

Forecasting with mixed-frequency data^{*}

Elena Andreou[†] Eric Ghysels[‡] Andros Kourtellos[§]

First Draft: January 2010

This Draft: May 30, 2010

^{*}This chapter was prepared for the *Oxford Handbook on Economic Forecasting* edited by Michael P. Clements and David F. Hendry. The authors would like to thank the Leventis foundation UCY grant 3411-32021. The first author also acknowledges support of the European Research Council under the European Community FP7/2007-2013 ERC grant 209116. We also benefitted from the valuable comments of two referees and Mike Clements.

[†]Department of Economics, University of Cyprus, P.O. Box 537, CY 1678 Nicosia, Cyprus, e-mail: elena.andreou@ucy.ac.cy.

[‡]Department of Economics, University of North Carolina, Gardner Hall CB 3305, Chapel Hill, NC 27599-3305, USA, and Department of Finance, Kenan-Flagler Business School, e-mail: eghysels@unc.edu.

[§]Department of Economics, University of Cyprus, P.O. Box 537, CY 1678 Nicosia, Cyprus, e-mail: andros@ucy.ac.cy

1 Introduction

Innovations in computer technology have made it possible to easily collect and store large data sets. One consequence of this is that many (financial) time series are recorded at very high sampling frequencies. Yet, many real activity series have maintained the traditional monthly or quarterly collection and release scheme. As a result, interest in forecasting with mixed-frequency data has emerged as an important topic.

Take, for example, the situation of macroeconomic forecasting involving a combination of past quarterly series and the choice between using past quarterly financial series - or instead using those same series sampled daily. Not using the readily available daily series has two important implications: (1) one loses information through temporal aggregation and (2) one foregoes the possibility of providing real-time daily, weekly or monthly updates of forecasts. Both topics have been addressed using state space models, which consist of a system with two types of equations, the measurement equations which link observed series to a latent state process, and the state equations which describe the state process dynamics. The setup treats the low-frequency data as “*missing data*” and the Kalman filter is a convenient computational device to extract the missing data.¹ The approach has many benefits, but also some drawbacks. State space models can be quite involved, as one must explicitly specify a linear dynamic model for all the series involved: the high-frequency data series, the latent high-frequency series treated as missing and the low-frequency observed processes. The system of equations therefore typically requires a lot of parameters, namely for the measurement equation, the state dynamics and their error processes. The steady state Kalman filter gain, however, yields a linear projection rule to (1) extract the current latent state, and (2) predict future observations as well as states. The Kalman filter can then be used to predict low frequency macro series, using both past high and low frequency observations. A number of recent papers also documented the gains of real-time forecast updating, sometimes also nowcasting when it applies to current quarter assessments.² These studies used again the state space setup.

An alternative approach to deal with data sampled at different frequencies has emerged

¹See for example, Harvey and Pierse (1984), Harvey (1989), Zdrozny (1990), Bernanke, Gertler, and Watson (1997), Mariano and Murasawa (2003), Mittnik and Zdrozny (2004), Aruoba, Diebold, and Scotti (2009), Bai, Ghysels, and Wright (2009), Kuzin, Marcellino, and Schumacher (2009), among others.

²Nowcasting is studied at length by Doz, Giannone, and Reichlin (2008), Doz, Giannone, and Reichlin (2006), Stock (2006), Angelini, Camba-Mendez, Giannone, Rünstler, and Reichlin (2008), Giannone, Reichlin, and Small (2008), Moench, Ng, and Potter (2009), among others.

in recent work by Ghysels, Santa-Clara, and Valkanov (2004), Ghysels, Santa-Clara, and Valkanov (2006) and Andreou, Ghysels, and Kourtellis (2010a), using so called MIDAS, meaning Mi(xed) Da(ta) S(ampling), regressions. It is a regression framework that is parsimonious - notably not requiring to model the dynamics of each and every daily predictor series - in contrast to the system of equations that require imposing many assumptions and estimating many parameters, for the measurement equation, the state dynamics and their error processes.³

The topic of mixing different sampling frequencies also emerges even when time series are available at the same frequency, but one is interested in multi-period forecasting. Take the example of an annual forecast with quarterly data. The first approach is to estimate a model with past annual data, and hence collapse the original multi-period setting into a single step forecast. The second approach is to estimate a quarterly forecasting model and then iterate forward the forecasts to a multi-period annual prediction. The forecasting literature refers to the first approach as “direct” and the second as “iterated”. (Marcellino, Stock, and Watson (2006)). Traditionally, the comparison has been made between direct and iterated forecasting, see e.g. Findley (1983), Findley (1985), Lin and Granger (1994), Clements and Hendry (1996), Bhansali (1999), and Chevillon and Hendry (2005)). Multi-period forecasts can also be constructed using a mixed-data sampling approach. A MIDAS model can use past quarterly data to produce directly multi-period forecasts. The MIDAS approach can be viewed as a middle ground between the direct and the iterated approaches. Namely, one preserves the past high frequency data, to directly produce multi-period forecasts.

There is a related literature on aggregation and forecasting in regression models (see, for instance, the surveys by Granger (1985) and Lütkepohl (2004) and more recent work by Hendry and Hubrich (2010), Hotta and Neto (2008) among others) as well as aggregation and volatility forecasting (see, for instance, the recent survey by Andersen, Bollerslev, Christoffersen, and Diebold (2006), and Ghysels and Sinko (2010)). While this literature recognizes the forecasting gains of disaggregation, the idea of using models where the variables are of mixed data sampling frequencies was first introduced in Ghysels, Santa-Clara, and Valkanov (2005) and since then there is a large and growing literature. Empirical applications involve regression and quantile regression models for forecasting macroeconomic variables as well as volatility models for understanding and forecasting financial risk.

The original work on MIDAS focused on volatility predictions; see for instance, Alper,

³See Armesto, Engemann, and Owyang (2010) for a user-friendly introduction to MIDAS regressions.

Fendoglu, and Saltoglu (2008), Chen and Ghysels (2009), Engle, Ghysels, and Sohn (2008), Forsberg and Ghysels (2006), Ghysels, Santa-Clara, and Valkanov (2004), León, Nave, and Rubio (2007), Clements, Galvão, and Kim (2008) among others. In addition a number of recent papers have documented the advantages of using such MIDAS regressions in terms of improving quarterly macro forecasts with monthly and daily data. For instance, Bai, Ghysels, and Wright (2009), Kuzin, Marcellino, and Schumacher (2009), Armesto, Hernandez-Murillo, Owyang, and Piger (2009), Clements and Galvão (2009), Clements and Galvão (2008), Galvão (2006), Schumacher and Breitung (2008), Tay (2007), use monthly data to improve quarterly forecasts. Similarly, quarterly and monthly macroeconomic predictions are improved by daily financial series, see e.g. Ghysels and Wright (2009), Hamilton (2006), Tay (2006), Monteforte and Moretti (2009), and Andreou, Ghysels, and Kourtellos (2010b).

The remainder of the chapter is structured as follows. In section 2 we cover MIDAS regressions. Section 3 covers so called nowcasting and relationship with the Kalman filter and its relationship with MIDAS regressions. A final section discusses volatility models using mixed frequencies.

2 MIDAS Regressions

Suppose we are interested in forecasting quarterly GDP growth rate, Y_{t+1}^Q , using daily stock returns, $X_{N_D-j,t}^D$, in the j^{th} day counting backwards in quarter t .⁴ Hence, the last day of quarter t corresponds to $j = 0$ and is therefore $X_{N_D-j,t}^D$. The conventional approach, in its simplest form, aggregates the data at the quarterly frequency by computing simple averages to obtain $X_t^Q = (X_{N_D,t}^D + X_{N_D-1,t}^D + \dots + X_{1,t}^D)/N_D$ and then estimates a simple Distributed Lag (DL) model

$$Y_{t+1}^Q = \alpha + \beta X_t^Q + u_{t+1}, \quad (2.1)$$

where α and β are unknown parameters and u_{t+1} is an error term. The implicit assumption in traditional models such as (2.1) is that temporal aggregation is based on an equal weighting

⁴For notational brevity, we will be dealing with one-step ahead forecasts. All the models and methods we will be presenting can be easily extended to multi-step forecasting.

scheme of the high frequency data. An alternative naive approach would estimate the model

$$Y_{t+1}^Q = \mu + \alpha Y_t^Q + \sum_{j=0}^{N_D-1} \beta_j X_{N_D-j,t}^D + u_{t+1}, \quad (2.2)$$

where N_D denotes the daily lags or the number of trading days per quarter. This is an unappealing approach because of parameter proliferation: when $N_D = 66$, we have to estimate not only 66 parameters β_j but also μ and α , hence a total of 68 slope coefficients.⁵

Instead, the MIDAS regression models use a parsimonious and data-driven aggregation scheme based on a low dimensional high frequency lag polynomial, $W(L^{N_D}; \theta)$ such that $W(L^{N_D}; \theta_X^D) X_t^D = \sum_{j=0}^{N_D-1} w_j(\theta_X^D) X_{t-j}^D$. There are various alternatives for the polynomial specification. Two flexible specifications that parameterize the weights into a two parameter vector include the two parameter exponential Almon lag and the Beta lag. Ghysels, Sinko, and Valkanov (2006) provide a discussion on the two specifications as well as for step-functions. The exponential Almon lag is specified as:

$$w_j(\theta_1, \theta_2) = \frac{\exp\{\theta_1 j + \theta_2 j^2\}}{\sum_{j=1}^m \exp\{\theta_1 j + \theta_2 j^2\}}. \quad (2.3)$$

with $\theta = (\theta_1, \theta_2)$. The Beta function is given by

$$w_j(\theta_1, \theta_2) = \frac{f(j, \theta_1; \theta_2)}{\sum_{j=1}^M f(j, \theta_1; \theta_2)} \quad (2.4)$$

where $f(x, \theta_1; \theta_2) = x^{\alpha-1}(1-x)^{\beta-1}\Gamma(\alpha+\beta)/(\Gamma(\alpha)\Gamma(\beta))$ and $\Gamma(a) = \int_0^\infty e^{-x} x^{a-1} dx$.

This approach yields a distributed lag model as a linear projection of high frequency data X_t^D onto Y_t^Q using a MIDAS filter, $DL - MIDAS(q_X^D)$

$$Y_{t+1}^Q = \mu + \beta \sum_{j=0}^{q_X^D-1} \sum_{i=0}^{N_D-1} w_{i+j*N_D}(\theta^D) X_{N_D-i,t-j}^D + u_{t+1}, \quad (2.5)$$

where the second summation allows for daily lags to extend beyond the last day of quarter t , but to simplify notation, we will always take lags in blocks of quarterly sets of daily data, q_X^D . Note that equation (2.5) nests the simple DL model in equation (2.1) under flat-

⁵Typically we have about 66 observations for many daily financial data over a quarter since each month has 22 trading days.

weights. To see this, note that, if $\theta_1 = \theta_2 = 0$ then the exponential Almon lag polynomial function yields equal/flat aggregation weights. We assume that $w_{i+j*N_D}(\theta^D) \in (0, 1)$ and $\sum_{j=0}^{q_X^D-1} \sum_{i=0}^{N_D-1} w_{i+j*N_D}(\theta^D) = 1$, that allows the identification of the slope coefficient β in the MIDAS regression model. The parameters (μ, β, θ^D) are estimated by Nonlinear Least Squares (NLS).

Our understanding of MIDAS regression can be further enhanced by decomposing the conditional mean in equation as the sum of an aggregated term based on flat weights, X_t^Q , and a weighted sum of (higher order) differences of the high frequency variable. Following Andreou, Ghysels, and Kourtellis (2010a), in the case of $q_X^D = 1$ we can easily show that the MIDAS term in equation (2.5) can be written as

$$\begin{aligned} \sum_{i=0}^{N_D-1} w_i(\theta^D) X_{N_D-i,t}^D &= \frac{1}{N_D} (X_{N_D,t}^D + X_{N_D-1,t}^D + \dots + X_{1,t}^D) \\ &+ (w_0 - \frac{1}{N_D}) X_{N_D,t}^D + (w_1 - \frac{1}{N_D}) X_{N_D-1,t}^D + \dots \\ &+ (w_{N_D-2} - \frac{1}{N_D-2}) X_{2,t}^D + (\frac{N_D-1}{N_D} - w_0 - w_1 - \dots - w_{N_D-2}) X_{1,t}^D, \end{aligned} \quad (2.6)$$

where the last parenthesis uses the assumption that the weights sum to one. Substituting equation (2.6) into (2.5) we get

$$Y_{t+1}^Q = \mu + \beta X_t^Q + \beta \sum_{i=0}^{N_D-1} (w_i(\theta^D) - \frac{1}{N_D}) \Delta^{N_D-i} X_{N_D-i,t}^D \quad (2.7)$$

Equation (2.7) shows that the traditional temporal aggregation approach, which imposes flat weights $w_i = 1/N_D$ and only accounts for X_t^Q , yields a nonlinear omitted variable term in the regression model (2.1). The nonlinearity of the omitted term is due to the nonlinear weighting schemes of MIDAS regression models such as the exponential Almon lag polynomial in (2.3).

In order to study the effects of misspecification imposed by the flat aggregation scheme, in the case of $q_X^D = 1$, let us denote the omitted term in (2.7) as $X_t^B(\theta) = \sum_{i=0}^{N_D-1} (w_i(\theta^D) - \frac{1}{N_D}) \Delta^{N_D-i} X_{N_D-i,t}^D$. Equation (2.7) implies that for a general, non-flat weighting scheme the traditional temporal aggregation approach may result in an omitted variable bias if the

omitted term, $X_t^B(\theta)$, is correlated with X_t^Q . This implies that the LS estimation of equation (2.1) will generally give rise to a bias, which depends on the type of the high frequency process and on the shape of the weighting scheme $W(L^{N_D}; \theta)$. For instance, declining weights imply an omitted variable that exhibits memory decay or mean reversion, which will be associated with higher bias than an omitted variable with a near-flat weighting scheme. Moreover, the bias will be zero in two cases: (i) When the omitted term $X_t^B(\theta)$ is orthogonal to X_t^Q , even when the true model is the DL-MIDAS regression. (ii) When the true weighting scheme is flat, $\theta = 0$, and the true model is the traditional DL model.

Andreou, Ghysels, and Kourtellis (2010a) show that when the omitted term is correlated with the equally weighted aggregated term in (2.5), then the Asymptotic Mean Squared Error (AMSE) of the LS estimator of β , is relatively larger than the AMSE of the NLS estimator of β in (2.5). In this case, one can easily show that the DL-MIDAS model in equation (2.5) based on NLS yields more accurate forecasts than forecasts based on LS estimation of the simple DL model in equation (2.1) assuming flat weights. In the section below we show the analytical expressions of the decomposition of the DL-MIDAS model in one important example namely when the high frequency predictor is an AR(1) process.

2.1 DL-MIDAS model with AR(1) high frequency predictor

Let abstract from the example of obtaining quarterly forecasts using daily observations and let us consider, in general, the high frequency univariate process $\{X_{t/m}^{(m)}\}$ sampled at some arbitrary high frequency m between t and $t - 1$ (which can be daily as in previous section or even intradaily), follow a stationary AR(1) given by $X_t^{(m)} = c_0 + \phi X_{t-1/m}^{(m)} + e_{t/m}$, $e_{t/m} \sim i.i.d.(0, \sigma_e^2)$.⁶ Then the MIDAS regression model with an AR(1) high frequency regressor is

$$Y_{t+1} = \beta_0 + \beta_1 X_t^A + \beta_1 X_t^B(\theta) + u_{t+1} \quad (2.8)$$

where the simple average term $X_t^A = \frac{1}{m} \sum_{j=1}^{m-1} X_{t-(j-1)/m}^{(m)}$ can be expressed as

$$X_t^A = \frac{1}{m} \left(\frac{c_0}{1-\phi} \left(m - \frac{1-\phi^m}{1-\phi} \right) + \left(\frac{1-\phi^m}{1-\phi} \right) X_{t-(m-1)} + \frac{1}{1-\phi} \sum_{j=1}^{m-1} (1-\phi^j) e_{t-(j-1)} \right) \quad (2.9)$$

⁶Following the previous section, this general notation also allows for the daily lags to extend beyond one quarter, $m \geq N_D$.

and the nonlinear term as

$$X_t^B(\theta) = \sum_{i=1}^{m-1} \left(w_i(\theta) - \frac{1}{m} \right) \left(c_0 \left(\frac{1 - \phi^{m-i}}{1 - \phi} \right) - (1 - \phi^{m-i}) x_{(t-(m-1))/m} + \sum_{j=0}^{m-i-1} \phi^j e_{t-j-(i-1)} \right). \quad (2.10)$$

Andreou, Ghysels, and Kourtellos (2010a) show that the LS estimator of β_1 in the linear DL-MIDAS model which omits $X_t^B(\theta)$ in (2.8), will be asymptotically biased. The analytical expression of the asymptotic bias for the slope coefficient given by

$$\begin{aligned} & \frac{\beta m \left(-\frac{1-\phi^m}{(1-\phi)} \frac{\sigma_e^2}{1-\phi^2} \sum_{i=1}^{m-1} \left(\left(w_i(\theta) - \frac{1}{m} \right) (1 - \phi^{m-i}) \right) \right)}{\left(\frac{1-\phi^m}{1-\phi} \right)^2 \frac{\sigma_e^2}{1-\phi^2} + \frac{\sigma_e^2}{(1-\phi)^2} \sum_{j=1}^{m-1} (1 - \phi^j)^2} \\ & + \frac{\frac{\beta}{(1-\phi)} \sum_{j=1}^{m-1} \left((1 - \phi^j) \left(\sum_{i=1}^j \left(w_i(\theta) - \frac{1}{m} \right) \phi^{j-i} \right) \right)}{\left(\frac{1-\phi^m}{1-\phi} \right)^2 \frac{\sigma_e^2}{1-\phi^2} + \frac{\sigma_e^2}{(1-\phi)^2} \sum_{j=1}^{m-1} (1 - \phi^j)^2} \end{aligned} \quad (2.11)$$

Andreou, Ghysels, and Kourtellos (2010a) evaluate numerically the analytical expression in (2.11) as a function of the aggregation horizon, m , for different values of the parameters, θ , ϕ , and σ_e^2 , in order to gain further insights about the behavior of the asymptotic bias. They find that the asymptotic bias for a persistent AR(1) process, $\phi = 0.9$ and $\sigma_e^2 = 1$, for the different weighting schemes: $\theta = (0, -0.05)$, $\theta = (0, -0.005)$, and $\theta = (0, -0.0005)$. In all cases they find that the bias becomes negative and increases in magnitude with m , where $m = 3$ and 100. As m becomes large the bias appears to stabilize at some negative value, which depends on the weighting scheme. This value is larger in absolute terms for faster decaying weights. As expected the bias of β_1 is larger for higher degrees of persistence, ϕ . In addition, simulation evidence reported in Table 1 of Andreou, Ghysels, and Kourtellos (2010a) shows that the Mean Square Error (MSE) gains from estimating a MIDAS regression model instead of a flat aggregation model are relevant even for $m = 3$ and realistic sample sizes, T , depending on the high frequency process and the pattern of the aggregation weights.

2.2 ADL-MIDAS regressions

When Y_{t+1}^Q is serially correlated, as it is typically the case for time series variables, the simple model in equation (2.1) is extended to a dynamic linear regression or autoregressive distributed lag (ADL) model. Again the conventional approach, in its simplest form, aggregates the high frequency data at the low frequency by computing simple averages and

estimates a simple linear regression of Y_{t+1}^Q on X_t^Q . Take for instance the ADL(1,1)

$$Y_{t+1}^Q = \mu + \alpha Y_t^Q + \beta X_t^Q + u_{t+1}, \quad (2.12)$$

where α and β are unknown parameters and u_{t+1} is an error term. In a similar manner the *ADL – MIDAS*(p_Y^Q, q_X^D) is:

$$Y_{t+1}^Q = \mu + \sum_{j=0}^{p_Y^Q-1} \alpha_{j+1} Y_{t-j}^Q + \beta \sum_{j=0}^{q_X^D-1} \sum_{i=0}^{N_D-1} w_{i+j*N_D}(\theta^D) X_{N_D-i,t-j}^D + u_{t+1} \quad (2.13)$$

Note that the number of daily lags is a multiple of the number of trading days in a quarter, N_D . As above the slope coefficient β in the MIDAS regression is identified via the scaling of the weights, such that they add up to one. The above model specification generates notation very similar to ARMA models, e.g. ADL-MIDAS(1,1) or ADL-MIDAS(AIC,AIC).

A MIDAS regression specification related to the ADL-MIDAS was proposed by Clements and Galvão (2008). Namely consider:

$$Y_{t+1}^Q = \mu + \alpha Y_t^Q + \beta \sum_{j=0}^{q_X^D-1} \sum_{i=0}^{N_D-1} w_{i+j*N_D}(\theta^D) X_{N_D-i,t-j}^D + u_{t+1}, \quad (2.14)$$

which can be written as a constrained DL-MIDAS regression with autocorrelated errors:

$$Y_{t+1}^Q = \mu(1 - \alpha)^{-1} + \beta \sum_{j=0}^{q_X^D-1} \sum_{i=0}^{N_D-1} w_{i+j*N_D}(\theta^D)(1 - L^Q)^{-1} X_{N_D-i,t-j}^D + \tilde{u}_{t+1}, \quad (2.15)$$

where L^Q is a quarterly lag operator and $\tilde{u}_{t+1} = (1 - \alpha L^Q)^{-1} u_t$. Clements and Galvão (2008) also consider a specification closely related to the ADL-MIDAS with a common factor restriction:

$$Y_{t+1}^Q = \mu + \alpha Y_t^Q + \beta \sum_{j=0}^{q_X^D-1} \sum_{i=0}^{N_D-1} w_{i+j*N_D}(\theta^D)(1 - \alpha L^Q)^{-1} X_{N_D-i,t-j}^D + u_{t+1}. \quad (2.16)$$

An undesirable property of both specifications is that the lag polynomial is characterized by geometrically declining spikes at distance N_D due to the interaction of the high frequency with the low frequency polynomials; see Ghysels, Sinko, and Valkanov (2006). Moreover,

while the ADL-MIDAS can be simply estimated by NLS, the DL-MIDAS with autocorrelated errors requires a more involved estimation such as nonlinear feasible GLS.

The comparison with temporal aggregation prompts us to consider two MIDAS regression models that allow for quarterly lags. First, define the following filtered parameter-driven *quarterly* variable

$$X_t^Q(\theta_X^D) \equiv \sum_{i=0}^{N^D-1} w_i(\theta_X^D) X_{N^D-i,t}^D, \quad (2.17)$$

Then, we can define the *ADL – MIDAS – M*(p_Y^Q, p_X^Q) model, where $-M$ refers to the fact that the model involves a multiplicative weighting scheme, namely:

$$Y_{t+1}^Q = \mu + \sum_{k=0}^{p_Y^Q-1} \alpha_k Y_{t-k}^Q + \sum_{k=0}^{p_X^Q-1} \beta_k X_{t-k}^Q(\theta_X^D) + u_{t+1}. \quad (2.18)$$

To alleviate the increasing number of parameters in equation (2.18) we can restrict the coefficients of the quarterly lags using another layer of a lag polynomial. This yields the *ADL – MIDAS – M*($p_Y^Q[r], p_X^Q[r]$) model:

$$Y_{t+1}^Q = \mu + \alpha \sum_{k=0}^{p_Y^Q-1} w_k(\theta_Y^Q) Y_{t-k}^Q + \beta \sum_{k=0}^{p_X^Q-1} w_k(\theta_X^Q) X_{t-k}^Q(\theta_X^D) + u_{t+1}, \quad (2.19)$$

where $[r]$ stands for the restricted case of estimating the model in equation *ADL – MIDAS – M*($p_Y^Q[r], p_X^Q[r]$). Both equations (2.18) and (2.19) apply MIDAS aggregation to the daily data of one quarter but they differ in the way they treat the quarterly lags. More precisely, while equation (2.18) does not restrict the coefficients of the quarterly lags, equation (2.19) restricts the coefficients of the quarterly lags - hence the notation $p_X^Q[r]$ - by hyper-parameterizing these coefficients using a multiplicative MIDAS polynomial.⁷ Both specifications nest the equally weighted aggregation scheme.

An interesting generalization of the ADL-MIDAS and ADL-MIDAS-M in equations (2.18) and (2.13), respectively, emerges when y_t is observed at a monthly frequency but one is

⁷The multiplicative MIDAS scheme was originally suggested for purpose of dealing with intra-daily seasonality in high frequency data, see Chen and Ghysels (2009).

interested at quarterly forecasts (e.g. CPI Inflation or Industrial Production). In this case we can easily generalize these models to allow for a MIDAS filter for the lagged dependent variable using another MIDAS polynomial.

2.3 Factors and other regressors in ADL-MIDAS models

A large body of recent work has developed factor model techniques that are tailored to exploit a large cross-sectional dimension; see for instance, Bai and Ng (2002), Bai (2003), Forni, Hallin, Lippi, and Reichlin (2000), Forni, Hallin, Lippi, and Reichlin (2005), Stock and Watson (1989), Stock and Watson (2003), among many others. These factors, which are usually estimated at quarterly frequency using a large cross-section of time-series are used as predictors in ADL models. Following this literature Andreou, Ghysels, and Kourtellis (2010b) investigate whether one can improve quarterly factor model forecasts by augmenting such models with daily financial variables and in particular daily financial factors. Such factors (at either low or high frequency) can be obtained by the following Dynamic Factor Model (DFM)

$$\begin{aligned} X_t &= \Lambda_t F_t + e_t \\ F_t &= \Phi_t F_{t-1} + \eta_t \\ e_{it} &= a_{it}(L)e_{it-1} + \varepsilon_{it}, \quad i = 1, 2, \dots, n, \end{aligned} \tag{2.20}$$

where $X_t = (X_{1t}, \dots, X_{nt})'$, F_t is the r -vector of static factors, Λ_t is a $n \times r$ matrix of factor loadings, $e_t = (e_{1t}, \dots, e_{nt})'$ is an n -vector of idiosyncratic disturbances, which can be serially correlated and (weakly) cross-sectionally correlated.⁸ The factor model representation in (2.20) allows for the possibility that the factor loadings change over time (compared to the standard DFMs) which may address potential instabilities during our sample period. The extracted common factors could be robust to instabilities in individual time series, if such instability is small and sufficiently dissimilar among individual variables; see Stock and Watson (2002) for formal conditions. Following the above assumptions the time-varying

⁸The static representation in equation (2.20) can be derived from the DFM assuming finite lag lengths and VAR factor dynamics in the DFM in which case F_t contains the lags (and possibly leads) of the dynamic factors. Although generally the number of factors from a DFM and those from a static one differ, we have that $r = d(s + 1)$ where r and d are the numbers of static and dynamic factors, respectively, and s is the order of the dynamic factor loadings. Moreover, empirically static and dynamic factors produce rather similar forecasts (Bai and Ng (2008)).

DFM can be estimated using principal components, which delivers consistent estimates of the common factors if $N \rightarrow \infty$ and $T \rightarrow \infty$.⁹ The number of factors can be chosen by the information criteria proposed; see for example Bai and Ng (2002).

These factors are then employed to augment the aforementioned MIDAS regression models. For instance, in the case of quarterly factors equation (2.13) generalizes to the *FADL – MIDAS*(p_Y^Q, p_F^Q, k_X^D) model given by

$$\begin{aligned}
Y_{t+1}^Q = & \mu + \sum_{k=0}^{p_Y^Q-1} \alpha_k Y_{t-k}^Q + \sum_{k=0}^{p_F^Q-1} \beta_k F_{t-k}^Q \\
& + \gamma \sum_{j=0}^{p_X^D-1} \sum_{i=0}^{N_D-1} w_{i+j*N^D} (\theta_X^D) X_{N_D-i, t-j}^D + u_{t+1}
\end{aligned} \tag{2.21}$$

Additionally, using the DFM in equation (2.20), Andreou, Ghysels, and Kourtellis (2010b) construct daily factors, denoted by F_t^D , which pool information from a large cross-section of daily financial data. This approach allows us to specify ADL-MIDAS models with both quarterly and daily factors that incorporate information across different frequencies while at the same time retain parsimony. For example, consider the FADL-MIDAS model in equation (2.22) using the daily factor as the daily predictor, $X_t^D = F_t^D$. Using a similar approach Clements and Galvão (2009) use leading indicators as predictors for quarterly macroeconomic variables, which is estimated using DL-MIDAS models with AR errors as in equation (2.15).

Note that we can also formulate a *FADL – MIDAS – M*(p_Y^Q, p_F^Q, p_X^Q) model, which involves the multiplicative MIDAS weighting scheme, hence generalizing equation (2.18). Notice also that equation (2.22) simplifies to the traditional factor model with additional regressors when the MIDAS features are turned off - i.e. say a flat aggregation scheme is used. When the lagged dependent variable is excluded then we have a projection on daily data, combined with aggregate factors. This brings us to the following benchmark models of *FADL*(p_Y^Q, p_F^Q, p_X^Q)

$$Y_{t+1}^Q = \mu + \sum_{k=0}^{p_Y^Q-1} \alpha_k Y_{t-k}^Q + \sum_{k=0}^{p_F^Q-1} \beta_k F_{t-k}^Q + \sum_{k=0}^{p_X^Q-1} \gamma_k X_{t-k}^Q + u_{t+1} \tag{2.22}$$

⁹Although the parametric AR assumption for F_t and e_{it} is not needed to estimate the factors, such assumptions can be useful when discussing forecasts using factors.

and $FAR(p_Y^Q, p_X^Q)$ when the regressor X^Q is not present

$$Y_{t+1}^Q = \mu + \sum_{k=0}^{p_Y^Q-1} \alpha_k Y_{t-k}^Q + \sum_{k=0}^{p_F^Q-1} \beta_k F_{t-k}^Q + u_{t+1} \quad (2.23)$$

3 Nowcasting and the Kalman Filter

We noted in the Introduction that the Kalman filter is a convenient computational device to extract missing data. We also noted that the approach has many benefits, but also some drawbacks. State space models can be quite computationally involved, as one must explicitly specify a linear dynamic model for all the series. State space models are therefore also prone to specification errors.

In this section we discuss how the regression in (2.13) relates to the more traditional approach involving the Kalman filter. It is natural to discuss this also in conjunction with the so called nowcasting literature discussed notably by Giannone, Reichlin, and Small (2008), among others. We start with a subsection on MIDAS with leads. The latter can be compared to nowcasting - although we consider the term MIDAS with leads more appropriate than nowcasting. The difference between nowcasting and MIDAS with leads can be explained with a simple example. Nowcasting refers to within-period updates of forecasts. An example would be weekly updates of current quarter GDP forecasts. MIDAS with leads can be viewed as - say again weekly updates - of not only current quarter GDP forecasts, but any future horizon GDP forecast (i.e. over several future quarters). Of course, when MIDAS with leads applies to updates of current quarter forecasts - it coincides with the exercise of nowcasting. The Kalman filter is typically used for nowcasting. We start with a discussion of MIDAS with leads and then cover connections with the Kalman filter.

3.1 MIDAS with leads

Giannone, Reichlin, and Small (2008), among others, have formalized the process of updating the nowcast and forecasts as new releases of data become available. This process can be mimicked via MIDAS regression models with *leads*. Say we are one or two months into quarter $t + 1$. Namely, we consider the MIDAS models with leads in order to incorporate real-time information available mainly on financial variables. Our objective is to forecast

quarterly economic activity and in practice we often have a monthly release of macroeconomic data within the quarter and the equivalent of at least 44 trading days of financial data observed with no measurement error. This means that if we stand on the first day of the last month of the quarter and wish to make a forecast for the current quarter we could use and around 44 leads of daily data for financial markets that trade on weekdays.

Consider the Factor ADL model with MIDAS in equation (2.22), which allows for J_X^D daily leads for the daily predictor, expressed in multiples of months, $J_X^D = 1, 2, \dots, J$. Then we can specify the *FADL – MIDAS*($p_Y^Q, p_F^Q, p_X^D, J_X^D$) model

$$\begin{aligned}
Y_{t+1}^Q = & \mu + \sum_{k=0}^{p_Y^Q-1} \alpha_k Y_{t-k}^Q + \sum_{k=0}^{p_F^Q-1} \beta_k F_{t-k}^Q + \gamma \left[\sum_{i=0}^{J_X^D-1} w_i(\theta_X^D) X_{J_X^D-i, t+1}^D \right. \\
& \left. + \sum_{j=0}^{p_X^D-1} \sum_{i=0}^{N_D-1} w_{i+j*N_D}(\theta_X^D) X_{N_D-i, t-j}^D \right] + u_{t+1},
\end{aligned} \tag{3.1}$$

Note that equation (3.1) differs from FADL-MIDAS model in (2.22) in that it includes the leads term $\sum_{i=0}^{J_X^D-1} w_i(\theta_X^D) X_{J_X^D-i, t+1}^D$, which uses daily information during period $t + 1$ to provide end of the quarter forecast of Y_{t+1} .

MIDAS with leads differs from the MIDAS regressions involving “leading indicator” series, as in Clements and Galvão (2009). The latter use of MIDAS regressions such as in equation (2.16) without leads appearing in the MIDAS polynomial, but with (monthly) leading indicator series *aligned* with quarterly GDP growth data.

3.2 Comparison with the Kalman filter

Bai, Ghysels, and Wright (2009) and Kuzin, Marcellino, and Schumacher (2009) discuss in detail the connections between the Kalman filter and MIDAS regressions. It is the purpose of this subsection to summarize their findings. The first important observation is that a MIDAS regression can be viewed as a reduced form representation of the linear projection that emerges from a state space model approach - where by reduced form we mean that the MIDAS regression does not require the specification of a full state space system of equations.

For illustrative purposes, consider a simple dynamic single factor model:

$$F_{i,t} = \rho F_{(i-1),t} + \eta_{i,t} \quad \forall t = 1, \dots, T, \quad i = 2, \dots, N_D \quad (3.2)$$

and $F_{1,t} = \rho F_{N_D,t-1} + \eta_{1,t}$. Moreover, let $\eta_{i,t}$ be i.i.d. Gaussian with mean zero and variance σ_η^2 . Suppose now the daily data $X_{i,t}^D$ relates to the factors as follows:

$$X_{i,t}^D = \gamma F_{i,t} + u_{i,t} \quad i \neq N_D \quad (3.3)$$

with $u_{i,t}$ i.i.d. Gaussian with mean zero and variance σ_u^2 . Finally, at the end of each quarter, we have:

$$X_{N_D,t}^D = \gamma F_{N_D,t} + u_{N_D,t} \quad Y_t^Q = F_{N_D,t} + v_{N_D,t} \quad (3.4)$$

with $v_{i,t}$ i.i.d. Gaussian with mean zero and variance σ_v^2 . This highly stylized state space model with mixed sampling and minimal parametric specification (involving five parameters collected in $\theta^S \equiv (\rho, \gamma, \sigma_\eta^2, \sigma_u^2, \sigma_v^2)$). Bai, Ghysels, and Wright (2009) show that the steady state Kalman filter corresponds to the following *ADL – MIDAS*(∞, ∞) :

$$E_t[Y_{t+1}^Q] = \sum_{j=0}^{\infty} \alpha_{j+1}(\theta^S) Y_{t-j}^Q + \beta \sum_{j=0}^{\infty} \sum_{i=1}^{N_D} w_{i+j*n_D}(\theta^S) X_{i,t-j}^D \quad (3.5)$$

where E_t is linear projection using past quarterly and daily data combined. The weights have a structure very similar to the MIDAS regression appearing in (2.13) and a related one discussed below in equation (2.18). It is important to note that the Kalman filter requires to specify a complete system of equations, which we kept to an absolute minimum representation in the above motivating example. Nevertheless, we counted five parameters driving the weights in equation (3.5) compared to two for the Exponential Almon weighting scheme of the MIDAS regression. In some cases the MIDAS regression is an exact representation of the Kalman filter, in other cases it involves approximation errors that are typically small.¹⁰ The Kalman filter, while clearly optimal as far as linear projections goes, has two main disadvantages (1) it is more prone to specification errors as a full system of equations for Y , X , and latent factors is required and (2) as already noted it requires a lot more parameters to achieve the same goal. Bai, Ghysels, and Wright (2009) show that the weighting scheme in equations (2.18) and (2.19) corresponds to the structure of a steady state Kalman filter

¹⁰Bai, Ghysels, and Wright (2009) discusses both the cases where the mapping is exact and the approximation errors in cases where the MIDAS does not coincide with the Kalman filter.

linear projection with mixed sampling frequencies. Namely,

$$E_t[Y_{t+1}^Q] = \sum_{j=0}^{\infty} \alpha_{j+1}(\theta^S) Y_{t-j}^Q + \beta \sum_{j=0}^{\infty} \sum_{i=1}^{N_D} w_k(\theta^S) X_{t-k}^Q(\theta^S) \quad (3.6)$$

with $X_{t-k}^Q(\theta^S)$ similar to $X_t^Q(\theta_X^D)$ appearing in equation (2.17). The downside of the MIDAS specification in equations (2.18) and (2.19) is that it is less parsimonious than the single weighting scheme in equation (2.13). Yet, it typically involves less parameters than the multiplicative scheme emerging from the Kalman filter appearing in driven by θ^S . Note also that equation (2.19) is more parsimonious than equation (2.18), and at the same time also more restrictive.

4 Forecasting volatility

There is a large literature on forecasting volatility and in particular using high frequency (intraday) data (see for instance Andersen, Bollerslev, Christoffersen, and Diebold (2006)). Here, we focus on the issues pertaining to mixed frequencies - typically created by multi-step volatility forecasting. In this respect, the MIDAS approach complements the literature on forecasting volatility in several important directions. Note that a related topic to multi-step volatility forecasting is that of forecasting Value-at-Risk (VaR) within the risk management literature. In the context of forecasting the 10-day VaR, required following the Basle accord, using daily or even intradaily information, MIDAS models can be used to produce directly multi-step forecasts (see for instance, Chen and Ghysels (2009)).

In order to analyze the role of MIDAS in forecasting volatility let us introduce the relevant notation. Let $V_{t+1,t}$ be a measure of volatility in the next period. We focus on predicting future conditional variance, measured as increments in quadratic variation (or its log transformation), due to the large body of existing recent literature on this subject. The increments in the quadratic variation of the return process, $Q_{t+1,t}$, is not observed directly but can be measured with some discretization error. One such measure would be the sum of (future) m intra-daily squared returns denoted r^{ID} , namely $\sum_{j=1}^m [r_j^{ID}]^2$, which we will denote by $\tilde{Q}_{t+1,t}^{(m)}$ since it involves a discretization based on m intra-daily returns. The superscript in parentheses indicates the number of high-frequency data used to compute the variable.

We change slightly the notation in this section for the regressors. We used the notation $X_{j,t}^D$

to refer to daily data in quarter t . In this section we simply refer to X_t^D . A MIDAS volatility regression with daily predictors is:

$$\tilde{Q}_{t+1,t}^{(m)} = \mu + \phi \sum_{k=0}^{k^{max}} w(k, \theta) X_{t-k}^D + \varepsilon_t \quad (4.1)$$

The volatility specification (4.1) has a number of important features.

First, the volatility measure on the left-hand side, and the predictors on the right-hand side are sampled at different frequencies. As a result the volatility in equation (4.1), can be measured at different horizons (e.g. daily, weekly, and monthly frequencies), whereas the forecasting variables X_{t-k}^D are available at daily or higher frequencies. Thus, this specification allows us not only to forecast volatility with data sampled at different frequencies, but also to compare such forecasts and ultimately evaluate empirically the continuous asymptotic arguments. In addition, equation (4.1) provides a method to investigate whether the use of high-frequency data necessarily leads to better volatility forecasts at various horizons.

Second, the weight function or the polynomial lag parameters w not only share all the advantages (discussed in previous section), but they can be especially relevant in estimating a persistent process parsimoniously, such as volatility, where distant X_{t-k}^D are likely to have an impact on current volatility. In addition, the parameterization also allows us to compare MIDAS regressions at different frequencies as the number of parameters to estimate will be the same even though the weights on the data and the forecasting capabilities might differ across horizons. Most importantly one does not have to adjust measures of fit for the number of parameters and in most situations with one predictor one has a MIDAS model with either one or two parameters determining the pattern of the weights.

Third, MIDAS regressions typically do *not* exploit an autoregressive scheme, so that X_{t-k}^D is not necessarily related to lags of the left hand side variable. Instead, MIDAS regressions are first and foremost regressions and therefore the selection of X_{t-k}^D amounts to choosing the best predictor of future quadratic variation from the set of several possible measures of past fluctuations in returns. In other words, MIDAS can be considered as a reduced-form forecasting method of volatility, rather than a model of conditional variance. Various regressors can be used in the MIDAS equation (4.1) to examine whether future volatility is well predicted that synthesize alternative methods in the literature. Examples of X_{t-k}^D are past daily squared returns (that correspond to the ARCH-type of models with some

parameter restrictions, Engle (1982) and Bollerslev (1986)), absolute daily returns (that relate to the specifications of (see e.g. Ding, Granger, and Engle (1993)), realized daily volatility (e.g. Andersen, Bollerslev, and Diebold (2010)), realized daily power of (Barndorff-Nielsen and Shephard (2003) and Barndorff-Nielsen, Graversen, and Shephard (2004)), and daily range (e.g. Alizadeh, Brandt, and Diebold (2002) and Gallant, Hsu, and Tauchen (1999)). Since all of the regressors are used within a framework with the same number of parameters and the same maximum number of lags, the results from MIDAS regressions are directly comparable. Moreover, MIDAS regressions can also be extended to study the joint forecasting power of the regressors.

Related to the MIDAS volatility regression is the Heterogeneous Autoregressive Realized Volatility (HAR-RV) regressions proposed in Andersen, Bollerslev, and Diebold (2007) and Corsi (2009). The HAR-RV model is given by (dropping m as argument for future and past realized volatilities):

$$RV_{t+1,t} = \mu + \beta^D RV_t^D + \beta^W RV_t^W + \beta^M RV_t^M + \varepsilon_{t+1}, \quad (4.2)$$

which has a simple linear prediction regression using RV over heterogeneous interval sizes, daily (D), weekly (W) and monthly (M). As noted by Andersen, Bollerslev, and Diebold (2007) (footnote 16) the above equation is in a sense a MIDAS regression with “step functions” (in the terminology of Ghysels, Sinko, and Valkanov (2006)).¹¹ In this regards the HAR-RV can be related to the MIDAS-RV in (4.1) of Ghysels, Santa-Clara, and Valkanov (2006) and Forsberg and Ghysels (2006), using different weight functions such as the Beta, exponential Almon or step functions and different regressors not just autoregressive with mixed frequencies. Note also that both models exclude the jump component of quadratic variation. Simulation results reported in Forsberg and Ghysels (2006) also show that the difference between HAR and MIDAS models is very small for RV. For other regressors, such as the realized absolute variance, the MIDAS model performs slightly better.

The MIDAS approach can also be used to study various other interesting aspects of forecasting volatility. Chen and Ghysels (2009) provide a comprehensive study and a novel method to analyze the impact of news on forecasting volatility. The following semi-parametric regression model is proposed to predict future realized volatility (RV) with past

¹¹See also the discussion in Corsi (2009) - page 181 - on this topic.

high-frequency returns:

$$RV_{t+1,t} = \psi_0 + \sum_{j=1}^{\tau} \sum_{i=1}^m \psi_{i,j}(\theta) NIC(r_{j,t}^{ID}) + \varepsilon_{t+1} \quad (4.3)$$

where $\psi_{i,j}(\theta)$ is a polynomial lag structure parameterized by θ , $NIC(\cdot)$ is the news impact curve and $r_{t/m}$ is the log asset price difference (return) over some short time interval i of length m on day t . Note $i = 1, \dots, m$ of intervals on day t .

The regression model in (4.3) shows that each intra-daily return has an impact on future volatility measured by $NIC(r_{j,t}^{ID})$ and fading away through time with weights characterized by $\psi_{i,j}(\theta)$. One can consider (4.3) as the semi-parametric (SP) model that nests a number of volatility forecasting models and in particular the benchmark realized volatility forecasting equation below:

$$RV_{t+1,t} = \psi_0 + \sum_{j=0}^{\tau} \psi_j(\theta) RV_{t-j}^D + \varepsilon_{t+1} \quad (4.4)$$

The nesting of (4.4) can be seen for $k = 1, \dots$, when we set $\psi_{i,j} \equiv \psi_i \forall j = 1, \dots, \Delta^{-1}$, and $NIC(r) \equiv r^2$ in equation (4.3). This nesting emphasizes the role played by both the news impact curve NIC and the lag polynomial $\psi_{i,j}$.

The reason it is possible to nest the RV AR structure is due to the multiplicative specification for $\psi_{i,j}(\theta) \equiv \psi_j^D(\theta) \times \psi_i^{ID}(\theta)$, with the parameter θ containing subvectors that determine the two polynomials separately. The polynomial $\psi_j^D(\theta)$ is a daily weighting scheme, similar to $\psi_i(\theta)$ in the regression model appearing in (4.4). The polynomial $\psi_i^{ID}(\theta)$ relates to the intra-daily pattern. With equal intra-daily weights one has the RV measure when NIC is quadratic - as is the case in the Symmetric model. Chen and Ghysels (2009) adopt the following specification for the polynomials:

$$\psi_j^D(\theta) \psi_i^{ID}(\theta) = Beta(j, \tau, \theta_1, \theta_2) \times Beta(i, 1/m, \theta_3, \theta_4) \quad (4.5)$$

where τ and $1/m$ are the daily (D) and intradaily (ID) frequencies. The restriction is imposed that the intra-daily patterns wash out across the entire day, i.e. $\sum_i Beta(i, 1/m, \theta_3, \theta_4) = 1$, and also impose without loss of generality, a similar restriction on the daily polynomial, in order to identify a slope coefficient in the regressions.

The multiplicative specification (4.5) has several advantages. First, as noted before, it nests

the so called flat aggregation scheme, i.e. all intra-daily weights are equal, yields a daily model with RV when the news impact curve is quadratic. Or more formally, when $\theta_3 = \theta_4 = 1$, and $NIC(r) = r^2$ one recovers RV -based regression appearing in equation (4.4). Second, by estimating $Beta(i, 1/m, \theta_3, \theta_4)$ one lets the data decide on the proper aggregation scheme which is a generic issue pertaining in MIDAS regressions as discussed in Andreou, Ghysels, and Kourtellis (2010a). Obviously, the intra-daily part of the polynomial will pick up how news fades away throughout the day and this - in part - depends on the well known intra-daily seasonal pattern.

Finally, the MIDAS-NIC model can also nest existing parametric specifications of news impact curves adopted in the ARCH literature, namely, the daily symmetric one when $NIC(r) = br^2$, the asymmetric GJR model when $NIC(r) = br^2 + (cr^2)\mathbf{1}_{r < 0}$ (Glosten, Jagannathan, and Runkle (1993)) and the asymmetric GARCH model when $NIC(r) = (b(r - c)^2)$ (Engle (1990)).

5 Conclusion

In this chapter we reviewed the use of regression models that involve data sampled at different frequencies. The research area is still in its infancy as there many topics we did not cover such multivariate models, Granger causality with mixed frequency data (see however Ghysels, Sinko, and Valkanov (2009)), nonlinear models, to name a few. Finally, the chapter dealt almost exclusively with the use of high frequency data to improve forecasts of low frequency data. In some circumstances the reverse may be of interest. An example is the use of macroeconomic variables in daily or monthly volatility forecasting - as in for instance Engle, Ghysels, and Sohn (2008) - or the use of low frequency correlations to improve daily frequency correlation forecasts, as in Colacito, Engle, and Ghysels (2010).

Last but not least we would like to note the availability of a Matlab Toolbox for MIDAS regressions, see Sinko, Sockin, and Ghysels (2010).

References

- ALIZADEH, S., M. W. BRANDT, AND F. X. DIEBOLD (2002): “Range-based estimation of stochastic volatility models,” *Journal of Finance*, 57, 1047–1091.
- ALPER, C. E., S. FENDOGLU, AND B. SALTOGLU (2008): “Forecasting Stock Market Volatilities Using MIDAS Regressions: An Application to the Emerging Markets,” MPRA Paper No. 7460.
- ANDERSEN, T., T. BOLLERSLEV, P. CHRISTOFFERSEN, AND F. DIEBOLD (2006): “Volatility and Correlation Forecasting,” in *Handbook of Economic Forecasting*, ed. by E. Graham, C. Granger, and A. Timmermann, pp. 778–878. North Holland, Amsterdam.
- ANDERSEN, T., T. BOLLERSLEV, AND F. DIEBOLD (2010): “Parametric and Nonparametric Volatility Measurement,” in *Handbook of Financial Econometrics*, ed. by Y. Aït-Sahalia, and L. P. Hansen.
- ANDERSEN, T. G., T. BOLLERSLEV, AND F. DIEBOLD (2007): “Roughing It Up: Including Jump Components in the Measurement, Modeling and Forecasting of Return Volatility,” *The Review of Economics and Statistics*, 89, 701–720.
- ANDREOU, E., E. GHYSELS, AND A. KOURTELLOS (2010a): “Regression Models With Mixed Sampling Frequencies,” *Journal of Econometrics*, in-press.
- (2010b): “Should macroeconomic forecasters use daily financial data and how?,” Discussion paper University of North Carolina and University of Cyprus.
- ANGELINI, E., G. CAMBA-MENDEZ, D. GIANNONE, G. RÜNSTLER, AND L. REICHLIN (2008): “A Two-Step Estimator for Large Approximate Dynamic Factor Models Based on Kalman Filtering,” ECB Working Papers.
- ARMESTO, M., K. ENGEMANN, AND M. OWYANG (2010): “Forecasting with Mixed Frequencies,” The Federal Reserve Bank of St. Louis Review, (forthcoming).
- ARMESTO, M. T., R. HERNANDEZ-MURILLO, M. OWYANG, AND J. PIGER (2009): “Measuring the Information Content of the Beige Book: A Mixed Data Sampling Approach,” *Journal of Money, Credit and Banking*, 41, 35–55.

- ARUOBA, S. B., F. X. DIEBOLD, AND C. SCOTTI (2009): “Real-time measurement of business conditions,” *Journal of Business and Economic Statistics*, 27, 417–427.
- BAI, J. (2003): “Inferential theory for factor models of large dimensions,” *Econometrica*, pp. 135–171.
- BAI, J., E. GHYSELS, AND J. WRIGHT (2009): “State space models and MIDAS regressions,” Working paper, NY Fed, UNC and Johns Hopkins.
- BAI, J., AND S. NG (2002): “Determining the number of factors in approximate factor models,” *Econometrica*, 70, 191–221.
- (2008): “Large Dimensional Factor Analysis,” in *Foundations and Trends in Econometrics*, vol. 3, pp. 89–163.
- BARNDORFF-NIELSEN, O., S. GRAVERSEN, AND N. SHEPHARD (2004): “Power variation and stochastic volatility: a review and some new results,” *Journal of Applied Probability*, 41A, 133–143.
- BARNDORFF-NIELSEN, O., AND N. SHEPHARD (2003): “Power variation with stochastic volatility and jumps,” Discussion paper, Nuffield College.
- BERNANKE, B., M. GERTLER, AND M. WATSON (1997): “Systematic Monetary Policy and the Effects of Oil Price Shocks,” *Brookings Papers on Economic Activity*, 1, 91–157.
- BHANSALI, R. (1999): “Parameter Estimation and Model Selection for Multistep Prediction of a Time Series: A Review,” in *Asymptotics, Nonparametrics, and Time Series*, ed. by S. Ghosh. Marcel Dekker.
- BOLLERSLEV, T. (1986): “Generalized Autoregressive Conditional Heteroskedasticity,” *Journal of Econometrics*, 31, 307–327.
- CHEN, X., AND E. GHYSELS (2009): “News - good or bad - and its impact on predicting future volatility,” *Review of Financial Studies* (forthcoming).
- CHEVILLON, G., AND D. HENDRY (2005): “Non-parametric direct multi-step estimation for forecasting economic processes,” *International Journal of Forecasting*, 21, 201–218.
- CLEMENTS, M., AND D. HENDRY (1996): “Mutli-Step Estimation for Forecasting,” *Oxford Bulletin of Economics and Statistics*, 58, 657–684.

- CLEMENTS, M. P., AND A. B. GALVÃO (2008): “Macroeconomic Forecasting with Mixed Frequency Data: Forecasting US output growth,” *Journal of Business and Economic Statistics*, 26, 546–554.
- (2009): “Forecasting US output growth using Leading Indicators: An appraisal using MIDAS models,” *Journal of Applied Econometrics*, 24, 1187–1206.
- CLEMENTS, M. P., A. B. GALVÃO, AND J. H. KIM (2008): “Quantile forecasts of daily exchange rate returns from forecasts of realized volatility,” *Journal of Empirical Finance*, 15, 729–750.
- COLACITO, R., R. ENGLE, AND E. GHYSELS (2010): “A component model of dynamic correlations,” *Journal of Econometrics*, (forthcoming).
- CORSI, F. (2009): “A simple approximate long-memory model of realized volatility,” *Journal of Financial Econometrics*, 7, 174–196.
- DING, Z., C. GRANGER, AND R. ENGLE (1993): “A long memory property of stock market returns and a new models,” *Journal of Empirical Finance*, 1, 83–106.
- DOZ, C., D. GIANNONE, AND L. REICHLIN (2006): “A Two-Step Estimator for Large Approximate Dynamic Factor Models Based on Kalman Filtering,” CEPR Discussion Papers 6043.
- (2008): “A Quasi Maximum Likelihood Approach for Large Approximate Dynamic Factor Models,” CEPR Discussion Papers 5724.
- ENGLE, R. (1982): “Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation,” *Econometrica*, 50, 987–1008.
- ENGLE, R. F. (1990): “Discussion: Stock market volatility and the crash of ’87,” *Review of Financial Studies*, 3, 103–106.
- ENGLE, R. F., E. GHYSELS, AND B. SOHN (2008): “On the Economic Sources of Stock Market Volatility,” Discussion Paper NYU and UNC.
- FINDLEY, D. (1983): “On the Use of Multiple Models of Multi-period Forecasting,” *Proceedings of the Business and Statistics Section, American Statistical Association*, pp. 528–531.

- (1985): “Model Selection for Multi-Step-Ahead Forecasting,” in *Proceedings of the 7th Symposium on Identification and System Parameter Estimation*, ed. by S. Ghosh, and P. Young, pp. 1039–1044. Pergamon, Oxford.
- FORNI, M., M. HALLIN, M. LIPPI, AND L. REICHLIN (2000): “The generalized dynamic-factor model: Identification and estimation,” *Review of Economics and Statistics*, 82, 540–554.
- (2005): “The generalized dynamic factor model,” *Journal of the American Statistical Association*, 100, 830–840.
- FORSBERG, L., AND E. GHYSELS (2006): “Why do absolute returns predict volatility so well?,” *Journal of Financial Econometrics*, 6, 31–67.
- GALLANT, A. R., C.-T. HSU, AND G. TAUCHEN (1999): “Using daily range data to calibrate volatility diffusions and extract the forward integrated variance,” *Review of Economic Statistics*, 81, 617–631.
- GALVÃO, A. B. (2006): “Changes in Predictive Ability with Mixed Frequency Data,” Discussion Paper QUeen Mary.
- GHYSELS, E., P. SANTA-CLARA, AND R. VALKANOV (2004): “The MIDAS touch: Mixed Data Sampling Regressions,” Discussion paper UNC and UCLA.
- (2005): “There is a risk-return tradeoff after all,” *Journal of Financial Economics*, 76, 509–548.
- GHYSELS, E., P. SANTA-CLARA, AND R. VALKANOV (2006): “Predicting volatility: getting the most out of return data sampled at different frequencies,” *Journal of Econometrics*, 131, 59–95.
- GHYSELS, E., AND A. SINKO (2010): “Volatility Forecasting and Microstructure Noise,” *Journal of Econometrics*, (forthcoming).
- GHYSELS, E., A. SINKO, AND R. VALKANOV (2006): “MIDAS regressions: Further results and new directions,” *Econometric Reviews*, 26, 53–90.
- GHYSELS, E., A. SINKO, AND R. VALKANOV (2009): “Granger Causality Tests with Mixed Data Frequencies,” Discussion Paper, UNC.

- GHYSELS, E., AND J. WRIGHT (2009): “Forecasting professional forecasters,” *Journal of Business and Economic Statistics*, 27, 504–516.
- GIANNONE, D., L. REICHLIN, AND D. SMALL (2008): “Nowcasting: The real-time informational content of macroeconomic data,” *Journal of Monetary Economics*, 55, 665–676.
- GLOSTEN, L. R., R. JAGANNATHAN, AND D. E. RUNKLE (1993): “On the relation between the expected value and the volatility of the nominal excess return on stocks,” *Journal of Finance*, 48, 1779–1801.
- GRANGER, C. (1985): “Aggregation of time-series variables: A survey,” in *Disaggregation in econometric modelling*, ed. by T. Barker, and M. H. Pesaran, pp. 17–34. Routledge, London and New York.
- HAMILTON, J. D. (2006): “Daily Monetary Policy Shocks and the Delayed Response of New Home Sales,” working paper, UCSD.
- HARVEY, A. (1989): *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press, Cambridge.
- HARVEY, A. C., AND R. G. PIERSE (1984): “Estimating missing observations in economic time series,” *Journal of the American Statistical Association*, 79, 125–131.
- HENDRY, D. F., AND K. HUBRICH (2010): “Combining disaggregate forecasts or combining disaggregate information to forecast an aggregate,” *Journal of Business and Economic Statistics*, forthcoming.
- HOTTA, L. K., AND J. C. NETO (2008): “The effect of aggregation on prediction in autoregressive integrated moving-average models,” *Journal of Time Series Analysis*, 14, 261–269.
- KUZIN, V., M. MARCELLINO, AND C. SCHUMACHER (2009): “MIDAS versus mixed-frequency VAR: nowcasting GDP in the euro area,” Discussion Paper 07/2009 Deutsche Bundesbank.
- LEÓN, Á., J. M. NAVE, AND G. RUBIO (2007): “The relationship between risk and expected return in Europe,” *Journal of Banking and Finance*, 31, 495–512.

- LIN, J., AND C. GRANGER (1994): “Model Selection for Multiperiod Forecasts,” *Journal of Forecasting*, 13, 1–9.
- LÜTKEPOHL, H. (2004): “Forecasting with VARMA processes,” in *Handbook of Economic Forecasting*, ed. by E. Graham, C. Granger, and A. Timmermann, pp. 87–115. North Holland, Amsterdam.
- MARCELLINO, M., J. STOCK, AND M. WATSON (2006): “A comparison of direct and iterated multistep AR methods for forecasting macroeconomic time series,” *Journal of Econometrics*, 135, 499–526.
- MARIANO, R. S., AND Y. MURASAWA (2003): “A new coincident index of business cycles based on monthly and quarterly series,” *Journal of Applied Econometrics*, 18, 427–443.
- MITNIK, S., AND P. A. ZADROZNY (2004): *Forecasting quarterly German GDP at monthly intervals using monthly Ifo business conditions data*. Springer.
- MOENCH, E., S. NG, AND S. POTTER (2009): “Dynamic hierarchical factor models,” Staff Reports 412, Federal Reserve Bank of New York.
- MONTEFORTE, L., AND G. MORETTI (2009): “Real Time Forecasts of Inflation: The Role of Financial Variables,” Bank of Italy.
- 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, 386–398.
- SINKO, A., M. SOCKIN, AND E. GHYSELS (2010): “Matlab Toolbox for MIDAS Regressions,” Available at www.unc.edu/~eghysels/Software_datasets.html.
- STOCK, J. H. (2006): “Forecasting and Now-casting with Disparate Predictors: Dynamic Factor Models and Beyond,” Manuscript, Department of Economics, Harvard University.
- STOCK, J. H., AND M. W. WATSON (1989): “New indexes of coincident and leading economic indicators,” *NBER Macroeconomics Annual*, pp. 351–394.
- (2002): “Forecasting using principal components from a large number of predictors,” *Journal of the American Statistical Association*, 97, 1167–1179.

- (2003): “Forecasting output and inflation: the role of asset prices,” *Journal of Economic Literature*, pp. 788–829.
- TAY, A. (2006): “Financial Variables as Predictors of Real Output Growth,” Discussion Paper SMU.
- TAY, A. S. (2007): “Mixing Frequencies: Stock Returns as a Predictor of Real Output Growth,” Discussion Paper SMU.
- ZADROZNY, P. A. (1990): “Forecasting US GNP at monthly intervals with an estimated bivariate time series model,” *Federal Reserve Bank of Atlanta Economic Review*, 75, 2–15.