0.960	0.893	0.544
	$\beta_1 = 3$	0.611	0.986	0.921	0.985	0.857	0.888	0.943	0.824	0.503
	$\beta_1 = 6$	0.593	0.852	0.966	0.965	0.813	0.899	0.847	0.768	0.543

Notes: Data are simulated from a MIDAS model (4.4). The time series length is $T = 100$. We use an AR(p) model and calculate an out-of-sample one-step ahead forecast. The root mean squared error of this forecast, $RMSE(1, \mathcal{I}_t^A)$, is used as a benchmark. The competitors are the MIDAS model with exponential Almon lag and Beta weighting function and the mixed-frequency VAR ($RMSE_{\{i\}}(1, \mathcal{I}_t^M), i = \{Almon, Beta, MF-VAR\}$). Panel A reports the ratio $RMSE_{\{i\}}(1, \mathcal{I}_t^M)/RMSE_{AR}(1, \mathcal{I}_t^A)$ for frequency mixtures $m = 3, 4, 12$ where the information set is augmented and includes information up to $mT - 1$. Panel B reports the ratio $RMSE_{\{i\}}(500)/RMSE_{\{i\}}(100)$ for each model i for frequency mixtures $m = 3, 4, 12$.

Tab. 6.23: RMSE Ratios: Process IV

	Panel A: Ratio vs. HF-VAR			
	Standard information set			
	LF-VAR	Almon	Beta	MF-VAR
$\rho = 0.1$	1.019	1.043	1.031	1.194
$\rho = 0.5$	1.056	1.034	1.060	1.008
$\rho = 0.9$	1.160	1.030	1.044	1.020
$\rho = 0.95$	1.351	1.039	1.032	1.010
	Panel B: Ratio vs. LF-VAR			
	Standard information set			
		Almon	Beta	MF-VAR
$\rho = 0.1$		1.023	1.012	1.171
$\rho = 0.5$		0.979	1.003	0.954
$\rho = 0.9$		0.888	0.900	0.879
$\rho = 0.95$		0.769	0.764	0.747
	Panel C: Ratio $T = 500$ vs. $T = 100$			
	Standard information set			
$\rho = 0.1$	0.889	0.978	0.988	1.022
$\rho = 0.5$	0.903	1.017	0.991	1.139
$\rho = 0.9$	0.867	1.181	1.133	1.041
$\rho = 0.95$	0.729	1.055	1.034	1.014
	Panel D: Ratio vs. LF-VAR			
	Augmented information set			
$\rho = 0.1$		1.031	1.011	1.288
$\rho = 0.5$		0.883	0.866	0.828
$\rho = 0.9$		0.405	0.413	0.721
$\rho = 0.95$		0.281	0.284	0.632
	Panel E: Ratio $T = 500$ vs. $T = 100$			
	Augmented information set			
$\rho = 0.1$	0.889	0.964	0.983	0.989
$\rho = 0.5$	0.903	1.078	1.104	0.988
$\rho = 0.9$	0.867	0.929	0.922	0.993
$\rho = 0.95$	0.729	1.067	1.066	0.996

Notes: Data are simulated from a high-frequency VAR as given in equation (4.6). Panel A reports the ratio $RMSE(1, \mathcal{I}_t^M)/RMSE(1, \mathcal{I}_t)$ where the HF-VAR serves as a benchmark. In Panel B the LF-VAR is the benchmark model $(RMSE(1, \mathcal{I}_t^M)/RMSE(1, \mathcal{I}_t^A)$. Panel C reports the ratio $RMSE_{\{i\}}(500)/RMSE_{\{i\}}(100)$ for each model i. In Panel D the information set is augmented and includes information up to $mT - 1$. The LF-VAR is the benchmark model. Panel E reports the ratio $RMSE_{\{i\}}(500)/RMSE_{\{i\}}(100)$ when the information set is augmented.

7. EMPIRICAL EXAMPLES

7.1 Introduction

In this chapter we apply the mixed-frequency time series models to a real-data example. From the Monte Carlo study we know that is worth using mixed-frequency time series models to improve forecasting accuracy. We also showed there that intra-period information does help to estimate the current period value of the variable under investigation. This phenomenon is labelled 'nowcasting' in the literature (see Giannone, Reichlin, and Small (2008)). Nowcasting is especially important for monetary policy decisions in real time. Assessments of current and future economic conditions are often based on incomplete data. Most data monitored by central banks are released with a lag and are subsequently revised. In principle, any data release may potentially affect current-quarter estimation of GDP and the precision of the results. *A priori*, there is no reason to discard any information.

We analyze the forecasting performance of the MIDAS and the mixed-frequency VAR models with two target variables: German GDP and Industrial Production (IP). We extend two previous studies in several ways. First, Mittnik and Zadrozny (2005) only employed the Ifo indicators within the mixed-frequency VAR framework to forecast German GDP growth. There are no comparisons with other models and indicators. Similarly, Marcellino and Schumacher (2007) use factors as leading indicators and employ different versions of the MIDAS model to nowcast German GDP growth. Again, the authors do not compare their results with other single indicators and models. In this chapter we conduct a case study with both monthly and daily data. First we forecast German quarterly and yearly GDP growth with two Prominent, monthly, single indicators for the German economy: the Ifo Business Climate Index and the ZEW Index of Economic Sentiment. Both indicators are available early and are not subject to revision. Furthermore, they proved to be successful for forecasting in many empirical applications.[1]

Additionally to the two survey indicators, we extract static and dynamic factors from a large data set exhibiting publication lags. In recent years, large-dimensional, dynamic, factor models have become popular in empirical macroeconomics. Factor models are able to handle many variables without violating the objective of parsimony and the degrees of freedom problem often faced in regression-based analysis. The use of many variables reflects the practice of central banks of looking at everything, as pointed out for example by Bernanke and Boivin (2003). By using factor models one can circumvent the problem of which variables to include in a structural model. See Breitung and Eickmeier (2006) and Eickmeier and Ziegler (2008) for an exposition and a literature review.

In sum we compare three indicators. First we want answer the question of how the mixed-frequency models perform with real data. Can the results from the Monte Carlo study be confirmed? Second, how does intra-quarterly information contribute to short- (nowcasting) and longer-horizon forecasts? Finally, we want to investigate whether it is worthwhile gathering many time series when single indicators are available earlier? If the single indicator provides comparable forecasting results, why should we care about adjusting for publication lags and revisions? Dreger and Schumacher (2004) demonstrate that factor models yield more accurate forecasting results but these are not statistical significant compared with forecasts generated with single indicators.

In a second step we conduct a nowcasting experiment. If the monthly information is informative and predictive for German GDP, what about weekly and daily data? We gather many daily time series (interest rates, stock returns, exchange rates) and extract daily factors from dynamic factor models. Factors are extracted for every day of the year, that is also on weekends and

[1] See the literature section and Robinzonov and Wohlrabe (2008) for further details.

holidays. We intended to use the EM algorithm which can handle missing observations as outlined in chapter 2, but this was computationally infeasible as the time series is too long. Therefore we interpolated the missing observations. We plug the obtained daily factors into the MIDAS model which generates the forecasts. We end up with a model which updates the estimate of the target variable as new daily information arrives.

Our approach is closely related to some other papers that investigate higher-frequency estimates of GDP. Evans (2005) constructs a model for updating the nowcast of long-run GDP as new information becomes available. However, his approach can handle only a limited number of series. This argument also applies to Aruoba, Diebold, and Scotti (2008), who rely on factor analysis but their study is limited to very few variables. The inclusion of more variables would raise the same problem as in the MF-VAR approach: the model estimation becomes infeasible. In general these approaches estimate an unobserved variable and do not produce forecasts in a strict sense. Giannone, Reichlin, and Small (2008) also rely on factor estimation and forecast GDP with a bridge equation approach, but their model operates on a monthly basis. Nevertheless, the authors demonstrate that as new information becomes available at a specific day within a quarter, the bridge model can be re-estimated and thus the current estimate of GDP can be updated.

All three papers focus on US data. We are the first to use daily data without temporal aggregation. Furthermore we are the first to apply high-frequency forecasting to German data. We update the forecast of GDP each day in a quarter.

This chapter is organized as follows: we start with representing some related single-frequency literature on forecasting German GDP. Then we provide details about the monthly data set. After the outline of details of the case study we present the results for the monthly data. Finally we nowcast the German economy on a daily basis for both GDP and IP.

7.2 Related Literature

In addition to the mixed-frequency approaches of Mittnik and Zadrozny (2005) and Marcellino and Schumacher (2007) we want to sketch some empirical applications of forecasting German GDP. Compared with the forecasts of US GDP, the literature for Germany is rather well arranged. Due to the unification of Germany 1991 there is a structural break in every macroeconomic time series at this date. Therefore more recently there are papers on forecasting German GDP, as the time series is now considered to be long enough to conduct comparative studies.

Camba-Mendez, Kapetanios, Smith, and Weale (2001) propose an automatic leading indicator approach (ALI) based on dynamic factor models. In a two-step approach, they show for Germany and other European countries that the on average ALI forecasts better than traditional VAR and BVAR models with traditional leading indicators. Kirschgässner and Savioz (2001) employ an ARX(4,4) model with a rolling forecasting scheme and find that the daily interest rate is the best financial predictor for four quarters ahead for the time period 1980:I-1989:IV. For the period 1992:III-1999:IV the money aggregate M1 proved to be the best predictor. Similar papers are Davis and Henry (1994), Krämer and Langfeld (1993), and Hagen and Kirchgässner (1996).

In a huge comparative study, Stock and Watson (2003) investigate the predictive power of asset prices for forecasting output and inflation for different OECD countries. The target variable is the approximate yearly growth rate from 1971 to 1999. With an direct ARX model they find that many asset prices do not prove to be better than univariate benchmark models.

Dreger and Schumacher (2004) compare the forecast performance of the Ifo Business Climate Index to static and dynamic factor models. The forecast variable depends on the forecast horizon. They forecast the growth in the GDP series between t and the period $t + h$.[2] In the direct ARX approach, GDP enters on the right-hand side of the equation as the approximate quarterly growth rate. They show that, based on the RMSE criterion, the dy-

[2] $y_{t+h} = \log(y_{t+h}/y_t)$

namic factor performs better than the Ifo indicator over all eight forecast horizons. But these advantages are not systematic, as shown in the insignificant results obtained with the Diebold-Mariano test.

Kholodilin and Siliverstovs (2006) proceed in a similar way. They compare the most common indicators for Germany with different diffusion indices (as well as their first differences). Their sample ranges from 1991:I to 2004:IV. The authors employ annualized quarterly and biannual and approximate yearly growth rates. For each horizon, with a direct ARX model, they generate 28 out-of-sample forecasts within both a rolling and a recursive forecasting scheme. For model selection they use both the BIC criterion as well as the automatic econometric model selection program PcGets. They find a relatively poor performance for the indicators compared with the naive benchmark (random walk). In general the first differences of the indicator are better than the levels. The authors detect a structural break in the GDP series around 2001 and claim that no indicator was able to predict and accommodate the structural break.

Robinzonov and Wohlrabe (2008) provide an excellent survey of the literature that deals with forecasting German industrial production.

7.3 A Case Study Based on Monthly Data

7.3.1 The Monthly Data Set

The time series to be forecast is the real GDP (calendar and seasonal adjusted) in Germany from 1991:I to 2006:IV. We do not consider data before 1991 to circumvent the problem of structural change, or level shifts in the data due to German unification. In contrast to Clements and Galvao (2008) we use final, not real-time data. This is a contradiction with the opinion of Pesaran and Timmermann (2005) that any real-time econometric model should make use of real-time data in all stages, so as not to overstate the degree of predictability, as shown by Diebold and Rudebusch (1991). The effects of data vintages on model specification and forecast evaluation have been addressed in a number of papers (see among others Croushore and Stark (2003) and Koenig, Dolmas, and Piger (2003)). We use final data for two reasons. First, we want to focus on the model competition between MIDAS, mixed-frequency VAR and single frequency models. Second, Schumacher and Breitung (2008) showed that for Germany, data revisions do not greatly affect forecasts. We calculate both quarterly and yearly growth rates (log differences).

We focus on two popular leading indicators for the German economy: the Ifo Business Climate Index and the ZEW Index of Economic Sentiment. The Ifo Business Climate Index is based on about 7,000 monthly survey responses from firms in manufacturing, construction, and wholesale and retail trade. The firms are asked to give their assessments of the current business situation and their expectations for the next six months. The balance value of the current business situation is the difference between the percentages of the responses "good" and "poor"; the balance value of expectations is the difference between the percentages of the responses "more favourable" and "more unfavourable". The replies are weighted in proportion to the importance of the industry and then aggregated. The business climate is a transformed mean of the balances of the business situation and expectations. For further information see Goldrian (2007). The ZEW Indicator of Economic Sentiment is published monthly. Up to 350 financial experts take part in the survey. The indicator reflects the difference between the share of analysts that are optimistic and the share of analysts that are pessimistic about the expected economic development in Germany in the next six months (see Hüfner and Schröder (2002)).

In addition to the two leading indicators, we want to include factors extracted from dynamic factor models. We use an extended data set by Marcellino and Schumacher (2007). We employ 151 monthly indicators from 1992:01 until 2006:12. The data set is final data for several reasons. First, a real-time data set for as many time series is not available for Germany. Second, Schumacher and Breitung (2008) show that the forecast performance changes little when real-time data are used instead of final data.[3]

To consider the availability of time series at the end of the sample period due to different publication lags (ragged-edge), we follow Banbura and Runstler (2007) and replicate the jarred-edge structure from the final vintage of data that are available.[4] The data used in this study, were downloaded on 21 June 2008.[5] Starting from the original data D_{mT}, we reconstruct the data sets, which were available earlier $\tau < mT$, by shifting the pattern of publication lags embodied in D_{mt} recursively back in time. That is, observation $x_{i,\tau-h}$ is eliminated in D_{mT}, if and only if observation $x_{i,mT-h}$ is missing in D_{mT}. Thus, the unbalanced data set uses the original pattern of publication lags as from 30 June 2008.

As in Zadrozny (2008) and Clements and Galvao (2008) the series are normalized (demeaned and divided by the standard deviation).[6]

7.3.2 Design of the Case Study

We conduct a recursive forecasting exercise so that the model is specified and re-estimated on an ever-increasing sample size using the vintage of data available at the time the of the forecast. For each time period, we estimate the factors with the EM algorithm outlined in chapter 2. The basic vintage ranges from 1991:I to 1999:IV. We report the forecasting performance for one to four quarters ahead ($h = 1, \dots, 4$). Furthermore we outline the performance of the models with information that becomes available inside the quarter (nowcasting). Suppose that we are in December 1999 and we want to forecast GDP growth in the first quarter of 2000 (data are available in March). We use all available information in December to calculate the projection ($h = 1$). Moving to January, new values for the indicators and the factors become

[3] The finding of Schumacher and Breitung (2008) is in contrast to Koenig, Dolmas, and Piger (2003) who find that real-time data matter for the United States.

[4] The different publication lags can be considered as a real-time aspect of the data. But we use only final data, that is the data used are not subject to revision.

[5] www.bundesbank.de

[6] As already noted, standardization speeds up computation considerably, whereas structural relationships are not affected by standardization.

available. The projection for the current quarter is updated. We denote this forecast horizon with $h = 2/3$, as two months of observations are missing in the current quarter. For $h = 1/3$, only one month of observations is missing.

To specify the number of factors, we follow Marcellino and Schumacher (2007) and consider two approaches. First, we determine the number of static and dynamic factors, r and q respectively, using information criteria from Bai and Ng (2002) and Bai and Ng (2007).[7] Second, we compute all possible combinations of r and q and evaluate them in the forecasting exercise. In our application, we consider a maximum of $r = 6$ and all combinations of r and q with $q \leq r$. For reasons of brevity we show only the results for the best combination. In line with the literature in Banerjee and Marcellino (2006) and Stock and Watson (2002b), only a few factors proved to be useful for forecasting.

To investigate the forecasting performance of the mixed-frequency VAR and the MIDAS model, we estimate some benchmark models. In the univariate (quarterly-frequency) case we fit an AR(p) model where the lag length is determined by the Schwarz criterion (BIC). In the multivariate case we estimate a bivariate VAR(p). Furthermore we estimate bridge models. Bridge equations build a bridge over the gap between monthly and quarterly frequencies. As new in-sample information becomes available, the remaining missing values are forecast within a separate time series model (mostly an AR(p) model). Numerous bridge models have been used in applied work.[8] We model a standard ARX(p, r) model as the bridge equation, where the indicator enters the equation with a contemporaneous value. The missing observations within a quarter are forecast via an AR(p) model. The AR, ARX, VAR and MF-VAR models produce iterated (or plug-in) forecasts, whereas MIDAS produces direct forecasts. The estimation of the benchmark models are also useful to see how much cross-variable feedback contributes to the forecasts.

The mixed-frequency VAR models are estimated in MATLAB via maximum

[7] See also Marcellino and Schumacher (2007) for details.
[8] See chapter 3 for details.

likelihood. For model selection we use the corrected Akaike Information criterion reported in Hurvich and Tsay (1989) and used in our Monte Carlo study. We selected the model which minimized the criterion over a range of $p_{max} = 4$ when the residuals showed no significant serial correlation. For the MIDAS model we allow for a maximum of $5m = 15$ months of lags. In the Monte Carlo study we demonstrated that the exponential Almon lag and the Beta weighting functions produce in some cases different forecasting results. Furthermore, in chapter 4 we outlined that restrictions may be an option. Therefore we restrict the exponential Almon lag weighting function to ensure declining weights, whereas the Beta weighting function is estimated unrestricted.

For each model class and horizon we report the RMSE as defined in the Monte Carlo study. The tabulation of RMSE for all models and forecasting combinations does not allow us to conclude whether the better RMSE results from a model i compared with model j (or indicator) are statistically significant, that is systematic. In the literature the Diebold-Mariano test (DM) (Diebold and Mariano (1995)) is widely used to test for systematic differences in the forecasting errors. It should be noted, that the standard asymptotic theory for the Diebold-Mariano test statistics is invalid whenever two models involved are nested (among others see Clark and McCracken (2001)). A model 1 is nested in model 2 if model 2 reduces to model 1 when some parameters in model 2 are set to zero. We employ the test statistic proposed by Clark and West (2007) which accounts for nested models and employs a non-standard limiting distribution. The authors state that this test can also be used for non-nested models.

Assume that model 1 is the parsimonious model. Model 2 is the larger model that nests model 1. Clark and West (2007) suggest running the following regression:

$$\hat{f}_{t+h} = (y_{t+h} - \hat{y}_{1t,t+h})^2 - [(y_{t+h} - \hat{y}_{2t,t+h})^2 - (\hat{y}_{1t,t+h} - \hat{y}_{2t,t+h})^2]. \tag{7.1}$$

This tests for equal mean-squared prediction errors by regressing $\hat{f}_{t+h}$ on a constant and using the resulting t-statistic for a zero coefficient. We reject if this statistic is greater than +1.282 (for a one-sided 0.10 test) or +1.645 (for a one-sided 0.05 test). For the one-step-ahead forecasting errors, the usual least squares error can be used. For autocorrelated forecasting errors, we use the Newey-West heteroscedasticity-autocorrelation consistent standard errors.

7.3.3 Results

We start by comparing integer forecast horizons ($h = 1, \ldots, 4$) for each indicator within each model class. Table 7.1 reports the RMSEs for each integer forecast horizon for both quarterly and yearly growth rates. The lowest RMSE for each horizon is highlighted in bold face. The results are interesting in several ways. The expected strict increase of the RMSE with an increasing forecast horizon cannot be detected for the mixed-frequency models. In some cases the RMSE decreases going from one horizon to the other. For example the RMSE for the MF-VAR model with the ZEW indicator decreases from 0.619 for $h = 3$ to 0.587 for $h = 4$. We have no plausible explanation for these results. Apparently, the forecasts for different horizons exhibit different dynamics, which are captured by the mixed-frequency approaches.

Comparing the AR and VAR model we can state that the chosen indicators have predictive power, as the obtained RMSE from the VAR model are lower than the from the AR model at all horizons and for both target variables. If we compare the different approaches with one indicator, we can state that the mixed-frequency models always produce lower RMSE than AR and VAR. This statement is only valid for the quarterly growth rates. In case of yearly growth rates only the Ifo indicator in combination with mixed-frequency models outperforms the AR and VAR benchmark, whereas the ZEW and the factor obtain higher RMSEs. Given the overall assessment, the MF-VAR model performs best. The comparison of Almon and Beta weighting functions remains inconclusive. There are cases where Almon is better than Beta (using the same indicator) and vice versa. The Ifo exhibits very good predictive power in combination with the MF-VAR. Five out of the eight lowest RMSEs over different horizons and target variables are obtained by MF-VAR(Ifo). Comparing the factor results with the Ifo and ZEW indicators we can conclude that there are only some cases where the factor approach is better than the other two indicators.

Tables 7.3 to 7.6 in the appendix to this chapter report the Clark and West (2007) test statistics for each model/indicator combination. The tables are read as follows: a t-statistic greater than $+1.282$ (for a one-sided 0.10 test) or $+1.645$ (for a one-sided 0.05 test) indicates that Model 2 (rows) has a significantly smaller RMSE than Model 1 (columns) and vice versa. For example, consider quarterly growth rates and compare the AR model (Model 1) with the MF-VAR model with Ifo as an indicator (Model 2). If we look in the first column and the 11th row we can read a test statistic of 3.071 indicating that the RMSE of MF-VAR are statistically smaller than the AR model. The tables demonstrate that the mixed-frequency models are usually statistically significantly better than their single-frequency benchmarks. Moreover, the MF-VAR outperforms the MIDAS models statistically significantly in many cases.

Now we go a step further and allow for intra-quarterly information. Table 7.2 provides the corresponding results. Instead of the VAR model we now provide the RMSE of the bridge equation approach. A forecast horizon of $h = 1/3$ means that only one observation of the quarter is missing, $h = 2/3$ that two months of information are missing and so forth. We want to focus on the nowcasting aspect, that is $h = 1/3$ and $h = 2/3$. Again, there is no strict increase in the RMSEs with an increasing forecast horizon. First we can confirm the *a priori* expected result that the more information that is available the more accurate will be the forecasts. For the quarterly growth rates, the MIDAS model provides the lowest RMSEs for intra-quarterly forecasts. For yearly growth rates, we find a different result. In this case the bridge equation approach (with the ZEW indicator) provides the greatest forecasting accuracy. But this difference is not always statistically significant. The corresponding Clark and West (2007) test statistics for $h = 1/3$

and $h = 2/3$ can be found in the Tables 7.7 and 7.8. Our results differ to the conclusions of Clements and Galvao (2008) who find that in the short run there is little to choose between the bridge and MIDAS approaches. We find that in the short run the mixed-frequency approaches do not necessarily outperform the popular bridge approach.

The forecasting results from the factor approach are ambiguous. For short-term forecasting the survey indicators provide lower forecast errors independently of the model class used. Our results are in line with Dreger and Schumacher (2004) who find that factor models do not lead to statistically lower RMSE's when compared with single survey indicators. Given these results, we can answer our question in the introduction to this chapter: it is not worthwhile collecting many time series for forecasting when timely indicators are available, at least according to this example.

Tab. 7.1: Forecasting comparisons for German GDP growth

	AR(p)	LF-VAR(p)			MIDAS Almon			MIDAS Beta			MF-VAR(p)		
Quarterly Growth Rates													
		Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor
$h=1$	0.705	0.685	0.657	0.658	0.557	0.588	0.616	0.557	0.591	0.623	**0.514**	0.605	0.604
$h=2$	0.672	0.650	0.706	0.644	0.662	0.610	0.583	0.645	0.610	0.668	0.611	**0.602**	0.622
$h=3$	0.691	0.687	0.664	0.602	0.663	0.631	0.661	0.671	0.634	0.623	0.619	0.619	**0.600**
$h=4$	0.706	0.711	0.665	0.644	0.667	0.644	0.646	0.657	0.644	0.638	**0.612**	0.587	0.614
Yearly Growth Rates													
		Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor
$h=1$	0.656	0.517	0.551	0.558	0.514	0.787	0.824	0.510	0.789	0.665	0.521	**0.503**	0.806
$h=2$	0.820	0.635	0.744	0.752	0.601	0.777	0.852	0.601	0.777	0.737	**0.521**	0.706	0.813
$h=3$	0.929	0.757	0.856	0.860	0.779	0.781	0.811	0.779	0.781	0.831	**0.546**	0.714	0.772
$h=4$	0.996	0.938	0.927	0.890	0.904	0.845	0.827	0.874	0.845	0.840	**0.630**	0.724	0.789

Notes: This table displays the RMSE for different indicators and models. The numbers in bold face denote the lowest RMSE for each horizon.

Tab. 7.2: Forecasting comparisons for German GDP growth: Accounting for intra-quarterly information

	Quarterly Growth Rates											
	Bridge-Equation			MIDAS Almon			MIDAS Beta			MF-VAR(p)		
	Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor
$h = 1/3$	0.677	0.643	0.642	**0.500**	0.520	0.641	0.500	0.542	0.557	0.522	0.581	0.567
$h = 2/3$	0.652	0.677	0.642	0.531	0.552	0.623	0.531	0.573	0.621	**0.504**	0.607	0.574
$h = 1$	0.699	0.694	0.635	0.557	0.588	0.616	0.557	0.591	0.623	**0.514**	0.605	0.604
$h = 4/3$	-	-	-	0.642	0.617	0.693	0.623	0.617	0.670	**0.553**	0.592	0.566
$h = 5/3$	-	-	-	0.655	0.614	0.599	0.615	0.614	0.575	**0.537**	0.624	0.566
$h = 2$	-	-	-	0.662	0.610	0.583	0.645	0.610	0.668	0.611	**0.602**	0.622
$h = 7/3$	-	-	-	0.671	0.601	0.663	0.681	0.601	0.740	0.579	0.589	**0.571**
$h = 8/3$	-	-	-	0.686	0.622	0.714	0.672	0.622	0.752	0.606	0.602	**0.571**
$h = 3$	-	-	-	0.663	0.631	0.661	0.671	0.634	0.623	0.619	0.619	**0.600**
$h = 10/3$	-	-	-	0.666	0.638	0.656	0.669	0.644	0.656	0.593	0.613	**0.574**
$h = 11/3$	-	-	-	0.669	0.655	0.688	0.677	0.644	0.604	**0.573**	0.628	0.574
$h = 4$	-	-	-	0.667	0.644	0.646	0.657	0.644	0.638	0.612	**0.587**	0.614
	Yearly Growth Rates											
	Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor
$h = 1/3$	0.473	**0.413**	0.499	0.527	0.752	0.725	0.526	0.780	0.634	0.426	0.512	0.696
$h = 2/3$	0.461	**0.418**	0.491	0.543	0.792	0.817	0.534	0.788	0.700	0.479	0.545	0.774
$h = 1$	0.429	**0.415**	0.481	0.514	0.787	0.824	0.510	0.789	0.665	0.521	0.503	0.806
$h = 4/3$	-	-	-	0.538	0.766	0.730	0.538	0.774	0.694	**0.419**	0.629	0.732
$h = 5/3$	-	-	-	0.595	0.779	0.900	0.595	0.790	0.745	**0.419**	0.671	0.755
$h = 2$	-	-	-	0.601	0.777	0.852	0.601	0.777	0.737	**0.521**	0.706	0.813
$h = 7/3$	-	-	-	0.731	0.729	0.836	0.731	0.729	0.776	**0.463**	0.622	0.708
$h = 8/3$	-	-	-	0.765	0.747	0.868	0.765	0.747	0.833	**0.435**	0.640	0.713
$h = 3$	-	-	-	0.779	0.781	0.811	0.779	0.781	0.831	**0.546**	0.714	0.772
$h = 10/3$	-	-	-	0.872	0.821	0.905	0.829	0.821	0.904	**0.556**	0.673	0.713
$h = 11/3$	-	-	-	0.898	0.829	0.913	0.825	0.829	0.874	**0.562**	0.647	0.714
$h = 4$	-	-	-	0.904	0.845	0.827	0.874	0.845	0.840	**0.630**	0.724	0.789

Notes: This table displays the RMSE for different indicators and models. The forecast horizon ranges from $h = 1/3$ (one month of information is missing to forecast current GDP growth) to $h = 4$. The numbers in bold face denote the lowest RMSE for each horizon.

7.4 Tracking the German Economy on a Daily Basis: A NOWcasting Experiment

The results from the previous sections are quite promising, as in-quarterly information improves the estimate for the current quarter. It is natural to ask: why not use even higher-frequency data, such as weekly and daily data? In this section we want to investigate, whether daily data are useful to nowcast the current state of the economy. Additionally we ask whether the daily data help us to obtain more accurate forecasts for longer horizons, that is do daily data contain any long-run information? One possible advantage of daily data is that they are not subject to revision and are available early (there are no publication lags). A possible disadvantage is, that daily data are more erratic and exhibit volatile clustering. It is well known that financial data exhibit a strong (G)ARCH behaviour (see Bera and Higgins (1993)).

In contrast to monthly data, there are no daily indicators constructed to lead or interpreted to contemporaneously describe a specific low-frequency time series representing the state of the economy (GDP or industrial production). Thus, it is rather unlikely that one specific daily time series leads to more accurate forecasts than forecasts generated with monthly indicators.[9] Therefore, as in the previous section, we extract static and dynamic daily factors and plug them into the mixed-frequency time series models. Here another advantage of factor models becomes evident. The factor approach can possibly eliminate movements which are likely to include measurement errors and local shocks. This yields a more reliable signal for policy makers and prevents them from reacting to idiosyncratic movements.

In this section we focus on the MIDAS approach with exponential Almon and Beta weighting functions and leave out the mixed-frequency VAR model for several reasons. First, the MIDAS approach is far more parsimonious in handling such large frequency mixtures. As the true data generating process is unknown, the approach is likely to estimate many parameters within

[9] We experiment with some daily time series from our data set and find indeed, that this is the case.

the mixed-frequency VAR approach. Second, and even more important, as the mixed-frequency VAR assumes that the process operates at the highest frequency (daily), we would have to forecast up to a maximum of 359 days ahead.[10] For stationary transformations of variables, there exists a maximum horizon beyond which forecasts can provide no more information about the variable than is present in the unconditional mean. This point, called the 'content horizon' by Galbraith (2003), varies substantially across data series. Therefore we would expect to forecast only the conditional mean beyond some forecast horizon. And third, from the experience with the model, the larger the frequency-mixture the higher the computational burden. Thus, a recursive forecasting exercise would not be feasible in a reasonable time.

In this section we want to forecast German GDP growth as well as Industrial Production (IP), sampled monthly. IP has often been the target variable in forecasting exercises.[11] In almost all cases the leading indicators were also sampled monthly. Using mixed-frequency time series models we can investigate whether daily data can help to improve the forecasting accuracy of monthly IP.

7.4.1 The Daily Data Set

The prerequisite for the application of the MIDAS model is to have equidistant data. Therefore we have to adjust the data set as the quarters are defined by different numbers of days.[12] We define that a year consists of 360 days such that each quarter is defined by 90 days ($m = 90$). For forecasting the monthly IP we define a month consisting of 30 days ($m = 30$). Abstracting for a moment from a leap year, 365 days constitute a year. The first quarter contains exactly 90 days. For the second quarter we have 91 days.

[10] Assuming we are in the first days of the year and we want to forecast four quarters ahead.

[11] See Robinzonov and Wohlrabe (2008) for a literature review for German IP.

[12] Ghysels, Sinko, and Valkanov (2007) state that MIDAS models can handle unequally space data. They propose that instead of using the lag operator $L^{1/m}$ to use an operator L^{τ}, where τ is real-valued instead of a rational number. But this approach has not been applied to real data so far.

And the third and fourth quarters consist of 92 days. To obtain 90 days per quarter we (arbitrarily) delete the following days: 1 May, 31 July, 31 August, 25 and 26 December. In the case of a leap year (1992, 1996, 2000, 2004) we also delete 1 January. For the monthly case where the months have 31 days, we delete the last day. The missing values for February are forecast with an AR(p) process. We are aware of the fact that this is rather *ad hoc* and that we destroy sample information, but we do not think is will significantly influence our results.

Our daily data were obtained from the Deutsche Bundesbank and the IMF database, downloaded on the 12 August 2008. Finally we can use 61 daily series, among them interest rates, exchange rates and the equity price index (see the appendix to this chapter for details). We match the daily series on a grid of 360 days per year. We have many gaps in the different series due to weekends, holidays and so forth. Finally, we end up with $360 \times 15 = 5400$ (possible) observations.

Originally we downloaded 268 daily time series from both data bases. Our intention was to apply the EM algorithm to extract factors from a large data set with missing observations. Unfortunately, due to the time series length of 5400 matrix inversion with the EM algorithm became infeasible due to the large time series length.[13] Thus we linearly interpolated the missing observations for the 61 daily time series.[14] The remaining 201 daily time series either ended before December 2006 or started later than January 1992. Therefore we had to discard them, as linear interpolation was not applicable in these cases.

7.4.2 The Forecasting Set-up

The forecasting set-up is basically the same as in the monthly case. The first forecast is based on data from 1999:360 (the last day of the year). Then we recursively enlarge the information set on a daily basis. On each day we

[13] Even on a PC with Intel Core 2 Quad and 4GB RAM.
[14] Giannone, Reichlin, and Small (2008) linearly interpolated quarterly GDP on a monthly basis within the factor approach and claim that it works quite well.

forecast the next four available GDP values. For instance suppose we have a day in February, then we forecast the four quarters of the year. Being in June we forecast the remaining three quarters of the current year and the first quarter of the next year. Factors are extracted as outlined for the monthly case. For the MIDAS estimation we allowed for $4 \times 90 = 360$ daily lags. We restricted the Almon weighting function to ensure declining weights, whereas the Beta weighting remains unrestricted. The reason is as we demonstrate in the previous chapters: both weighting functions often deliver similar results. Furthermore we experimented with an unrestricted Almon weighting function, but we obtained in some cases infinite forecasts due to the numbers being too large in the weighting function.

7.4.3 Results

In contrast with the monthly case we only present the results graphically. Figures 7.1 and 7.2 graph the results for quarterly and yearly growth rates respectively. The graphs can be read as follows: the blue and red lines correspond to the average RMSE of the exponential Almon and Beta weighting function for each daily forecast horizon. The forecast horizon (x-axis) ranges from one day to 360 days. The thick vertical line marks 90 days, which is the current quarter. The two horizontal lines denote the RMSE of the best model/indicator combination for $h = 1/3$ (best available monthly short-term forecast) and $h = 4$ (best long-term forecast) from the monthly forecasting exercise in the previous section.

In the ideal stylized case one would expect a decreasing RMSE for a decreasing forecast horizon. This implies a declining line from the upper-right to the lower-left corner of the figures. If daily data provide lower RMSEs than the their monthly counterparts, the lines should be below the horizontal lines. For short-term forecasting all RMSEs left of the vertical line should be below the lower horizontal line. For long-term forecasting the blue and red lines for $h > 270$ should be below the upper horizontal line.

The results for the quarterly growth rates are depicted in Figure 7.1. They demonstrate that the ideal expectation cannot be met. First, we can state the RMSE are erratic. Sometimes the RMSE jumps more than one standard deviation up or down from one day to another. Although there is no typical clustering, as in many high-frequency time series, the erratic behaviour seems to transmit to the RMSEs. Despite the erratic figures there is a decreasing linear trend for the short-term forecast horizon. Nevertheless, the monthly forecast are not outperformed. Both the exponential Almon lag and the Beta weighting function exhibit RMSEs which are about 0.1 higher than the best monthly short-term forecast (MIDAS model with the Ifo Indicator).

Similar results are found for yearly growth rates, as shown in Figure 7.2. The short- and long-term monthly forecasts are not outperformed by the daily forecasts. There is one notable difference to the quarterly growth rates. The RMSE produced by the Beta weighting function exhibits a clear downward trend and is less erratic than its Almon counterpart. This could be due to the restriction imposed on the exponential Almon lag weighting function, whereas the Beta counterpart is estimated unrestricted.

Figures 7.3 and 7.4 plot the corresponding results for German Industrial Production. The vertical line denotes now 30 days. The competitive monthly forecasts are generated with a VAR(p) with the three described monthly indicators. Generally the results are similar to the GDP results: standard monthly models are not outperformed in the short run.

Figure 7.3 plots the RMSE for monthly IP growth rates. In contrast to GDP there is no downward trend of the RMSE. They fluctuate around a mean. This is the main difference to the GDP results. For longer forecast horizons the RMSE of the daily forecasts are lower than the best monthly forecast (VAR with Factors). It seems that the daily data in this case contain long-run information.

A similar result can be stated for the yearly growth rates plotted in Figure 7.4. In this case we find a counter-intuitive result. The RMSE of both weighting functions decrease with an increasing forecast horizon. So far we do not have an explanation for this result.

Fig. 7.1: Tracking the German economy on a daily basis: Daily forecasts of quarterly GDP growth rates

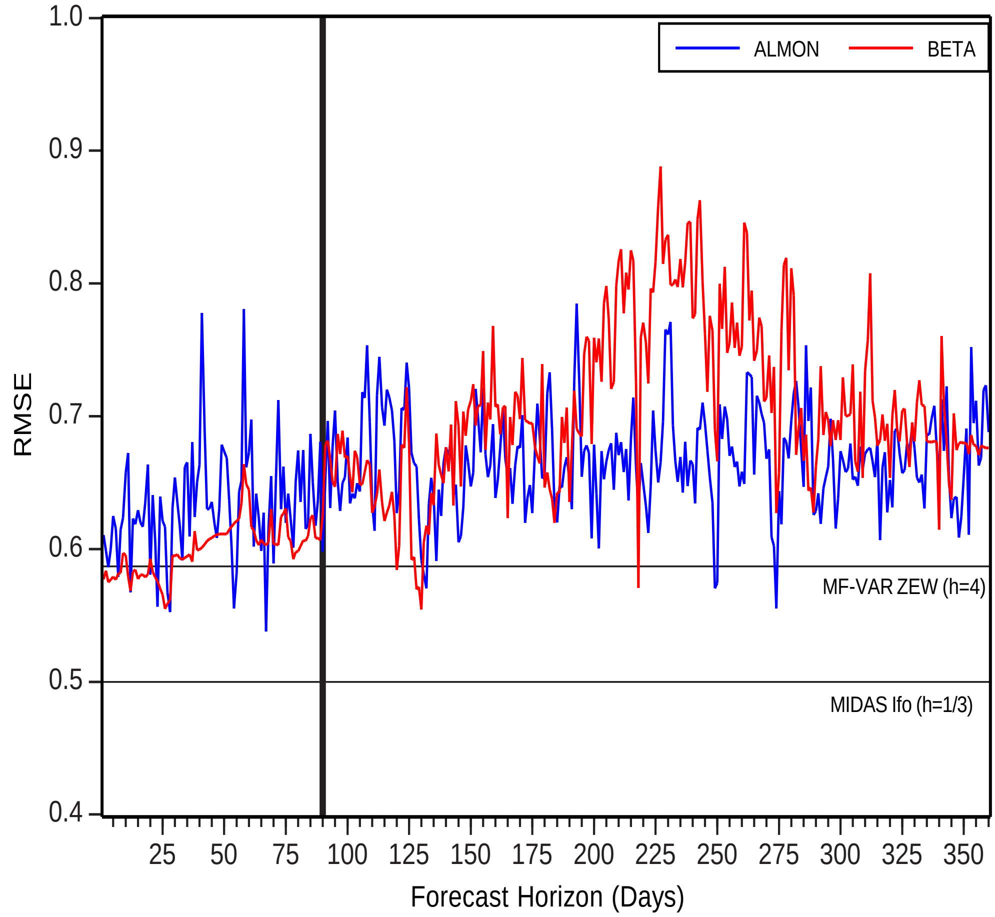

Fig. 7.2: Tracking the German economy on a daily basis: Daily forecasts of yearly GDP growth rates

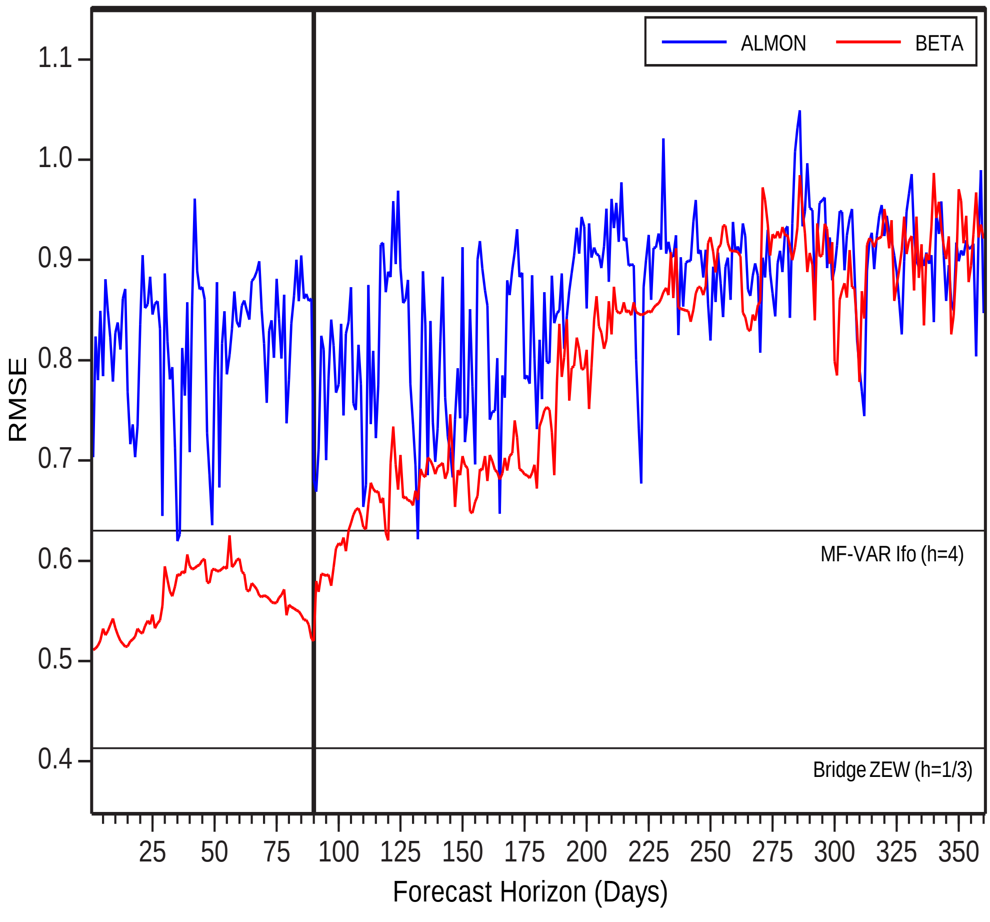

Fig. 7.3: Tracking the German economy on a daily basis: Daily forecasts of monthly IP growth rates

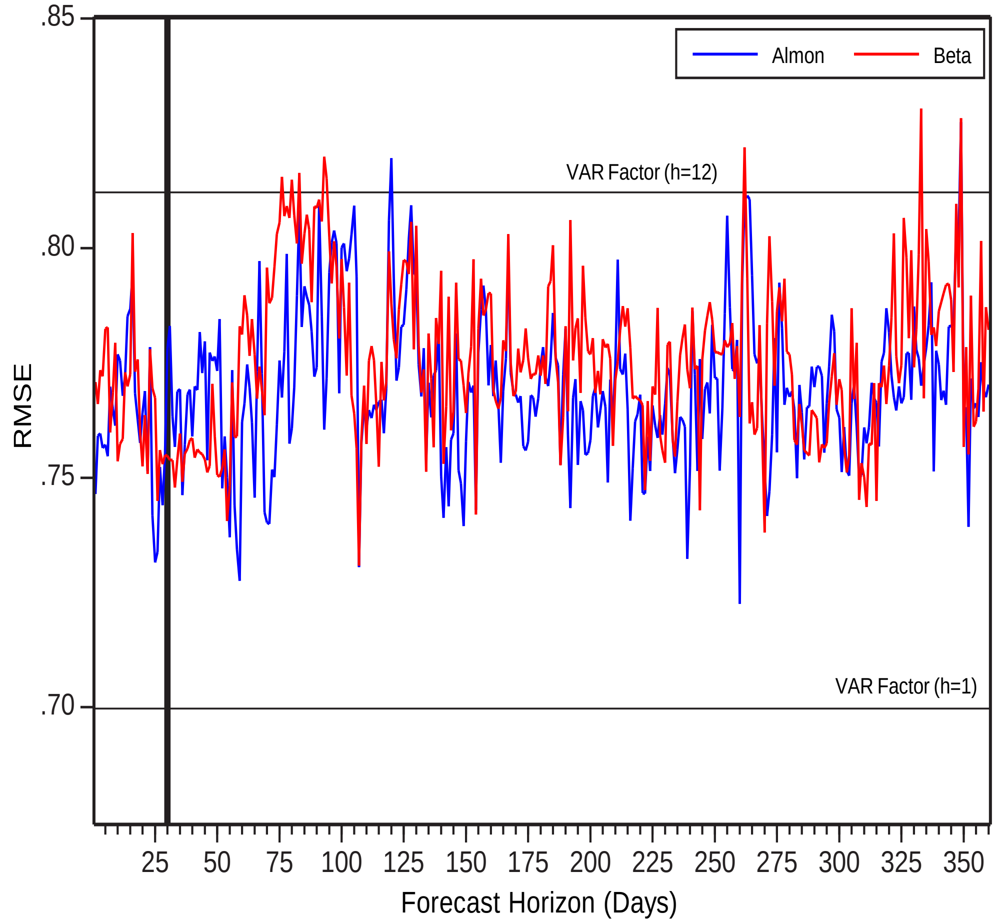

Fig. 7.4: Tracking the German economy on a daily basis: Daily forecasts of yearly IP growth rates

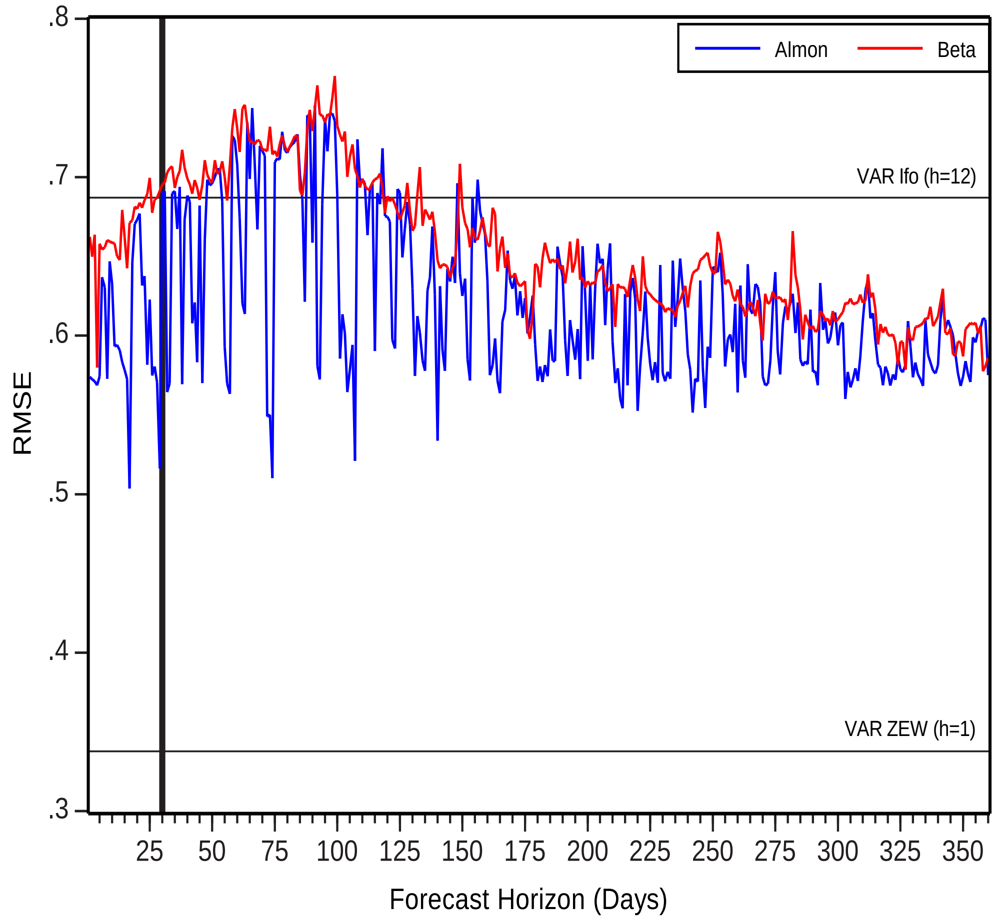

7.5 Summary and Discussion

In this chapter we compared the forecasting performances of the mixed-frequency time series models using real economic data. We started with forecasting quarterly German GDP growth with monthly indicators. We found for that for quarterly growth rate independently of the indicators, mixed-frequency models outperform single-frequency models (AR and VAR). This statement even holds for longer forecast horizons and is confirmed by statistical tests. Therefore we confirm that temporal aggregation destroys sample information and leads to the deterioration of forecasting accuracy. On average, the mixed-frequency VAR performs better than its MIDAS competitor.

Moving from monthly to daily predictors does not lead more accurate forecasts. We are the first to forecast GDP and IP growth on a daily basis. We extracted static and dynamic factors from a daily data set which is matched to an artificial year consisting of 360 days. The missing observations were interpolated. The extracted factors are plugged into the MIDAS model to generate forecasts. The obtained RMSEs exhibit erratic behaviour but with a decreasing tendency for lower forecasting horizons. We think that the erratic behaviour is transmitted from the (noisy) daily data via the extracted factors.

In the short run the daily forecasts do not outperform monthly generated forecasts. In case of GDP, the daily update was not better than the monthly counterpart at any forecasting horizon. In the case of industrial production, the daily data seem to contain long-run information, as the RMSE was lower than the one obtained from monthly mixed-frequency models.

The results are disappointing, but we think that this can be a starting point for fruitful future research. We feel confident that the factor approach for condensing a large information set is the best way to proceed, as we do not believe that one daily time series can be made operable to forecast time series such as IP and GDP. Daily updated forecasts may be better than monthly ones when the following issues can be resolved:

- A larger data base should be used. Many (important) daily time series had to be discarded as they were shorter than the estimation period. Linear interpolation was not applicable in these cases. To extract factors from time series with an arbitrary structure, one needs a computable algorithm able to handle long time series and many missing observations. Furthermore the computation should be done in a reasonable time.[15] The EM algorithm was not applicable as matrix inversion was not feasible.

- The factor extraction approach should account for the ARCH dynamics inherently founded in many daily time series. The erratic behaviour seems to transmit to the calculated accuracy measures.

- We employed the MIDAS approach to calculate the forecasts. It is easy to implement and the estimation with NLS proceeds very fast, but the MIDAS approach should be revised to allow for non-equidistant data. State-space approaches are more flexible but they are too computationally intensive at the moment. Moreover, calculations of long-run forecasts (more than 100 days ahead) with state-space models are not feasible at the moment.[16]

Suppose that we can resolve these issues, it nevertheless may be that lower-frequency (monthly) forecasts are not outperformed by higher-frequency forecasts. For example, daily data remain erratic even after some transformation, and thus the daily forecast update is also erratic. If and how central banks (for example) can account for this information should be discussed in future research. It will be interesting also to discuss the issue of how far we can go by using disaggregated data for forecasting lower-frequency variables. Aruoba, Diebold, and Scotti (2008) claim to push the factor approach to its high-frequency limit (daily basis). But maybe there is a frequency limit where we cannot improve forecasting accuracy for lower-frequency variables.

[15] The importance of this issue was also noted by Aruoba, Diebold, and Scotti (2008).
[16] In this case only the conditional mean would be forecast.

7.6 Appendix

This appendix states the additional tables and describes the daily and monthly time series for the German economy used in the forecasting and nowcasting exercise.

7.6.1 Additional Tables

Tab. 7.3: Comparisons of RMSE: Statistical Tests for $h = 1$

		AR(p)	LF-VAR(p)			MIDAS Almon			MIDAS Beta			MF-VAR(p)		
Quarterly Growth Rates		AR	Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor
AR														
LF-VAR(p)	Ifo	2.177												
	ZEW	1.662	1.133											
	Factor	2.088	2.746	1.586										
MIDAS Almon	Ifo	2.665	3.292	2.069	3.144									
	ZEW	2.072	3.092	1.748	2.863	-0.207								
	Factor	1.844	2.068	1.548	2.522	-1.189	0.395							
MIDAS Beta	Ifo	2.665	3.292	2.069	3.144	0.493	2.554	3.328						
	ZEW	2.052	3.103	1.737	2.845	-0.260	-1.498	1.780	-0.260					
	Factor	1.954	2.836	1.959	2.555	-0.386	0.580	1.164	-0.385	0.658				
MF-VAR(p)	Ifo	3.071	3.214	2.191	3.767	1.983	3.268	3.845	1.983	3.264	4.094			
	ZEW	2.672	2.723	1.808	2.832	-0.303	0.268	1.178	-0.303	0.322	2.390	-1.421		
	Factor	2.619	2.458	1.885	2.405	-0.889	0.772	1.006	-0.889	0.806	2.743	-1.780	1.127	
Yearly Growth Rates		AR	Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor
AR														
LF-VAR(p)	Ifo	2.892												
	ZEW	1.860	1.209											
	Factor	1.709	2.023	1.084										
MIDAS Almon	Ifo	3.428	1.883	2.835	3.067									
	ZEW	1.623	0.358	0.394	0.555	0.024								
	Factor	0.226	-0.528	-0.412	-0.222	-0.334	0.198							
MIDAS Beta	Ifo	3.413	1.918	2.876	3.103	1.416	2.962	2.683						
	ZEW	1.524	-0.004	0.025	0.387	-0.106	0.512	1.209	-0.205					
	Factor	1.592	0.498	0.500	0.519	0.659	2.254	1.485	0.545	2.822				
MF-VAR(p)	Ifo	4.549	1.608	2.197	2.400	-0.051	2.687	3.021	-0.146	2.272	1.945			
	ZEW	2.971	1.459	1.345	1.461	1.779	2.472	3.020	1.715	2.325	1.889	2.585		
	Factor	0.605	-0.394	-0.177	0.043	-0.231	0.527	1.636	-0.272	0.595	0.382	0.400	-2.154	

Notes: This table contains the results of the statistical test of equal RMSE of two models proposed by Clark and West (2007). A t-statistic greater than +1.282 (for a one sided 0.10 test) or +1.645 (for a one sided 0.05 test) indicates that Model 2 (rows) has a significant smaller RMSE than Model 1 (columns) and vice versa. Heteroscedastic and autocorrelation robust standard errors (Newey-West) are computed.

Tab. 7.4: Comparisons of RMSE: Statistical Tests for $h = 2$

		AR(p)	LF-VAR(p)			MIDAS Almon			MIDAS Beta			MF-VAR(p)		
Quarterly Growth Rates														
		AR	Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor
AR														
LF-VAR(p)	Ifo	0.411												
	ZEW	0.606	1.173											
	Factor	1.237	1.470	1.433										
MIDAS Almon	Ifo	-0.042	1.016	1.037	0.145									
	ZEW	0.595	1.425	1.174	0.226	0.602								
	Factor	1.337	1.563	1.745	0.744	1.669	1.515							
MIDAS Beta	Ifo	0.037	0.725	0.884	0.354	0.269	0.378	-0.046						
	ZEW	0.601	1.418	1.217	0.243	0.609	0.279	0.094	0.832					
	Factor	-0.535	0.780	0.088	-0.085	-0.941	-0.622	-0.669	0.070	-0.659				
MF-VAR(p)	Ifo	1.671	1.847	2.513	2.115	2.108	1.976	1.746	1.293	2.001	3.643			
	ZEW	1.786	1.982	3.682	1.567	1.856	2.940	1.692	1.568	2.999	2.629	0.273		
	Factor	1.513	1.659	2.584	1.305	2.156	2.024	1.382	1.173	2.054	3.940	-0.623	-0.200	
Yearly Growth Rates														
		AR	Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor
AR														
LF-VAR(p)	Ifo	3.128												
	ZEW	1.691	0.462											
	Factor	1.578	1.385	1.090										
MIDAS Almon	Ifo	2.776	1.159	1.734	1.904									
	ZEW	1.413	-0.682	-1.080	0.420	0.222								
	Factor	0.971	0.137	0.159	-0.587	-0.096	0.821							
MIDAS Beta	Ifo	2.776	1.159	1.734	1.904	-0.710	1.875	2.325						
	ZEW	1.367	-0.662	-1.181	0.453	0.200	0.108	1.693	0.200					
	Factor	1.133	-0.092	0.244	-0.537	-0.278	0.955	1.430	-0.278	0.944				
MF-VAR(p)	Ifo	2.852	1.752	1.732	2.142	2.262	1.739	2.605	2.261	1.765	2.059			
	ZEW	3.108	1.270	1.305	1.561	0.874	1.768	2.053	0.874	1.799	2.147	1.263		
	Factor	1.948	0.267	0.606	0.525	-0.142	1.165	0.933	-0.142	1.184	1.499	-0.026	-1.216	

Notes: This table contains the results of the statistical test of equal RMSE of two models proposed by Clark and West (2007). A t-statistic greater than $+1.282$ (for a one sided 0.10 test) or $+1.645$ (for a one sided 0.05 test) indicates that Model 2 (rows) has a significant smaller RMSE than Model 1 (columns) and vice versa. Heteroscedastic and autocorrelation robust standard errors (Newey-West) are computed.

Tab. 7.5: Comparisons of RMSE: Statistical Tests for $h = 3$

		AR(p)	LF-VAR(p)			MIDAS Almon			MIDAS Beta			MF-VAR(p)		
Quarterly Growth Rates		AR	Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor
AR														
LF-VAR(p)	Ifo	0.167												
	ZEW	1.743	1.678											
	Factor	2.168	1.915	2.145										
MIDAS Almon	Ifo	1.189	3.409	0.641	-0.546									
	ZEW	2.302	2.218	2.632	-0.841	1.033								
	Factor	1.624	2.300	0.270	-1.177	0.188	-1.579							
MIDAS Beta	Ifo	0.220	1.300	0.026	-0.675	-0.966	-0.782	0.028						
	ZEW	2.412	2.123	2.595	-0.756	1.084	0.444	2.035	1.442					
	Factor	1.133	1.920	0.232	-0.927	0.484	-0.749	0.158	1.035	-0.672				
MF-VAR(p)	Ifo	2.380	2.290	3.575	0.342	1.412	1.516	2.093	1.621	1.337	2.726			
	ZEW	2.890	2.894	3.559	0.535	2.353	1.583	2.278	2.229	1.607	2.207	0.893		
	Factor	2.799	2.500	4.154	0.377	1.545	2.375	2.499	1.742	1.988	2.630	0.647	1.173	
Yearly Growth Rates		AR	Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor
AR														
LF-VAR(p)	Ifo	1.995												
	ZEW	0.735	-0.309											
	Factor	1.429	1.329	1.168										
MIDAS Almon	Ifo	1.852	1.434	1.462	1.064									
	ZEW	2.013	1.124	1.627	0.653	0.155								
	Factor	2.049	1.003	2.488	1.013	0.709	1.438							
MIDAS Beta	Ifo	1.852	1.434	1.462	1.064	0.814	1.137	0.663						
	ZEW	2.013	1.124	1.627	0.653	0.155	-2.669	0.466	0.155					
	Factor	1.478	0.521	1.773	0.427	-0.514	-0.177	-1.110	-0.514	-0.177				
MF-VAR(p)	Ifo	2.457	2.279	1.800	2.580	2.350	2.010	1.802	2.350	2.010	2.548			
	ZEW	2.648	2.069	1.700	3.025	2.205	2.332	1.644	2.205	2.332	3.290	0.674		
	Factor	2.376	1.617	1.552	2.195	1.591	2.133	1.283	1.591	2.133	3.292	-0.168	-1.673	

Notes: This table contains the results of the statistical test of equal RMSE of two models proposed by Clark and West (2007). A t-statistic greater than $+1.282$ (for a one sided 0.10 test) or $+1.645$ (for a one sided 0.05 test) indicates that Model 2 (rows) has a significant smaller RMSE than Model 1 (columns) and vice versa. Heteroscedastic and autocorrelation robust standard errors (Newey-West) are computed.

Tab. 7.6: Comparisons of RMSE: Statistical Tests for $h = 4$

		AR(p)	LF-VAR(p)			MIDAS Almon			MIDAS Beta			MF-VAR(p)		
Quarterly Growth Rates														
		AR	Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor
AR														
LF-VAR(p)	Ifo	-0.290												
	ZEW	1.791	2.120											
	Factor	1.933	1.940	0.826										
MIDAS Almon	Ifo	0.493	2.484	-0.564	-0.519									
	ZEW	0.913	2.275	-0.587	-0.578	0.864								
	Factor	1.803	1.970	-0.203	-0.713	0.843	0.825							
MIDAS Beta	Ifo	1.175	2.323	0.314	0.137	1.203	1.137	0.441						
	ZEW	0.897	2.404	-0.454	-0.446	0.942	0.507	0.030	0.169					
	Factor	1.323	1.877	0.712	0.675	1.378	1.129	0.897	1.280	1.156				
MF-VAR(p)	Ifo	3.116	2.207	2.447	2.140	1.805	1.973	2.403	1.992	1.903	2.564			
	ZEW	2.005	3.116	2.183	1.523	2.328	2.276	1.533	2.665	2.342	4.164	0.956		
	Factor	2.744	2.224	2.693	2.190	1.689	1.844	2.423	1.899	1.784	2.853	0.202	-0.197	
Yearly Growth Rates														
		AR	Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor
AR														
LF-VAR(p)	Ifo	1.207												
	ZEW	0.513	0.093											
	Factor	1.586	1.999	1.300										
MIDAS Almon	Ifo	1.641	1.623	1.382	-0.043									
	ZEW	1.648	1.676	1.401	0.445	0.460								
	Factor	2.036	1.487	1.563	0.040	-0.117	-0.225							
MIDAS Beta	Ifo	2.369	2.005	1.827	0.732	0.660	0.415	1.739						
	ZEW	1.648	1.676	1.401	0.445	0.460	0.736	1.009	0.208					
	Factor	1.935	1.432	1.516	-0.045	-0.261	-0.381	-1.174	-0.400	-0.381				
MF-VAR(p)	Ifo	2.104	2.216	1.724	2.163	1.953	2.336	1.891	1.654	2.336	1.924			
	ZEW	2.191	2.430	1.730	3.450	2.763	2.570	2.773	2.006	2.570	2.798	1.258		
	Factor	1.986	2.142	1.580	2.661	2.101	2.273	2.351	1.480	2.273	2.402	0.592	-2.305	

Notes: This table contains the results of the statistical test of equal RMSE of two models proposed by Clark and West (2007). A t-statistic greater than +1.282 (for a one sided 0.10 test) or +1.645 (for a one sided 0.05 test) indicates that Model 2 (rows) has a significant smaller RMSE than Model 1 (columns) and vice versa. Heteroscedastic and autocorrelation robust standard errors (Newey-West) are computed.

Tab. 7.7: Comparisons of RMSE: Statistical Tests for $h = 1/3$

		Bridge Equation			MIDAS Almon			MIDAS Beta			MF-VAR(p)		
Quarterly Growth Rates		Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor
Bridge Equation	Ifo												
	ZEW	1.837											
	Factor	2.362	2.233										
MIDAS Almon	Ifo	4.219	2.841	3.958									
	ZEW	3.191	2.644	3.217	0.166								
	Factor	3.890	1.727	2.750	2.077	1.696							
MIDAS Beta	Ifo	4.219	2.841	3.958	1.857	2.358	2.509						
	ZEW	3.120	2.351	3.175	0.052	-0.318	2.734	0.052					
	Factor	4.236	2.004	3.215	1.538	1.268	2.129	1.538	1.392				
MF-VAR(p)	Ifo	3.422	2.545	3.220	0.217	1.931	2.368	0.217	2.449	1.783			
	ZEW	2.364	2.084	2.246	-0.370	-0.633	2.871	-0.370	0.019	1.456	0.447		
	Factor	2.720	2.212	2.306	-0.153	0.593	2.579	-0.153	0.798	1.502	2.174	1.684	
Yearly Growth Rates		Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor
Bridge	Ifo												
	ZEW	1.551											
	Factor	0.136	-2.219										
MIDAS Almon	Ifo	-0.681	1.644	1.911									
	ZEW	-1.489	-2.285	-1.522	0.850								
	Factor	0.568	0.338	1.667	2.149	1.548							
MIDAS Beta	Ifo	-0.691	1.628	1.884	0.466	2.672	1.851						
	ZEW	-1.958	-2.540	-2.128	1.125	1.219	0.868	1.128					
	Factor	-0.199	-0.386	0.730	1.646	2.279	1.424	1.666	1.996				
MF-VAR(p)	Ifo	1.116	1.765	2.050	1.556	2.665	1.823	1.626	2.190	2.096			
	ZEW	0.665	-0.082	1.403	2.037	2.648	1.648	2.065	2.356	1.796	1.045		
	Factor	-0.713	-1.369	-0.466	1.292	1.195	1.149	1.295	1.210	0.397	-0.300	-1.666	

Notes: This table contains the results of the statistical test of equal RMSE of two models proposed by Clark and West (2007). A t-statistic greater than $+1.282$ (for a one sided 0.10 test) or $+1.645$ (for a one sided 0.05 test) indicates that Model 2 (rows) has a significant smaller RMSE than Model 1 (columns) and vice versa. Heteroscedastic and autocorrelation robust standard errors (Newey-West) are computed.

Tab. 7.8: Comparisons of RMSE: Statistical Tests for $h = 2/3$

		Bridge Equation			MIDAS Almon			MIDAS Beta			MF-VAR(p)		
Quarterly Growth Rates		Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor
Bridge Equation	Ifo												
	ZEW	1.235											
	Factor	2.170	2.304										
MIDAS Almon	Ifo	3.732	2.638	4.131									
	ZEW	3.242	2.374	3.296	0.091								
	Factor	2.586	2.567	1.856	-1.190	-0.718							
MIDAS Beta	Ifo	3.732	2.638	4.131	0.615	2.276	3.094						
	ZEW	3.084	2.113	3.095	-0.295	-1.022	1.890	-0.295					
	Factor	2.189	1.977	2.034	-0.715	-0.045	2.095	-0.715	0.202				
MF-VAR(p)	Ifo	4.126	2.723	3.982	1.058	2.591	2.829	1.058	2.942	3.970			
	ZEW	1.864	1.505	2.044	-1.203	-1.450	1.509	-1.203	-0.698	1.980	-1.020		
	Factor	2.541	1.806	2.060	-0.943	0.187	1.591	-0.944	0.555	2.690	-0.588	2.503	
Yearly Growth Rates		Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor	Ifo	ZEW	Factor
Bridge Equation	Ifo												
	ZEW	1.564											
	Factor	0.080	-2.189										
MIDAS Almon	Ifo	-0.014	2.141	2.261									
	ZEW	-2.001	-2.350	-1.706	0.564								
	Factor	-1.900	-1.969	-1.538	-0.146	0.108							
MIDAS Beta	Ifo	0.098	2.119	2.264	1.603	2.679	2.800						
	ZEW	-2.181	-2.644	-2.143	0.896	1.490	2.080	0.741					
	Factor	-0.966	-1.254	-0.938	1.725	2.050	1.925	1.562	2.447				
MF-VAR(p)	Ifo	0.694	1.967	2.176	1.425	2.523	2.770	1.332	2.151	1.846			
	ZEW	0.549	-0.655	1.067	2.431	2.300	2.609	2.365	2.304	2.167	2.294		
	Factor	-1.288	-1.437	-0.890	0.278	0.648	1.394	0.182	0.694	0.741	-0.337	-1.308	

Notes: This table contains the results of the statistical test of equal RMSE of two models proposed by Clark and West (2007). A t-statistic greater than $+1.282$ (for a one sided 0.10 test) or $+1.645$ (for a one sided 0.05 test) indicates that Model 2 (rows) has a significant smaller RMSE than Model 1 (columns) and vice versa. Heteroscedastic and autocorrelation robust standard errors (Newey-West) are computed.

7.6.2 The Monthly Data Set

We employ and extended the same data set as used in Marcellino and Schumacher (2007). The monthly data set for Germany contains 111 time series over the sample period 1992:01 until 2006:12. The source of the time series is the Bundesbank database. The download date is 21 June 2008. In this data set there are missing values at the end of the sample.

Natural logarithms were taken for all time series except interest rates. Stationarity was enforced by appropriately differencing the time series. Most of the downloaded time series are already seasonally adjusted. Remaining time series with seasonal fluctuations were adjusted using Census-X12 prior to the forecasting exercise. We also corrected for outliers. Large outliers are defined as observations that differ from the sample median by more than six fold the sample inter-quartile range (see Watson (2003)). The identified observation is set equal to the respective outside boundary of the inter-quartile.

Included monthly time series

Money market rates Overnight money Monthly average
Money market rates One-month funds Monthly average
Money market rates Three-month funds Monthly average
Money market rates Six-month funds Monthly average
Money market rates Twelve-month funds Monthly average
Money market rates Fibor one-month funds Monthly average
Money market rates Fibor three-month funds Monthly average
Money market rates Fibor six-month funds Monthly average
Money market rates Fibor nine-month funds Monthly average
Money market rates Fibor twelve-month funds Monthly average
Money stock M3 3-month moving average (centred)
Money Stock M1
Money Stock M2
Money Stock M3
Turnover Intermediate goods
Turnover Capital goods

Turnover Durable goods
Turnover Non-durable goods
Turnover Industry
Turnover Consumer goods
Orders received Value
Orders received Value Intermediate goods
Orders received Value Capital goods
Orders received Value Consumer goods
Employment
Orders received Volume
Orders received Volume Intermediate goods
Orders received Volume Capital goods
Orders received Volume Consumer goods
Unemployment
Unemployment rate (unemployment as a percentage of the civilian labour force)
Job vacancies
Orders received Construction sector
Orders received Structural engineering
Orders received Housing construction
Orders received Industrial construction
Orders received Public sector construction
Orders received Civil engineering
Orders received Industrial clients
Orders received Public sector clients
Consumer price index
Consumer price index
Producer price index of farm products
Production Production sector including construction
Production Production sector excluding construction
Production Construction sector Total
Production Structural engineering
Production Civil engineering

Production Energy
Production Industry
Producer prices for industrial products (domestic sales)
Producer price index of industrial products / Total
Imports Price index
Exports Price index
Terms of trade
Retail turnover Value
Retail turnover Value Including retail of motor vehicles and including petrol stations
Retail turnover Volume
Retail turnover Volume Including retail of motor vehicles and including petrol stations
Yields on debt securities outstanding issued by residents Public debt securities
Yields on debt securities outstanding issued by residents Total
Yields on debt securities outstanding issued by residents Mortgage
Yields on debt securities outstanding issued by residents Public
CDAX price index / End 1987 = 100 End of month
Yields on debt securities outstanding issued by residents Debt securities issued by special purpose credit institutions / Monthly average
Yields on debt securities outstanding issued by residents Other bank debt securities / Monthly average
Yields on debt securities outstanding issued by residents Corporate bonds / Monthly average
Yield on foreign DM/EURO bonds outstanding issued by a German-managed syndicates / Monthly average
Yields on debt securities outstanding issued by residents / Listed Federal securities / Monthly average
CDAX performance index End 1987 = 100 End of month
REX price index End of month
REX performance index End 1987 = 100 End of month
Yields on debt securities outstanding issued by residents Mean residual ma-

turity of more than 1 and up to 2 years
Yields on debt securities outstanding issued by residents Mean residual maturity of more than 2 and up to 3 years
Yields on debt securities outstanding issued by residents Mean residual maturity of more than 3 and up to 4 years
Yields on debt securities outstanding issued by residents Mean residual maturity of more than 5 and up to 6 years
Yields on debt securities outstanding issued by residents Mean residual maturity of more than 6 and up to 7 years
Yields on debt securities outstanding issued by residents Mean residual maturity of more than 7 years /Monthly average
Yields on debt securities outstanding issued by residents Bank debt securities Mean residual maturity of more than 1 and up to 2 years
Yields on debt securities outstanding issued by residents Bank debt securities Mean residual maturity of more than 2 and up to 3 years
Yields on debt securities outstanding issued by residents Bank debt securities Mean residual maturity of more than 3 and up to 4 years
Yields on debt securities outstanding issued by residents Bank debt securities Mean residual maturity of more than 4 and up to 5 years
Yields on debt securities outstanding issued by residents Bank debt securities Mean residual maturity of more than 5 and up to 6 years
Yields on debt securities outstanding issued by residents Bank debt securities Mean residual maturity of more than 6 and up to 7 years
Yields on debt securities outstanding issued by residents Bank debt securities Mean residual maturity of more than 7 years
Yields on debt securities outstanding issued by residents Public debt securities Mean residual maturity of more than 1 and up to 2 years
Yields on debt securities outstanding issued by residents Public debt securities Mean residual maturity of more than 2 and up to 3 years
Yields on debt securities outstanding issued by residents Public debt securities Mean residual maturity of more than 3 and up to 4 years
Yields on debt securities outstanding issued by residents Public debt securities Mean residual maturity of more than 4 and up to 5 years

Yields on debt securities outstanding issued by residents Public debt securities Mean residual maturity of more than 5 and up to 6 years
Yields on debt securities outstanding issued by residents Public debt securities Mean residual maturity of more than 6 and up to 7 years
Yields on debt securities outstanding issued by residents Public debt securities Mean residual maturity of more than 7 years
Yields on debt securities outstanding issued by residents Bank debt securities
DAX price index End 1987 = 1000 End of month
DAX performance index End 1987 = 1000 End of month
Yields on debt securities outstanding issued by residents Listed Federal securities Residual maturity of more than 15 and up to 30 years
Price of gold in London / morning fixing 1 ounce of fine gold = USD ...
Yields on debt securities outstanding issued by residents Mean residual maturity of more than 7 and up to 8 years
Yields on debt securities outstanding issued by residents Mean residual maturity of more than 8 and up to 9 years
Yields on debt securities outstanding issued by residents Mean residual maturity of more than 9 and up to 10 years
Yields on debt securities outstanding issued by residents Public debt securities Mean residual maturity of more than 7 and up to 8 years
Yields on debt securities outstanding issued by residents Public debt securities Mean residual maturity of more than 8 and up to 9 years
Yields on debt securities outstanding issued by residents Public debt securities Mean residual maturity of more than 9 and up to 10 years
Yields on debt securities outstanding issued by residents Bank debt securities Mean residual maturity of more than 7 and up to 8 years
Yields on debt securities outstanding issued by residents Bank debt securities Mean residual maturity of more than 8 and up to 9 years
Yields on debt securities outstanding issued by residents Bank debt securities Mean residual maturity of more than 9 and up to 10 years
Yields on listed Federal securities (only bonds eligible as underlying instruments for future contracts are included
Yields on debt securities outstanding issued by residents Mortgage Pfand-

briefe Mean residual maturity of more than 1 and up to 2 years
Yields on debt securities outstanding issued by residents Mortgage Pfandbriefe Mean residual maturity of more than 2 and up to 3 years
Yields on debt securities outstanding issued by residents Mortgage Pfandbriefe Mean residual maturity of more than 3 and up to 4 years
Yields on debt securities outstanding issued by residents Mortgage Pfandbriefe Mean residual maturity of more than 4 and up to 5 years
Yields on debt securities outstanding issued by residents Mortgage Pfandbriefe Mean residual maturity of more than 5 and up to 6 years
Yields on debt securities outstanding issued by residents Mortgage Pfandbriefe Mean residual maturity of more than 6 and up to 7 years
Yields on debt securities outstanding issued by residents Mortgage Pfandbriefe Mean residual maturity of more than 7 years / Monthly average
Yields on debt securities outstanding issued by residents Mortgage Pfandbriefe Mean residual maturity of more than 7 and up to 8 years
Yields on debt securities outstanding issued by residents Mortgage Pfandbriefe Mean residual maturity of more than 8 and up to 9 years
Yields on debt securities outstanding issued by residents Mortgage Pfandbriefe Mean residual maturity of more than 9 and up to 10 years
Yields on debt securities outstanding issued by residents Public Pfandbriefe Mean residual maturity of more than 1 and up to 2 years
Yields on debt securities outstanding issued by residents Public Pfandbriefe Mean residual maturity of more than 2 and up to 3 years
Yields on debt securities outstanding issued by residents Public Pfandbriefe Mean residual maturity of more than 3 and up to 4 years
Yields on debt securities outstanding issued by residents Public Pfandbriefe Mean residual maturity of more than 4 and up to 5 years
Yields on debt securities outstanding issued by residents Public Pfandbriefe Mean residual maturity of more than 5 and up to 6 years
Yields on debt securities outstanding issued by residents Public Pfandbriefe Mean residual maturity of more than 6 and up to 7 years
Yields on debt securities outstanding issued by residents Public Pfandbriefe Mean residual maturity of more than 7 years

Yields on debt securities outstanding issued by residents Public Pfandbriefe Mean residual maturity of more than 7 and up to 8 years
Yields on debt securities outstanding issued by residents Public Pfandbriefe Mean residual maturity of more than 8 and up to 9 years
Yields on debt securities outstanding issued by residents Public Pfandbriefe Mean residual maturity of more than 9 and up to 10 years
Price of gold in London afternoon fixing 1 ounce of fine gold = USD ...
Term structure of interest rates on listed Federal securities (method by Svensson) residual maturity of 1 years
Term structure of interest rates on listed Federal securities (method by Svensson) residual maturity of 2 years
Term structure of interest rates on listed Federal securities (method by Svensson) residual maturity of 3 years
Term structure of interest rates on listed Federal securities (method by Svensson) residual maturity of 4 years
Term structure of interest rates on listed Federal securities (method by Svensson) residual maturity of 5 years
Term structure of interest rates on listed Federal securities (method by Svensson) residual maturity of 6 years
Term structure of interest rates on listed Federal securities (method by Svensson) residual maturity of 7 years
Term structure of interest rates on listed Federal securities (method by Svensson) residual maturity of 8 years
Term structure of interest rates on listed Federal securities (method by Svensson) residual maturity of 9 years
Term structure of interest rates on listed Federal securities (method by Svensson) residual maturity of 10 years

Ifo business situation capital goods producers
Ifo business situation producers durable goods
Ifo business situation producers non-durable goods
Ifo business situation retail trade

Ifo business situation wholesale trade
Ifo business expectations next six months capital goods producers
Ifo business expectations next six months producers durable goods
Ifo business expectations next six months producers non-durable goods
Ifo business expectations next six months retail trade
Ifo business expectations next six months wholesale trade
Ifo stocks of finished goods capital goods producers
Ifo stocks of finished goods producers durable goods
Ifo stocks of finished goods producers non-durable goods

7.6.3 The Daily Data Set

The daily data were downloaded on the 12 August 2008 from the Deutsche Bundesbank and the IMF data base. In contrast to the monthly data, we did not seasonally adjust the data. Natural logarithms were taken for all time series except interest rates. Stationarity was enforced by appropriately differencing the time series. We also corrected for outliers as in the monthly case.

Included daily time series

Money market rates reported by Frankfurt banks / Overnight money
Yields on debt securities outstanding issued by residents / Public debt securities
Yields on debt securities outstanding issued by residents
Yields on debt securities outstanding issued by residents / Mortgage Pfandbriefe
Yields on debt securities outstanding issued by residents / Public Pfandbriefe
Yields on debt securities outstanding issued by residents / Corporate bonds
Yield on foreign DM/EURO bonds outstanding issued by a German managed syndicate
Yields on debt securities outstanding issued by residents / Listed Federal securities

Yields on listed Federal securities (only bonds eligible as underlying instruments for future contracts are included)
Yields on debt securities outstanding issued by residents / Bank debt securities
Price of gold in London / morning fixing / 1 ounce of fine gold = USD ..
Price of gold in London / afternoon fixing / 1 ounce of fine gold = USD ...
Yields on debt securities outstanding issued by residents / Listed Federal securities / Residual maturity of more than 3 and up to 5 years
Yields on debt securities outstanding issued by residents / Listed Federal securities / Residual maturity of more than 5 and up to 8 years
Yields on debt securities outstanding issued by residents / Listed Federal securities / Residual maturity of more than 8 and up to 15 years
Yields on outstanding debt securities issued by residents / Listed Federal securities / Residual maturity of more than 15 and up to 30 years
DAX
Government Benchmarks, Bid, 10 Year, Yield, Close, EUR
Government Benchmarks, Bid, 2 Year, Yield, Close, EUR
Government Benchmarks, Bid, 3 Year, Yield, Close, EUR
Government Benchmarks, Bid, 7 Year, Yield, Close, EUR
Interest Rate Swaps, Ask, 1 Year, Close, EUR
Interest Rate Swaps, Ask, 10 Year, Close, EUR
Interest Rate Swaps, Ask, 2 Year, Close, EUR
Interest Rate Swaps, Ask, 3 Year, Close, EUR
Interest Rate Swaps, Ask, 30 Year, Close, EUR
Interest Rate Swaps, Ask, 4 Year, Close, EUR
Interest Rate Swaps, Ask, 5 Year, Close, EUR
Interest Rate Swaps, Ask, 6 Year, Close, EUR
Interest Rate Swaps, Ask, 7 Year, Close, EUR
Interest Rate Swaps, Ask, 8 Year, Close, EUR
Interest Rate Swaps, Ask, 9 Year, Close, EUR
Spot Rates, USD/DEM, Close, DEM
Deposit Rates, 1 Month, Close, EUR
Deposit Rates, 1 Week, Close, EUR

Deposit Rates, 12 Month, Close, EUR
Deposit Rates, 2 Month, Close, EUR
Deposit Rates, 3 Month, Close, EUR
Deposit Rates, 6 Month, Close, EUR
Deposit Rates, 9 Month, Close, EUR
Deposit Rates, T/N, Close, EUR
Interbank Rates, BBA LIBOR, 1 Month, Fixing, EUR
Interbank Rates, BBA LIBOR, 1 Week, Fixing, EUR
Interbank Rates, BBA LIBOR, 10 Month, Fixing, EUR
Interbank Rates, BBA LIBOR, 11 Month, Fixing, EUR
Interbank Rates, BBA LIBOR, 12 Month, Fixing, EUR
Interbank Rates, BBA LIBOR, 2 Month, Fixing, EUR
Interbank Rates, BBA LIBOR, 3 Month, Fixing, EUR
Interbank Rates, BBA LIBOR, 4 Month, Fixing, EUR
Interbank Rates, BBA LIBOR, 5 Month, Fixing, EUR
Interbank Rates, BBA LIBOR, 6 Month, Fixing, EUR
Interbank Rates, BBA LIBOR, 7 Month, Fixing, EUR
Interbank Rates, BBA LIBOR, 8 Month, Fixing, EUR
Interbank Rates, BBA LIBOR, 9 Month, Fixing, EUR
Interbank Rates, EONIA, O/N, Fixing, EUR
Interbank Rates, EURIBOR, 12 Month, Fixing, EUR
Interbank Rates, EURIBOR, 3 Month, Fixing, EUR
Policy Rates, ECB Lombard Rate (Ceiling), EUR

8. SUMMARY AND CONCLUSIONS

The continuous inflow of information concerning economic variables is closely monitored by central banks, governments and companies. Knowledge of the current and future states of the economy is essential as a basis for public and private decisions. However, there is no 'official' definition of the state of economy, nor is there a representative economic time series. Nevertheless, GDP is widely acknowledged as the most comprehensive indicator for the economic performance of countries. Therefore, possession of accurate estimates of current and future values of GDP is essential for economic agents. But there are considerable challenges: GDP is only sampled at quarterly intervals, published with delay, and then subject to several revisions.

Therefore, there is a requirement for methods which are able to estimate and forecast (final) GDP and other low-frequency variables by using higher-frequency information. The existing solutions in the literature are barely connected, not described, compared, and tested comprehensively against each other. In this thesis, we provide new research to fill this gap.

In a first step we outline the techniques capable of handling mixed-frequency data. The predominant approach in applied forecasting is to transform the data into a single frequency. Early approaches to combine different frequencies were linkage models and bridge equations. Linkage models pool forecasts from different frequencies to improve the forecasting accuracy of the lower frequency. They have often been used in the late 1980s and early 1990s. In contrast, bridge equations are still quite popular in applied work, especially at central banks. Bridge equations allow to update the current estimate of an economic variable as soon as new information becomes available. Yet it is essentially still a single-frequency approach.

Recently, two pure mixed-frequency models were developed: the state-space VARMA and the MIDAS (MIxed DAta Sampling) approach. We review these two techniques extensively and compare them against each other.

In an extensive literature review we link these approaches to the existing empirical applications. In general, we can conclude that mixed-frequency data matter for forecasting. Many articles demonstrate that accounting for timely information improves the forecasting accuracy of lower-frequency variables, notably the GDP. Recently, factor models have played a dominant role in forecasting. The extraction of factors allows forecasters to condense a large data set into a few factors and therefore to account for the aspect of parsimony in econometrics. Recent factor models also account for mixed-frequency data. Nevertheless, the focus is still on monthly data.

Currently, there are only few practical applications of the two mixed-frequency time series approaches. But they clearly demonstrate their advantage in forecasting over single-frequency approaches. In this thesis we confirm and generalize these early findings in a systematic way. Mixed-frequency approaches are feasible, provide more accurate forecasts, and are able to accommodate any given data structure.

As model selection is essential for any forecasting calculation, we investigate how to best specify mixed-frequency models. Although standard model selection criteria can in principle be applied, this aspect has been neglected in the literature to date. The estimation of mixed-frequency VAR models is computationally intensive, especially for models with many variables. Nevertheless, we find that small models selected on the basis of the BIC criterion are sufficiently accurate for forecasting purposes.

In the MIDAS context we elicit whether the weighting function should be restricted (for example to ensure declining weights) and how many lags should be included for estimation. We find that the number of included lags, the restriction aspect, and the forecast performance are interrelated. In some cases restrictions cause a deterioration in the forecasting performance. For strongly persistent target variables the restrictions may indeed improve forecasting accuracy. Applying the BIC criterion can determine the optimal number of lags included. Since the additional inclusion of lags does not increase the number of estimated parameters, there is a risk of including too many lags and thereby reducing the estimation sample. As a rule of thumb, we suggest using four or five times the frequency mixture as a starting point for lag selection. These results apply to both the exponential Almon lag and the Beta weighting function.

The evaluation of the appropriate model specification prepares the ground for a systematic forecasting comparison in our Monte Carlo study. Employing four data-generation processes displaying realistic economic data structures, we find that mixed-frequency time series models are at least as good as their single-frequency counterparts. In many cases the use of untransformed mixed-frequency data clearly outperforms the temporal aggregated data results. However, none of of the mixed-frequency approaches clearly dominates. For processes with strongly autoregressive components the MF-VAR model has an advantage over the MIDAS technique. Given the structure and length of standard macroeconomic time series, forecasts generated by both approaches are robust, and improvements in forecasting accuracy by increased sample size are rather small. Given the heteroscedastic nature of many economic and financial series, the relative gain in forecasting accuracy relative to single-frequency models is smaller compared to the homoscedastic case. The biggest advantage of mixed-frequency models is achieved by adjusting the information set. Allowing for intra-period information is only feasible for mixed-frequency models (except for bridge equations), and clearly improves forecasting accuracy. Relative gains up to 80 percent can be reached.

Therefore, the mixed-frequency VAR and the MIDAS approach provide useful tools for updating current estimates of macroeconomic variables, even these are published with delay and are subject to revisions.

We verify the finding from the Monte Carlo study in a real-data example using German GDP. Thus, we are the first to compare both approaches with real data. The Ifo Business Climate Index, the ZEW Index of Economic Sentiment, and factors extracted from a large database were used as indicators.

Independently of the indicators, the mixed-frequency VAR and the MIDAS model outperformed single-frequency benchmark models both in the short and long run.

Eventually we went a step further by replacing monthly with daily indicators. Static and dynamic factors are extracted from a daily database and used within the MIDAS framework as indicators. Although our approach does not outperform short-run monthly generated forecasts, we consider this experiment as a useful starting point for future research where daily or even higher-frequency data can help to nowcast the economy in almost real-time. We propose three steps for future research. First, one should account for non-equidistantly sampled time series. Second, the computational burden must be reduced considerably, since state-space models are able to handle any data patterns but they quickly become infeasible with large data sets. Third, one needs to account for heteroscedastic and erratic data. In our investigation we find that these patterns transmit to the forecast evaluation measures. Therefore they are not reliable in assessing the forecasting performance in almost real time.

This thesis can trigger the discussion of how far we can push the usage of high-frequency data for the purpose of forecasting. Maybe there is a frequency-mixture limit, where it is inadvisable to use still high-frequency data to generate forecasts. To phrase it differently: Are temporally aggregated forecasts are superior beyond a certain level of aggregation?

A further issue for future research is, whether the advantage of mixed-frequency models over single-frequency models still holds when the forecast horizon is increased. Our empirical example provides evidence in the affirmative.

BIBLIOGRAPHY

Abeysinghe, T. (1998): "Forecasting Singapore's quarterly GDP with monthly external trade," *International Journal of Forecasting*, 14(4), 505–513.

——— (2000): "Modeling variables of different frequencies," *International Journal of Forecasting*, 16(1), 117–119.

Altissimo, F., R. Cristadoro, M. Forni, M. Lippi, and G. Veronese (2006): "New Eurocoin: Tracking economic growth in real time," CEPR Discussion Paper No. 5633.

Amemiya, T., and R. Wu (1972): "The effect of aggregation on prediction in the autoregressive model," *Journal of the American Statistical Association*, 67(339), 628–632.

Andersen, T., T. Bollerslev, F. Diebold, and P. Labys (2003): "Modeling and forecasting realized volatility," *Econometrica*, 71(2), 579–625.

Andersen, T., T. Bollerslev, F. Diebold, and C. Vega (2003): "Micro effects of macro announcements: Real-time price discovery in foreign exchange," *American Economic Review*, 93(1), 38–62.

Anderson, B., and J. Moore (1979): *Optimal filtering*. Englewood Cliffs, NJ: Prentice Hall.

Andreou, E., E. Ghysels, and A. Kourtellos (2007): "Regression models with mixed sampling frequencies," *mimeo*, Chapel Hill, N.C.

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.

Ansley, C., and R. Kohn (1985): "Estimation, filtering, and smoothing in state space models with incompletely specified initial conditions," *Annals of Statistics*, 13(4), 1286–1316.

Aruoba, S., F. Diebold, and C. Scotti (2008): "Real-Time measurement of business conditions, Second version," PIER Working Paper Archive No. 11, Penn Institute for Economic Research, Department of Economics, University of Pennsylvania.

Ashley, R., C. Granger, and R. Schmalensee (1980): "Advertising and aggregate consumption: An analysis of causality," *Econometrica*, 48(5), 1149–1168.

Baffigi, A., R. Golinelli, and G. Parigi (2004): "Bridge models to forecast the Euro area GDP," *International Journal of Forecasting*, 20(3), 447–460.

Bai, J. (2003): "Inferential theory for factor models of large dimensions," *Econometrica*, 71(1), 135–171.

Bai, J., and S. Ng (2002): "Determining the number of factors in approximate factor models," *Econometrica*, 70(1), 191–221.

——— (2007): "Determining the number of primitive shocks in factor models," *Journal of Business and Economic Statistics*, 25(1), 52–60.

Banbura, M., and G. Runstler (2007): "A Look into the factor model black box: Publication lags and the role of hard and soft data in forecasting GDP," ECB Working Paper No. 751.

Banerjee, A., and M. Marcellino (2006): "Are there any reliable leading indicators for US inflation and GDP growth?," *International Journal of Forecasting*, 22(1), 137–151.

Bera, A., and M. Higgins (1993): "ARCH Models: Properties, estimation and testing," *Journal of Economic Surveys*, 7(4), 305–366.

Bernanke, B., and J. Boivin (2003): "Monetary policy in a data-rich environment," *Journal of Monetary Economics*, 50(3), 525–546.

Bernanke, B., M. Gertler, and M. Watson (1997): "Systematic monetary policy and the effects of oil price shocks," *Brookings Papers on Economic Activity*, 1997(1), 91–142.

Bodo, G., A. Cividini, and L. Signorini (1991): "Forecasting the Italian industrial production index in real time," *Journal of Forecasting*, 10(3), 285–99.

Bollerslev, T. (1986): "Generalized autoregressive conditional heteroskedasticity," *Journal of Econometrics*, 31(3), 307–327.

Boot, J., W. Feibes, and J. Lisman (1967): "Further methods of derivation of quarterly figures from annual data," *Applied Statistics*, 16(1), 65–75.

Braun, S. (1990): "Estimation of current-quarter Gross National Product by pooling preliminary labor market data," *Journal of Business and Economic Statistics*, 8(3), 293–304.

Breitung, J., and S. Eickmeier (2006): "Dynamic factor models," *Journal of the German Statistical Society*, 90(1), 27–40.

Camba-Mendez, G., G. Kapetanios, R. Smith, and M. Weale (2001): "An automatic leading indicator of economic activity: forecasting GDP growth for European countries," *The Econometrics Journal*, 4(1), 56–90.

Carrasco, M., and J. Florens (2000): "Generalization of GMM to a continuum of moment conditions," *Econometric Theory*, 16(6), 797–834.

Chen, B., and P. A. Zadrozny (1998): "An extended Yule-Walker method for estimating vector autoregressive models with mixed-frequency data," in *Advances in econometrics: Messay data, missing observations, outliers, and mixed-frequency data*, ed. by T. Fomby, and R. Hill, pp. 47–73. Greenwich, CT: JAI Press.

Chen, X., and E. Ghysels (2008): "News - good or bad - and its impact on volatility predictions over multiple horizons," *mimeo*, Chapel Hill, N.C.

Chow, G., and A. Lin (1971): "Best linear unbiased interpolation, distribution, and extrapolation of time series by related series," *The Review of Economics and Statistics*, 53(4), 372–375.

Chow, G., and A. Lin (1976): "Best linear unbiased estimation of missing observations in an economic time series," *Journal of the American Statistical Association*, 71(355), 719–721.

Christiano, L. (1989): "P*: Not the inflation forecaster's holy grail," *Federal Reserve Bank of Minneapolis Quarterly Review*, 13(4), 3–18.

Clark, T., and M. McCracken (2001): "Tests of equal forecast accuracy and encompassing for nested models," *Journal of Econometrics*, 105(11), 85–110.

Clark, T., and K. West (2007): "Approximetely normal tests for equal predictive accuracy in nested models," *Journal of Econometrics*, 138(1), 291–311.

Clements, M. P., and A. B. Galvao (2005): "Macroeconomic forecasting with mixed-frequency data: Forecasting US output growth and inflation," Warwick working paper.

——— (2008): "Macroeconomic forecasting with mixed-frequency data: Forecasting US output growth," *Journal of Business and Economic Statistics*, forthcoming.

Cohen, K., W. Müller, and M. Padberg (1971): "Autoregressive approaches to disaggregation of time series data," *Applied Statistics*, 20(2), 119–129.

Corrado, C., and M. Greene (1988): "Reducing uncertainty in short-term projections: Linkages of monthly and quarterly data," *Journal of Forecasting*, 7(2), 77–102.

Corsi, F. (2003): "A simple long memory model of realized volatility," Unpublished manuscript, University of Southern Switzerland.

Crone, T., and A. Clayton-Matthews (2005): "Consistent economic indexes for the 50 States," *Review of Economics and Statistics*, 87(4), 593–603.

Croushore, D., and T. Stark (2003): "A real-time data set for macroeconomists: Does the data vintage matter?," *Review of Economics and Statistics*, 85(3), 605–617.

Cuche, N., and M. Hess (2000): "Estimating monthly GDP in a general Kalman filter framework: Evidence from Switzerland," *Economic & Financial Modelling*, 7, 153–194.

Davis, E., and S. Henry (1994): "The use of financial spreads as indicator varialbes: Evidence for the United Kingdom and Germany," *IMF Staff Papers*, 41(3), 517–525.

Denton, F. (1971): "Adjustment of monthly or quarterly series to annual totals: An approach based on quadratic minimization," *Journal of the American Statistical Association*, 66(333), 99–102.

Dhrymes, P. (1971): *Distributed lags; problems of estimation and formulation.* Holden-Day.

Diebold, F., and G. Rudebusch (1991): "Forecasting output with the composite leading index: A real-time analysis," *Journal of the American Statistical Association*, 86(415), 603–610.

Diebold, F. X., and R. S. Mariano (1995): "Comparing predictive accuracy," *Journal of Business & Economic Studies*, 13(3), 253–263.

Diron, M. (2006): "Short-term forecasts of Euro area real GDP growth - An Assessment of real-time performance based on vintage data," ECB Working Paper No. 622.

Donihue, M., and E. Howrey (1992): "Using mixed frequency data to improve macroeconomic forecasts of inventory investment," *International Journal of Production Economics*, 26(1-3), 33–41.

Doz, C., D. Giannone, and L. Reichlin (2006): "A quasi maximum likelihood approach for large approximate dynamic factor models," ECB Working Paper No. 674.

Dreger, C., and C. Schumacher (2004): "Estimating large-scale factor models for economic activity in Germany: Do they outperform simpler models?," *Jahrbucher fur Nationalokonomie und Statistik*, 224(6), 731–750.

Dufour, J., and E. Renault (1998): "Short-run and long-run causality in time series: Theory," *Econometrica*, 66(5), 1099–1125.

Durbin, J., and S. Koopman (2001): *Time series analyses by state space methods*. Oxford University Press, New York.

Eickmeier, S., and C. Ziegler (2008): "How successful are dynamic factor models at forecasting output and inflation? A meta-analytic approach," *Journal of Forecasting*, 27(1), 237–265.

Engle, R., and V. Ng (1993): "Measuring and testing the impact of news on volatility," *Journal of Finance*, 48(5), 1749–1778.

Evans, M. (2005): "Where are we now? Real-time estimates of the macroeconomy," *International Journal of Central Banking*, 1(2), 127–175.

Fair, R., and R. Shiller (1990): "Comparing information in forecasts from econometric models," *The American Economic Review*, 80(3), 375–389.

FERNANDEZ, R. B. (1981): "A methodological note on the estimation of time series," *Review of Economics and Statistics*, 63(3), 471–476.

FITZGERALD, T., AND P. MILLER (1989): "A simple way to estimate current-quarter GNP," *Federal Reserve Bank of Minneapolis Quartely Review*, 13(Fall), 27–31.

FORSBERG, L., AND E. GHYSELS (2007): "Why do absolute returns predict volatility so well?," *Journal of Financial Econometrics*, 5(1), 31–67.

FRENCH, K., W. SCHWERT, AND R. STAMBAUGH (1987): "Expected returns and volatility," *Journal of Financial Economics*, 19(1), 3–29.

FRIEDMAN, M. (1962): "The interpolation of time series by related series," *Journal of the American Statistical Association*, 57(300), 729–757.

FROMM, G., AND L. KLEIN (1973): "A comparison of eleven econometric models of the United States," *American Economic Review*, 63(2), 385–93.

FUHRER, J., AND J. HALTMAIER (1988): "Minimum variance pooling of forecasts at different levels of aggregation.," *Journal of Forecasting*, 7(1), 63–73.

GALBRAITH, J. (2003): "Content horizons for univariate time-series forecasts," *International Journal of Forecasting*, 19(1), 43–55.

GHYSELS, E., P. SANTA-CLARA, AND R. VALKANOV (2004): "The MIDAS touch: MIxed DAta Sampling regression models," *mimeo*, Chapel Hill, N.C.

——— (2005): "There is a risk-return trade-off after all," *Journal of Financial Economics*, 76(3), 509–548.

——— (2006): "Predicting volatility: Getting the most out of return data sampled at different frequencies," *Journal of Econometrics*, 131(1-2), 59–95.

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

Ghysels, E., and R. Valkanov (2006): "Linear time series processes with mixed data sSampling and MIDAS regressions models," *mimeo*, Chapel Hill, N.C.

Ghysels, E., and J. Wright (2008): "Forecasting professional forecasters," *Journal of Business and Economic Statistics*, forthcoming.

Giacomini, R., and H. White (2006): "Tests of conditional predictive ability," *Econometrica*, 74(6), 1545–1578.

Giannone, D., L. Reichlin, and D. Small (2006): "Nowcasting GDP and inflation: The real-time informational content of macroeconomic data releases," ECB Working paper No. 633.

——— (2008): "Nowcasting: The real-time informational content of macroeconomic data," *Journal of Monetary Economics*, 55(4), 665–676.

Goldrian, G. (2007): *Handbook of survey-based business cycle analysis*. Edward Elgar.

Golinelli, R., and G. Parigi (2007): "The use of monthly indicators to forecast quarterly GDP in the short-run: An application to the G7 countries," *Journal of Forecasting*, 26, 77–94.

——— (2008): "Realtime squared: A real-time data set for real-time GDP forecasting," *International Journal of Forecasting*, 24(3), 368–385.

Gomez, V., and A. Maravall (1994): "Estimation, prediction, and interpolation for nonstationary series with the Kalman filter," *Journal of the American Statistical Association*, 89(426), 611–624.

Granger, C. (1993): "On the limitations of comparing mean squared forecast errors: Comment," *Journal of Forecasting*, 12(8), 651–652.

Greene, M., E. Howrey, and S. Hymans (1986): "The use of outside information in econometric forecasting," in *Model Reliability*, ed. by D. Belsely, and E. Kuh, pp. 90–116. MIT Press.

Hagen, H., and G. Kirchgässner (1996): "Interest based forecasts for german economic growth," *Weltwirtschaftliches Archiv*, 132(4), 763–773.

Hamilton, J. D. (1994): *Time series analysis*. Princeton University Press, Princeton, New Jersey.

Hannan, E. (1963a): "Regression for time series," in *Proceedings of a symposium on time series analysis*, ed. by M. Rosenblatt, pp. 17–37. Wiley.

——— (1963b): "Regression for time series with errors of measurement," *Biometrika*, 50(3-4), 293–302.

Hansen, P., and A. Lunde (2006): "Realized variance and market microstructure noise," *Journal of Business and Economic Statistics*, 24(2), 127–161.

Harvey, A., and R. Pierse (1984): "Estimating missing observations in economic time series," *Journal of the American Statistical Association*, 79(385), 125–131.

Hogrefe, J. (2008): "Forecasting data revisions of GDP: A mixed frequency approach," *Advances in Statistical Analysis*, 92(3), 271–296.

Howrey, E., S. Hymans, and M. Donihue (1991): "Merging monthly and quarterly forecasts: Experience with MQEM," *Journal of Forecasting*, 10(3), 255–268.

Hurvich, C., and C. Tsay (1989): "Regression and time series model selection in small samples," *Biometrika*, 76(2), 297–307.

Hyung, N., and C. Granger (2008): "Linking series generated at different frequencies," *Journal of Forecasting*, 27(2), 95–108.

JONES, R. (1980): "Maximum likelihood fitting of ARMA models to time series with missing observations," *Technometrics*, 22(3), 389–395.

JUDGE, G., W. GRIFFITH, C. HILL, H. LÜTKEPOHL, AND T. LEE (1985): *The theory and practice of econometrics.* Wiley: New York.

KAPETANIOS, G., AND M. MARCELLINO (2006): "A Parametric Estimation Method for Dynamic Factor Models of Large Dimensions," CEPR Discussion Paper No. 5620.

KHOLODILIN, K., AND B. SILIVERSTOVS (2006): "On the forecasting properties of the alternative leading indicators for the German GDP: Recent evidence," *Jahrbücher für Nationalökonomie und Statistik*, 226(3), 234–259.

KIRSCHGÄSSNER, G., AND M. SAVIOZ (2001): "Monetary policy and forecasts for real GDP growth: An empirical investigation for the Federal Republic of Germany," *German Economic Review*, 2(4), 339–365.

KITCHEN, J., AND R. MONACO (2003): "Real-time forecasting in practice: The US Treasury Staff's real-time GDP forecast system," *Business Economics*, pp. 10–19.

KLEIN, L., AND J. PARK (1993): "Economic forecasting at high-frequency intervals," *Journal of Forecasting*, 12(3-4), 301–319.

KLEIN, L., AND E. SOJO (1989): "Combinations of high and low frequency data in macroeconometric models," in *Economics in theory and practice: An eclectic approach*, ed. by L. Klein, and J. Marquez, pp. 3–16. Kluwer Academic Publishers.

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.

KOHN, R., AND C. ANSLEY (1986): "Estimation, prediction, and interpolation for ARIMA models with missing data," *Journal of the American Statistical Association*, 81(395), 751–761.

Kong, D., H. Liu, and L. Wang (2008): "Is there a risk-return trade-off? Evidences from Chinese stock markets," *Frontiers of Economics in China*, 3(1), 1–23.

Krämer, J., and E. Langfeld (1993): "Die Zinsdifferenz als Frühindikator für die westdeutsche Konjunktur," *Die Weltwirtschaft*, 1, 34–43.

Lanning, S. (1986): "Missing observations: A simultaneous approach versues interpolation by related series," *Journal of Economic and Social Measurement*, 14(1), 155–163.

León, Á., J. Nave, and G. Rubio (2007): "The relationship between risk and expected return in Europe," *Journal of Banking and Finance*, 31(2), 495–512.

Linton, O., and E. Mammen (2005): "Estimating semiparametric ARCH(∞) models by kernel smoothing methods," *Econometrica*, 73(3), 771–836.

Liou, R., and C. Shen (1996): "The use of high frequency data to improve macroeconometric forecast," *International Economic Journal*, 10(2), 65–83.

Lisman, J., and J. Sandee (1964): "Derivation of quarterly figures from annual data," *Applied Statistics*, 13(2), 87–90.

Litterman, R. (1983): "A random walk, Markov model for the distribution of time series," *Journal of Business and Economic Statistics*, 1(2), 169–173.

Liu, H., and S. Hall (2001): "Creating high-frequency national accounts with state-space modelling: A Monte Carlo experiment," *Journal of Forecasting*, 20(6), 441–449.

Liu, T. (1969): "A monthly recursive econometric model of United States: A test of feasibility," *Review of Economics and Statistics*, 51(1), 1–13.

Liu, T., and E. Hwa (1974): "A monthly econometric model of the US economy," *International Economic Review*, 15(2), 328–365.

Lütkepohl, H. (1987): *Forecasting aggregated vector ARMA processes.* Springer.

Mankiw, N., and M. Shapiro (1986): "News or noise? An analysis of GNP revisions," NBER Working Paper No. 1939.

Marcellino, M., and C. Schumacher (2007): "Factor nowcasting of German GDP with ragged-edge data. A model comparison using MIDAS projections," Bundesbank Discussion Paper, Series 1 No. 34.

Marcellino, M., J. Stock, and M. Watson (2006): "A comparison of direct and iterated multistep AR methods for forecasting macroeconomic zime series," *Journal of Econometrics*, 135(1-2), 499–526.

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.

McNees, S. (1986): "Forecasting accuracy of alternative techniques: A comparison of US macroeconomic forecasts," *Journal of Business and Economic Statistics*, 4(1), 5–15.

Meijering, E. (2002): "A chronology of interpolation: from ancient astronomy to modernsignal and image processing," *Proceedings of the IEEE*, 90(3), 319–342.

Merton, R. (1973): "An intertemporal capital asset pricing model," *Econometrica*, 41(5), 867–887.

Miller, P., and D. Chin (1996): "Using monthly data to improve quarterly model forecasts," *Federal Reserve Bank of Minneapolis Quarterly Review*, 20(2), 16–33.

Mitchell, J., R. Smith, M. Weale, S. Wright, and E. Salazar (2005): "An indicator of monthly GDP and an early estimate of quarterly GDP growth," *The Economic Journal*, 115(501), 108–129.

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. Wollmershäuser, pp. 19–48. Heidelberg: Physica-Verlag.

Nieto, F. (2007): "*Ex post* and *ex ante* prediction of unobserved multivariate time series: a structural-model based approach," *Journal of Forecasting*, 26(1), 53.

Nijman, T., and F. Palm (1990): "Parameter identification in ARMA processes in the presence of regular but incomplete sampling," *Journal of Time Series Analysis*, 11(3), 239–248.

Nunes, L. (2005): "Nowcasting quarterly GDP growth in a monthly coincident indicator model," *Journal of Forecasting*, 24(8), 575.

Parigi, G., and G. Schlitzer (1995): "Quarterly forecasts of the italian business cycle by means of monthly indicators," *Journal of Forecasting*, 14(2), 117–141.

Perez, J. (2007): "Leading indicators for euro area government deficits," *International Journal of Forecasting*, 23(2), 259–275.

Pesaran, H., and A. Timmermann (2005): "Real-time econometrics," *Econometric Theory*, 21(1), 212–231.

Proietti, T. (2006): "Temporal disaggregation by state space methods: Dynamic regression methods revisited," *Econometrics Journal*, 9(3), 357–372.

Proietti, T., and F. Moauro (2006): "Dynamic factor analysis with non-linear temporal aggregation constraints," *Journal of the Royal Statistical Society Series C*, 55(2), 281–300.

RATHJENS, P., AND R. ROBINS (1993): "Forecasting quarterly data using monthly information," *Journal of Forecasting*, 12(3-4), 321–330.

ROBERDS, W. (1988): "A quarterly Bayesian VAR model of the US economy," Working Paper No. 88-2, Federal Reserve Bank of Atlanta.

ROBINZONOV, N., AND K. WOHLRABE (2008): "Freedom of choice in macroeconomic forecasting: An illustration with German industrial production and linear models," ifo working paper No. 57, Munich.

RÜNSTER, G., AND F. SÉDILLOT (2003): "Short-term estimate of Euro area real GDP growth by means of monthly indicators," European Central Bank Working Paper No. 276.

SCHUMACHER, C., AND J. BREITUNG (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.

SEONG, B., S. K. AHN, AND P. A. ZADROZNY (2007): "Cointegration analysis with mixed-frequency data," CESifo Working Paper No. 1939, CESifo GmbH.

SHEN, C. (1996): "Forecasting macroeconomic variables using data of different periodicities," *International Journal of Forecasting*, 12(2), 269–282.

SILVESTRINI, A., AND D. VEREDAS (2008): "Temporal aggregation of univeriate and multivariate time series models: A survey," *Journal of Economic Surveys*, 22(3), 458–497.

STOCK, J., AND M. WATSON (1989): "New indexes of coincident and leading indicators," *NBER Macroeconomics Annual*, 4, 351–394.

——— (1991): "A probability model of the coincident economic indicators," in *Leading economic indicators: New approaches and forecasting records*, ed. by K. Lahiri, and G. Moore, pp. 63–85. Cambridge University Press.

——— (2002a): "Forecasting using principal components from a large number of predictors," *Journal of the American Statistical Association*, 97(460), 1167–1179.

——— (2002b): "Macroeconomic forecasting using diffusion indexes," *Journal of Business and Economic Statistics*, 20(2), 147–62.

——— (2003): "Forecasting output and inflation: The role of asset prices," *Journal of Economic Literature*, 41(3), 788–829.

Stram, D., and W. Wei (1986): "A methodological note on the disaggregation of time series totals," *Journal of Time Series Analysis*, 7(4), 293–302.

Tiao, G. (1972): "Asymptotic behaviour of temporal aggregates of time series," *Biometrika*, 59(3), 525–531.

Timmermann, A. (2006): "Forecast combinations," in *Handbook of Economic Forecasting*, ed. by C. Granger, G. Elliot, and A. Timmermann, pp. 135–196. Amsterdam, North Holland.

Trehan, B. (1989): "Forecasting growth in current quarter real GNP," *Federal Reserve Bank of San Francisco Economic Review*, (Winter), 39–51.

——— (1992): "Predicting contemporaneous output," *Federal Reserve Bank of San Francisco Economic Review*, (2), 3–11.

Trehan, B., and R. Ingenito (1996): "Using monthly data to predict quarterly output," *Federal Reserve Bank of San Francisco Economic Review*, (3), 3–11.

Watson, M. (2003): "Macroeconomic forecasting using many predictors," in *Advances in economics and econometrics, theory and applications, Eight World Congress of the Econometric Society*, ed. by H. L. Dewatripont, M., and S. Turnovsky, pp. 87–115. Cambridge University Press.

ZADROZNY, P. A. (1988): "Gaussian-Likelihood of countinuous-time ARMAX models when data are strocks and flows at different frequencies," *Econometric Theory*, 4(1), 108–124.

——— (1990): "Estimating a multivariate ARMA model with mixed-frequency data: An application to forecating U.S. GNP an monthly intervals," Federal Reserve Bank of Atlanta Working Paper Series No. 90-6.

——— (2008): "Estimating a multivariate ARMA model with mixed-frequency data: An application to forecating U.S. GNP at monthly intervals," *CESifo Working Paper*, (forthcoming).

Eidesstattliche Versicherung

Ich versichere hiermit eidesstattlich, dass ich die vorliegende Arbeit selbständig und ohne fremde Hilfe verfasst habe. Die aus fremden Quellen direkt oder indirekt übernommenen Gedanken sowie mir gegebene Anregungen sind als solche kenntlich gemacht. Die Arbeit wurde bisher keiner anderen Prüfungsbehörde vorgelegt und auch noch nicht veröffentlicht.

München, 23. September 2008

Klaus Wohlrabe

Curriculum Vitae

06/2004 – 09/2008	Ph.D. Student at the University of Munich, Junior Researcher at the Ifo Institute for Economic Research
09/1998 – 05/2004	Diploma in Economics, University of Dresden
07/1997	Abitur, Wirtschaftsgymnasium Dresden
17.12.1977	Born in Quedlinburg, Germany