AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

Macroeconomic Forecasting with Mixed Frequency Data:
Forecasting US output growth and inflation.

Michael P. Clements and Ana Beatriz Galvão

No 773

WARWICK ECONOMIC RESEARCH PAPERS

DEPARTMENT OF ECONOMICS

THE UNIVERSITY OF
WARWICK

Macroeconomic Forecasting with Mixed Frequency Data: Forecasting US output growth and inflation.

Michael P. Clements*

Department of Economics

University of Warwick

M.P.Clements@warwick.ac.uk

Ana Beatriz Galvão

Bank of Portugal

abferreira@bportugal.pt

First version September 2005. This version July 2006.

Abstract

Although many macroeconomic series such as US real output growth are sampled quarterly, many potentially useful predictors are observed at a higher frequency. We look at whether a recently developed mixed data-frequency sampling (MIDAS) approach can improve forecasts of output growth and inflation. We carry out a number of related real-time forecast comparisons using various indicators as explanatory variables. We find that MIDAS model forecasts of output growth are more accurate at horizons less than one quarter using coincident indicators; that MIDAS models are an effective way of combining information from multiple indicators; and that the forecast accuracy of the unemployment-rate Phillips curve for inflation is enhanced using the MIDAS approach.

Keywords: Data frequency, multiple predictors, combination, real-time forecasting.

JEL classification: C51, C53.

*We are grateful to the Editor, an Associate Editor, and two anonymous referees for helpful comments and guidance. We also acknowledge helpful comments from participants at the International Symposium of Forecasting in Santander, June 2006. Michael P. Clements is the corresponding author: Department of Economics, University of Warwick, CV4 7AL Coventry, UK; Phone: 44-2476-522990.

Macroeconomic Forecasting with Mixed Frequency Data: Forecasting US output growth and inflation

Abstract

Although many macroeconomic series such as US real output growth are sampled quarterly, many potentially useful predictors are observed at a higher frequency. We look at whether a recently developed mixed data-frequency sampling (MIDAS) approach can improve forecasts of output growth and inflation. We carry out a number of related real-time forecast comparisons using various indicators as explanatory variables. We find that MIDAS model forecasts of output growth are more accurate at horizons less than one quarter using coincident indicators; that MIDAS models are an effective way of combining information from multiple indicators; and that the forecast accuracy of the unemployment-rate Phillips curve for inflation is enhanced using the MIDAS approach.

Keywords: Data frequency, multiple predictors, combination, real-time forecasting.

JEL classification: C51, C53.

First version September 2005. This version July 2006.

1 Introduction

The unavailability of key macroeconomic variables such as GDP (or GNP) at frequencies higher than quarterly has led to many macroeconometric models being specified on quarterly data, as well as attempts (following Chow and Lin (1971)) to estimate monthly GDP series from monthly data related to GDP and the quarterly GDP series. We assume that the aim is to forecast a quarterly macroeconomic series, but allowing for the potentially useful information available at higher frequencies, especially monthly. For forecasting output growth, for example, there are a range of leading and coincident indicator variables at a monthly frequency. Because of the requirement that all the variables in the model are sampled at the same frequency, a number of strategies are commonly used to attempt to profit from the monthly data. Perhaps the most common is filtering the monthly data to the quarterly frequency, either by averaging the months, or taking the last (or middle) month in the quarter. There are formal ways of exploiting the monthly information by combining forecasts from monthly models with the forecasts from quarterly models (see, e.g., Miller and Chin (1996)) ranging to less formal approaches that seek to use the higher-frequency information by making off-line adjustments to the forecasts.

In recent research Ghysels, Santa-Clara and Valkanov (2004a), Ghysels, Sinko and Valkanov (2004c) have suggested the use of time-series regressions that allow the regressand and regressors to be sampled at different frequencies: the so-called MIDAS (MIXed Data Sampling) approach. Typically, the regressand is sampled at the lower frequency. With a few exceptions (such as Ghysels and Wright (2005)), the approach has so far been applied to modelling high-frequency financial data, with Ghysels, Santa-Clara and Valkanov (2004b) predicting the weekly and monthly future volatility of equity returns using daily realized volatility and intraday squared and absolute returns.

In this paper we explore whether MIDAS specifications are useful for modelling and forecasting two key US macroeconomic variable, US output growth and the inflation rate. Much of the recent forecasting literature seeks to exploit information from large sets of variables, either by combining the predictions from simple models (each using a small subset of the possible explanatory variables) as exemplified by Stock and Watson (1999), or by using factor models (see e.g., Forni, Hallin, Lippi and Reichlin (2000)). We show that combining information from multiple potential predictors is one way in which MIDAS models may contribute to improved macroeconomic forecasts. In addition, MIDAS models may yield more accurate short-term forecasts, especially when the lag between the quarterly macro series and the monthly series is short. Significant improvements in forecast accuracy are found from the direct use of information on the monthly indicator series inside the quarter being forecast.

Where possible, our findings are based on comparisons of MIDAS model forecasts with forecasts

from other models using a real-time approach to modelling and forecast evaluation, as discussed below. Competitor forecasts are generated from standard (quarterly-frequency) distributed lag and autoregressive distributed lag models using various indicators as explanatory variables. We consider a number of selection methods when there are multiple indicators, including ‘bagging’, as discussed by Inoue and Kilian (2005), in order to obtain rivals to MIDAS that offer a stern test.

The effects of data vintage on model specification and forecast evaluation have been addressed in a number of recent papers: Patterson (1995, 2003), Robertson and Tallman (1998), Orphanides (2001), Croushore and Stark (2001, 2003) and Koenig, Dolmas and Piger (2003), among others. We generate forecasts using the vintage of data available at the time the forecast was made. Faust, Rogers and Wright (2005) show that data revisions (at least for the US) are largely unpredictable, so the practice of using latest vintage data would appear to be untenable (it would not have been possible to predict and therefore utilise these latest vintage values at the time). Faust, Rogers and Wright (2003) find that the forecasting performance of exchange rate models, for example, can be quite sensitive to whether the forecasts are based on real-time or revised data. To provide further evidence of the differences that arise between using real-time and latest-vintage data, in some cases we repeat the exercise using latest vintage data. A more difficult issue concerns which vintage of data to use for the actuals: whether the aim was to forecast an early release, or later revised releases of the data. For some of the forecast comparison exercises accuracy measures are calculated using both the first-release and the latest available data, so that the sensitivity of our findings to this assumption can be seen. The real-time perspective also leads us to consider models that either use individual component indicators of the composite leading indicator (CLI, produced by the Conference Board), or re-weighted combinations of these components, rather than using the CLI itself. The historical CLI series has been subject to revisions and re-weightings of the components which make it inappropriate for use in a real-time forecast exercise (see e.g., Diebold and Rudebusch (1991)).

The plan of the rest of our paper is as follows. In section 2 we briefly review the MIDAS approach of Ghysels *et al.* (2004a, 2004c). In section 3 we discuss a number of extensions and refinements that facilitate the application of MIDAS modelling to macroeconomic data, including the treatment of the autoregressive structure of macro variables and the estimation of an autoregressive MIDAS model. In section 4 we outline the ways in which we compare the MIDAS and quarterly-frequency models. We also test whether differences in predictive accuracy measured by Root Mean Square Forecast Errors (RMSFE) are statistically significant, and whether the quarterly model forecasts encompass those of the MIDAS models, using the approach of Clark and McCracken (2005). The empirical results are discussed in section 5. The discussion is divided into three parts: forecasting output growth with leading indicators (section 5.1), forecasting output growth with coincident

indicators (section 5.2) and forecasting inflation (section 5.3). Section 6 offers some concluding remarks.

2 MIDAS regression approach

The MIDAS models of Ghysels *et al.* (2004a, 2004c) are closely related to distributed lag models (see, e.g., Dhrymes (1971) and Sims (1974)). The response to the higher frequency explanatory variables is modelled using highly parsimonious distributed lag polynomials, as a way of preventing the proliferation of parameters that might otherwise result, and as a way of side-stepping difficult issues to do with lag-order selection. Parameter proliferation could be especially important in financial applications, where say, daily volatility is related to 5-minute interval intraday data (so that a day's worth of observations amounts to 288 data points), but parsimony is also likely to be important in typical macroeconomic applications, where quarterly data are related to monthly data, given the much smaller numbers of observations typically available. Modelling the coefficients on the lagged explanatory variables as a distributed lag function allows for long lags with only a small number of parameters needing to be estimated.

The basic MIDAS model for a single explanatory variable, and one-step ahead forecasting, is given by:

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

where $B\left(L^{1/m}; \boldsymbol{\theta}\right) = \sum_{k=1}^K b(k; \boldsymbol{\theta}) L^{(k-1)/m}$, and $L^{s/m} x_{t-1}^{(m)} = x_{t-1-s/m}^{(m)}$. Here, t indexes the basic time unit (in our case, quarters), and m is the higher sampling frequency ($m = 3$ when x is monthly and y is quarterly), and as shown $L^{1/m}$ operates at the higher frequency. The ‘Exponential Almon Lag’ of Ghysels *et al.* (2004a, 2004c) parameterizes $b(k; \boldsymbol{\theta})$ as:

$$b(k; \boldsymbol{\theta}) = \frac{\exp(\theta_1 k + \theta_2 k^2)}{\sum_{k=1}^K \exp(\theta_1 k + \theta_2 k^2)}$$

To make matters concrete, consider the basic MIDAS regression with $m = 3$ and $K = 12$:

$$y_t = \beta_0 + \beta_1 B\left(L^{1/3}; \boldsymbol{\theta}\right) x_{t-1}^{(3)} + \varepsilon_t^{(3)} \quad (2)$$

where $B\left(L^{1/3}; \boldsymbol{\theta}\right) = \sum_{k=1}^{12} b(k; \boldsymbol{\theta}) L^{(k-1)/3}$, and $L^{s/3} x_{t-1}^{(3)} = x_{t-1-s/3}^{(3)}$, so that:

$$y_t = \beta_0 + \beta_1 [b(1; \boldsymbol{\theta}) x_{t-1}^{(3)} + b(2; \boldsymbol{\theta}) x_{t-1-1/3}^{(3)} + \dots + b(12; \boldsymbol{\theta}) x_{t-4-2/3}^{(3)}] + \varepsilon_t^{(3)} \quad (3)$$

If y_t is quarterly output growth in 2005Q1, say, then $x_{t-1}^{(3)}$ is the value of x in December 2004, $x_{t-1-1/3}^{(3)}$ the value of x in November 2004, and so on, back to $x_{t-4-2/3}^{(3)}$ denoting January 2004.

This is a one step (i.e., one quarter) ahead forecasting model. For h -steps ahead, we would use:

$$\begin{aligned} y_t &= \beta_0 + \beta_1 B\left(L^{1/3}; \boldsymbol{\theta}\right) x_{t-h}^{(3)} + \varepsilon_t^{(3)} \\ y_t &= \beta_0 + \beta_1 \left[b(1; \boldsymbol{\theta}) x_{t-h}^{(3)} + b(2; \boldsymbol{\theta}) x_{t-h-1/3}^{(3)} + b(3; \boldsymbol{\theta}) x_{t-h-2/3}^{(3)} + b(4; \boldsymbol{\theta}) x_{t-h-1}^{(3)} + \dots \right] + \varepsilon_t^{(3)} \end{aligned} \quad (4)$$

where strictly $\boldsymbol{\beta} = (\beta_0 \ \beta_1)'$, $\boldsymbol{\theta}$ and $\varepsilon_t^{(3)}$ should be indexed by the forecast horizon. This approach to multi-step forecasting is in the spirit of ‘multi-step’ or ‘direct estimation’ (see the review by Bhansali (2002)), whereby y_t is related directly to information available at $t - h$, as opposed to generating forecasts from a one-step model such as (2). The obvious advantage for models with other variables as explanatory factors is that models to forecast these variables are not required.

Notice that the standard practice of calculating a quarterly series from the monthly indicator x amounts to imposing restrictions on $B(L^{1/3}; \boldsymbol{\theta})$. For example, in relation to (3), taking the last month in the quarter to produce a quarterly series amounts to setting $b(2; \boldsymbol{\theta}) = b(3; \boldsymbol{\theta}) = b(5; \boldsymbol{\theta}) = b(6; \boldsymbol{\theta}) = \dots = b(11; \boldsymbol{\theta}) = b(12; \boldsymbol{\theta}) = 0$. The empirical implications of these restrictions are discussed in section 5.

3 Application of MIDAS models to macroeconomic data

3.1 Forecasting with information inside the quarter

Macroeconomic forecasts are often produced a number of times during each quarter. For example, the staff of the Board of Governors of the Federal Reserve prepare forecasts for the meetings of the Open Market Committee. These meetings occur several times each quarter: see Karamouzis and Lombra (1989) and Joutz and Stekler (2000). Therefore, monthly data pertaining to the first part of the quarter will sometimes be available for forecasts made later in the quarter. The MIDAS framework can exploit this data. For example, suppose we consider forecasts made every quarter, but after the value of x in the first month of the quarter is known. In which case the MIDAS model (4) becomes:

$$y_t = \beta_0 + \beta_1 B\left(L^{1/3}; \boldsymbol{\theta}\right) x_{t-2/3}^{(3)} + \varepsilon_t^{(3)}$$

with $h = 2/3$ because $1/3$ of the information of the current quarter is employed. Forecasts with $h = 1/3$ are also possible using information on the first two months of the quarter being forecast. So the number of steps-ahead allowed in the MIDAS approach are $h = 1/m, 2/m, 3/m \dots$. For a given observed value of y , this approach produces m different forecasts based on different information assumptions about the regressors and also estimates of $\{\boldsymbol{\beta}, \boldsymbol{\theta}\}$. One might expect the relative

performance of the MIDAS model to be good because it uses all higher frequency information available at the time the forecast is made: see e.g., Montgomery, Zarnowitz, Tsay and Tiao (1998) on the use of monthly data for quarterly predictions in a different context.

3.2 Combining Leading Indicators

A potential application of the MIDAS approach in macroeconomic forecasting is in the combination of diverse monthly leading indicators to predict quarterly output growth. Birchenhall, Jessen, Osborn and Simpson (1999) propose a ‘general-to-specific’ selection algorithm as a manageable way of choosing a parsimonious model from a large set of explanatory variables. Clements and Galvão (2006) employ this approach to select the component leading indicators (and their associated lags) to be included at each step ahead. General-to-specific selection typically result in ‘holes’ in lag structures (e.g., say for $h = 1$ we have the first and third lags of x_i , for $h = 2$ the second and fourth) and may raise the suspicion of overfitting. However, Campos, Hendry and Krolzig (2003) and Hendry and Krolzig (2005) argue that these fears are often exaggerated, especially if ‘multi-path’ searches are used. The main advantage of MIDAS in this context is that a relatively parsimonious model can be estimated that incorporates all the component leading indicators (of the composite leading index of the Conference Board, say).

A MIDAS model that combines the information of nl monthly leading indicators to predict output growth, h -steps-ahead, would be written as:

$$y_t = \beta_0 + \sum_{i=1}^{nl} \beta_{1i} B_i(L^{1/m}; \boldsymbol{\theta}_i) x_{i,t-h}^{(m)} + \varepsilon_t^{(m)} \quad (5)$$

where the component indicators are indexed by i , and $m = 3$. Each leading indicator requires the estimation of only two parameters to describe the lag structure ($\boldsymbol{\theta}_i$) and one to weight their impact on y_t (β_{1i}). Because the number of parameters required for each additional leading indicator is small, it is feasible to estimate a model with all the leading indicators. The β_{1i} parameters define the weights attached to the leading indicators, and are specific to the forecast horizon.

An alternative approach to generating forecasts when there are multiple potential predictors is by ‘bagging’: see Inoue and Kilian (2005). Bagging is short for bootstrap aggregation (see Breiman (1996)). The idea is to average (‘aggregate’) forecasts from models estimated on bootstrap resamples. For forecasting US inflation, Inoue and Kilian (2005) find that bagging methods perform well compared to alternatives such as factor models, shrinkage estimators, forecast combination, and Bayesian model averaging. We briefly describe bagging, as we shall use it to generate competing forecasts to the MIDAS forecasts. Consider the h -step ahead forecasting model given by:

$$y_t = \mathbf{\Gamma}' \mathbf{x}_{t-h} + \varepsilon_t$$

where $\mathbf{x}_{t-h}$ is a vector containing lags of y and lags of the explanatory variables, subject to all being dated $t - h$ or earlier. The statistical rule for choosing the elements of $\mathbf{x}_{t-h}$ (denoted $x_{i,t-h}$) is that the corresponding variable is chosen if $|t_i| > c$, where t_i is the usual t -statistic of $\Gamma_i = 0$ computed using the OLS estimator and the Newey-West variance estimator (assuming autocorrelation of order $h - 1$ in the disturbances). The equation is then re-estimated with the selected regressors on data up to T , and the forecast generated as $\hat{y}_{T+h} = \hat{\gamma}' \mathbf{S}_T \mathbf{x}_T$. Here, $\mathbf{S}_T$ is the selection matrix that picks out the regressors chosen in the last step, and $\hat{\gamma}$ is the estimate of γ , where γ is defined implicitly by the regression of y_t on $\mathbf{S}_T \mathbf{x}_{t-h}$. Bagging consists in aggregating the forecasts from models selected and estimated on a number (B) of bootstrap resamples, where in each case the same selection rule is applied. The forecast obtained by bagging is then:

$$\hat{y}_{T+h}^B = \frac{1}{B} \sum_{i=1}^B \hat{\gamma}_i^{*'} \mathbf{S}_T^{*i} \mathbf{x}_T,$$

where $\hat{\gamma}_i^*$ and $\mathbf{S}_T^{*i}$ are respectively the coefficient estimates and selection matrices in each bootstrap sample. Inoue and Kilian (2005) provide details of the re-sampling scheme. Gains in forecasting accuracy may result from bagging if typical pre-test strategies (such as deleting all variables with small t -test values) are unstable, in the sense that the set of selected variables changes in response to small changes in the data set (e.g., the addition of a new observation as the estimation window moves through the data). Bagging should counter this instability by averaging the forecasts from the models selected on each bootstrap sample.

3.3 Autoregressive structure

Models to forecast US GNP growth are often simple autoregressive-distributed lag (ADL) models, incorporating leading indicators as explanatory variables: see e.g., Stock and Watson (2003). Including autoregressive dynamics in models that sample the explanatory variables at a higher frequency may also improve forecast accuracy. As noted by Ghysels *et al.* (2004c), this is not straightforward unless the dependent variable is also available at the higher frequency. When it is not, simply adding a lower frequency lag of y , y_{t-1} to (1), results in:

$$y_t = \beta_0 + \lambda y_{t-1} + \beta_1 B \left(L^{1/3}; \boldsymbol{\theta} \right) x_{t-1}^{(3)} + \varepsilon_t^{(3)}.$$

As Ghysels *et al.* (2004c) explain, this strategy is not in general appropriate, as from writing the model as:

$$y_t = \beta_0 (1 - \lambda)^{-1} + \beta_1 (1 - \lambda L)^{-1} B \left(L^{1/3}; \boldsymbol{\theta} \right) x_{t-1}^{(3)} + \tilde{\varepsilon}_t^{(3)}$$

($\tilde{\varepsilon}_t^{(3)} = (1 - \lambda L)^{-1} \varepsilon_t^{(3)}$) it is apparent that the polynomial on $x_{t-1}^{(3)}$ is the product of a polynomial in $L^{1/3}$, $B \left(L^{1/3}; \boldsymbol{\theta} \right)$, and a polynomial in L , $\sum \lambda^j L^j$. This mixture generates a ‘seasonal’ response

of y to $x^{(3)}$, irrespective of whether $x^{(3)}$ displays a seasonal pattern.

Our suggested solution is simply to introduce autoregressive dynamics in y_t as a common factor (see, e.g., Hendry and Mizon (1978)):

$$y_t = \beta_0 + \lambda y_{t-1} + \beta_1 B\left(L^{1/3}; \boldsymbol{\theta}\right) (1 - \lambda L) x_{t-1}^{(3)} + \varepsilon_t^{(3)}.$$

so that the response of y to $x^{(3)}$ remains non-seasonal. A multi-step analogue can be written as:

$$y_t = \lambda y_{t-h} + \beta_1 + \beta_2 B(L^{1/3}, \boldsymbol{\theta}) \left(1 - \lambda L^h\right) x_{t-h} + \epsilon_t, \quad (6)$$

which we term the MIDAS-AR.

3.4 Estimation of MIDAS models with autoregressive lags

In this section we consider estimation issues. We begin with a brief discussion of the estimation of the standard MIDAS model. Ghysels, Santa-Clara and Valkanov (2004b) show that non-linear least squares is a consistent estimator in this case. Suppose $m = 3$, so that:

$$y_t = \beta_0 + \beta_1 B\left(L^{1/3}; \boldsymbol{\theta}\right) x_{t-h}^{(3)} + \varepsilon_t^{(3)} \quad (7)$$

Using an exponential function to define the weights on past values of x , and assuming that $K = 24$ (two years of past monthly information):

$$\begin{aligned} B\left(L^{1/3}; \boldsymbol{\theta}\right) &= \sum_{k=1}^{24} b(k; \boldsymbol{\theta}) L^{(k-1)/3} \\ b(k; \boldsymbol{\theta}) &= \frac{\exp(\theta_1 k + \theta_2 k^2)}{\sum_{k=1}^{24} \exp(\theta_1 k + \theta_2 k^2)}. \end{aligned}$$

The dimension of the numerical optimization procedure to obtain the parameters $\boldsymbol{\beta}$ and $\boldsymbol{\theta}$ can be reduced by concentrating the least squares objective function with respect to $\boldsymbol{\beta}$. For a given $\boldsymbol{\theta}$, $\boldsymbol{\beta}$ can be obtained by the least squares formula:

$$\boldsymbol{\beta} = \left(\sum_{t=h}^T \mathbf{x}_{t-h}(\boldsymbol{\theta}) \mathbf{x}_{t-h}(\boldsymbol{\theta})' \right)^{-1} \left(\sum_{t=h}^T \mathbf{x}_{t-h}(\boldsymbol{\theta}) y_t \right)$$

where $\mathbf{x}_{t-h}(\boldsymbol{\theta}) = \left[1, B\left(L^{1/3}; \boldsymbol{\theta}\right) x_{t-h}^{(3)} \right]$ and $\boldsymbol{\beta} = \left(\beta_0, \beta_1 \right)'$.

The computation is done using the package constrained ML - Gauss 5 package CML 2.0 - selecting the BFGS algorithm. The restrictions imposed in the estimation are that $\theta_1 \leq 300$ and that $\theta_2 < 0$. We experiment with a number of initial values for $\boldsymbol{\theta}$ in order to counter any

dependence of the optimization routine on the initial values. When the MIDAS with multiple regressors of equation (5) is estimated, initial values of θ_i for each weighting function are taken from MIDAS estimation with single regressors.

To estimate the autoregressive MIDAS model (6), we take the residuals ($\hat{\varepsilon}_t$) of the standard MIDAS in equation (7), and estimate an initial value for λ , say $\hat{\lambda}_0$, from $\hat{\lambda}_0 = (\sum \hat{\varepsilon}_{t-h}^2)^{-1} \sum \hat{\varepsilon}_t \hat{\varepsilon}_{t-h}$. We 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 $\hat{\theta}_1$ is obtained by applying nonlinear least squares to:

$$y_t^* = \beta_1 + \beta_2 B(L^{1/3}, \theta) x_{t-h}^* + \varepsilon_t.$$

A new value of λ , $\hat{\lambda}_1$, is obtained from the residuals of this regression. Then using $\hat{\lambda}_1$ and $\hat{\theta}_1$ as initial values, we run BFGS to get the estimates $\hat{\lambda}$ and $\hat{\theta}$ that minimize the sum of squared residuals.

4 Evaluation of the forecast accuracy of MIDAS models

The forecasts from MIDAS are compared to those from standard quarterly-frequency models by comparing their Root Mean Square Forecast Errors (RMSFEs), as described in detail in the following section. We also test whether the differences in predictive accuracy are statistically significant. The tests take into account the uncertainty introduced by parameter estimation in addition to the underlying uncertainty in forecasting. West (1996) is a seminal paper on the impact of estimation uncertainty on tests of predictive accuracy, while West (2006) provides a general review of the literature.

The comparisons we make involve nested models. Compared to an autoregressive model, MIDAS models may improve forecasting by including explanatory variables and by allowing temporal disaggregation of those variables. The MIDAS-AR therefore nests the AR in two dimensions. The investigation of whether a model that excludes a particular variable is as good as the unrestricted model (that includes that variable) is the usual form of nested model comparison in the literature (i.e., that an ADL nests an AR). In addition, if we compare MIDAS-AR and ADL models with the same indicator variable then the temporal disaggregation of the MIDAS model gives rise to nesting by the MIDAS-AR of the ADL. The testing approach that has been developed to enable tests of predictive ability for standard nested structures can also be applied to the case of nesting due to temporal disaggregation.

As our evaluation is of multi-step forecasts computed by direct estimation, as well as involving nested models, we follow the approach of Clark and McCracken (2005). They consider approaches to testing for equal forecast accuracy and forecast encompassing (see, e.g., Chong and Hendry (1986),

Clements and Hendry (1993)) for multi-step forecasts from nested models. Clark and McCracken (2005) take the test of equal MSE of McCracken (2004) (labelled MSE-F) and the encompassing test of Clark and McCracken (2001) (ENC-F) and show that their limiting distributions depend on unknown nuisance parameters. They find that a bootstrap implementation of these tests (similar to that of Kilian (1999)) appears to work well, and that ENC-F has superior power.

Appendix I briefly records the bootstrap method we use to calculate p -values of the null hypotheses that the DL forecasts are as accurate as the MIDAS (MSE-F test) and that the DL forecasts encompass those from the MIDAS (ENC-F). As discussed in the literature relating to nested model comparisons, the hypothesis of interest is one-sided: the alternative to the null is that the nested model is less accurate than the larger model (MSE-F). In this framework, testing that the nesting model encompasses the smaller model makes no sense. A brief Monte Carlo exercise indicates that the tests have reasonable size and power properties.

5 Empirical section: Forecasting output growth and inflation

We present a number of related empirical forecast comparisons to gauge the usefulness of various aspects of MIDAS models for forecasting two key macroeconomic variables: output growth and inflation. The forecast comparisons are based on three separate exercises: forecasting output growth using leading indicators, forecast output growth using coincident indicators, and forecasting inflation.

5.1 Forecasting output growth using leading indicators

We assess the relative usefulness of MIDAS (compared to quarterly-frequency models) for extracting information from single monthly leading indicators (LIs) for forecasting output growth. That is, are there gains in forecast accuracy to using MIDAS rather than aggregating the monthly indicator series to obtain a quarterly series? We also wish to measure the size of any gains from using MIDAS models with monthly information on the quarter being forecast. After considering the indicators one at a time, we then assess the usefulness of the MIDAS approach in extracting information from the set of leading indicators taken together. As competitors, we use forecasts from quarterly-frequency ADL models, with regressors and lags chosen using the strategy of Birchenhall *et al.* (1999) based on the Schwarz information criterion, as well as the bagging approach to model selection (following Inoue and Kilian (2005)).

The leading indicator data are the component LIs of the Conference Board’s Composite Leading Indicator (CLI), as used by Stock and Watson (2003). The data is monthly from 1959:1 to 2002:9, and the ten component series are listed in Table 1, along with the transformation of the series applied

prior to the use of the series in modelling. The transformed series are also normalized (demeaned and divided by the standard deviation). The quarterly LI data are obtained by averaging the raw monthly series before applying the transformations indicated in Table 1, and for the quarterly series differencing refers to quarterly differencing.

5.1.1 Single-indicator models

A quarterly-frequency distributed lag (DL) model that employs x to forecast y h -steps ahead is:

$$y_t = \beta_0 + \sum_{i=0}^{p-1} \beta_i x_{t-h-i} + \varepsilon_t \quad (8)$$

Models are estimated for each h . For the DL model, p is determined by the Schwarz information criterion based on a maximum lag of 8, whereas for the MIDAS models we also allow up to two years' data on the monthly LIs (by setting $K = 24$). We perform a real-time recursive forecasting exercise so that the model is specified and re-estimated on ever increasing sample sizes, using the vintage of output data available at the time (see Croushore and Stark (2001) and the review by Croushore (2006)). Unfortunately real-time data are not readily available for all the LIs, and final data are used instead. Note that real-time and final-vintage data coincide for a number of the indicator variables such as interest rates.

To see how real-time output data is used in the forecasting exercise, consider the quarterly-frequency DL models. The first data set is the 1989:Q1 vintage, that contains data from 1959:Q1 to 1988:Q4 (the start date is common across vintages). A 1-step forecast of 1989:Q1 is generated from regressing output up to 1988:Q4 on the quarterly observations on the LI up to 1988:Q3. The 1988:Q4 and earlier values of the LI are then used with these coefficient estimates to forecast 1989:Q1. This forecast is evaluated using the observed value of 1989:Q1 from the latest vintage (2002:Q4), as well as the first-released data from the next quarterly vintage (in terms of our example, this is the 1989:Q2 vintage). A 2-step forecast, again made with the 1989:Q1 vintage, regresses output growth up to 1988:Q4 on the LI up to 1988:Q2, and uses these coefficient estimates along with LI data up to 1988:Q4 to forecast 1989:Q2. This is evaluated using the 1989:Q3 vintage value for 1989:Q2 (first-release) as well as the latest-vintage value. Three and four step forecasts are constructed in a similar fashion. We proceed in this way for each of the data vintages, so that for the last vintage of 2002:Q3, we have 1 to 4-step forecasts of 2001:Q3 to 2002:Q2. This gives 51 forecasts of each of the 4 horizons.

For the MIDAS models we proceed as above for the DL models but in addition we exploit the monthly LI data, and also consider intra-quarter horizons of one and two months, e.g., $h = 1/3, 2/3, \dots$. For $h = 1/3$, for example, the forecast of 1989:Q1 is generated from a model estimated

by regressing output data up to 1988:Q4 (from the 1989:Q1 real-time quarterly data set) on monthly LI data up to 1988:11. The monthly leading indicator series up to 1989:02 are then used with these coefficient estimates to forecast 1989:Q1. The forecast is evaluated as for the quarterly-frequency DL model forecasts, using first-release and latest-vintage data values.

The results are presented in Table 2. The first part of Table 2 presents MIDAS model RMSFEs, and ratios of the MIDAS model RMSFEs to those of the DL model forecasts, calculated assuming the actual values are given by the first-release data. The second part presents the same quantities assuming the actuals are the latest vintage values. A comparison of these two parts of the table shows that RMSFEs computed using latest-vintage data are larger than using first-release data, and that the ranking of the two models can switch between the two (e.g., for $h = 1/3$, using claims and ordersc, where the MIDAS forecasts are better evaluated using latest-vintage data).

It is evident that the relative performance of the two models differs between indicator, over horizon and depends on the definition of actuals. However, the findings suggest that: for some indicators MIDAS forecasts with a monthly horizon ($h = 1/3$) are more accurate than the $h = 1$ DL model forecasts (e.g., for vendor, expectations); that the MIDAS forecasts are more accurate at $h = 1$ for some indicators (e.g., the spread); and for some indicators this greater accuracy extends to longer horizons, such as MIDAS $h = 4/3$ against $h = 2$. Moreover, even when the null of the DL forecasts being at least as accurate as MIDAS is not rejected, the null of forecast encompassing is sometimes rejected, indicating that the MIDAS forecasts contain useful information for forecasting over and above that included in the DL model forecasts. (The MSE-F and ENC-F p -values are given in the columns headed simply MSE and ENC in all the tables). Nevertheless, MIDAS is clearly not always superior to the quarterly frequency DL model. In section 5.2 we perform a similar exercise with coincident indicators (instead of LIs), as it may be that the ability of MIDAS models to exploit intra-quarter monthly information will be more valuable for series which are more ‘coincident’ with the target series.

Table 3 shows the effects of allowing autoregressive output growth dynamics in the MIDAS models (as described in section 3.3) and in the competitor quarterly-frequency models. The table reports RMSFEs and ratios of RMSFEs for first-release actual values. Allowing autoregressive dynamics in the MIDAS generally reduces RMSFEs (see the ratios of MIDAS-AR to MIDAS), and for four of the ten indicators the autoregressive MIDAS models are significantly more accurate at $h = 1$ than the autoregressive-distributed lag quarterly-frequency models (ADL). The last two sets of comparisons in the table (the ADL against the AR, and the MIDAS-AR against the AR) question the predictive ability of leading-indicator information for output growth whether in quarterly or monthly form. The AR model is beaten for $h = 1$ only for building as the leading indicator when used within the MIDAS framework. However, the results of the forecast encompassing tests

caution against discarding the leading indicator information, as the null that the AR model forecast encompasses the models with LIs is rejected for seven of the ten indicators, suggesting combining the indicators' forecasts. We comment further on forecast combination in the following section.

5.1.2 Combining Leading indicators

In this section we evaluate whether MIDAS is an effective way of combining the information in monthly LIs to forecast quarterly output growth. The competitor to MIDAS is a quarterly-frequency distributed lag model that contains multiple LIs (M-DL). The M-DL is specified using the general-to-specific SIC strategy of Birchenhall *et al.* (1999), with the LIs and the lags at which they enter being chosen for each horizon and forecast origin, starting with $(K/3)$ lags for each LI.

The multiple MIDAS (M-MIDAS) specification uses all ten LIs, and is re-estimated on each data vintage and for each forecast horizon for that vintage. The LIs are again standardized (as in the previous section) before the estimation. The numerical minimization of the non-linear least squares function employs as initial values the estimates of θ_i obtained from the single indicator MIDAS models. The estimated β_i 's give the weights on the non-linear combination of past values of each of the LIs.

The forecast exercise is conducted as in the previous section. The RMSFEs for horizons 1 to 4 for the M-MIDAS and M-DL models are presented in Table 4. The results indicate large gains (20 - 30%) from using MIDAS to combine LIs compared to M-DL. M-MIDAS is only a little worse than the best single-indicator MIDAS model (building) at a one step horizon, and is superior to some of the single-indicator models. For $h = 4$, M-MIDAS is still better than M-DL, while some single-indicator models have significantly better performance. Because *ex ante* it would be unclear which of the LI-models to use, some form of combining to guard against a bad selection might be thought desirable. In that case, the M-MIDAS performance is quite promising.

There is a vast literature on the combination or pooling of forecasts from different models or sources: see, e.g., the reviews by Diebold and Lopez (1996), Newbold and Harvey (2002) and Timmermann (2006). As an alternative to combining information in modelling, we consider combining forecasts from a range of simple models (see Clements and Galvão (2006) who consider the relative usefulness of combining information in modelling versus combining forecasts for quarterly-frequency models of output growth). Table 4 includes the results for an equal-weighted combination of the single LI model forecasts (denoted 'Mean'), as simple averaging has often been found to be competitive with schemes that estimate optimal weights. We also record RMSFEs where the information set is restricted to the three *ex-post* best leading indicators of the single LI models. In which case, the M-MIDAS and M-DL models are formulated using only these three LIs, and the 'Mean' in

the second panel combines the forecasts from these three single LI models. We find that forecast combination is preferable to combining information in modelling, whether we consider all the LIs or the subset. Moreover, the pooled forecast RMSFEs are similar for the MIDAS and DL models.

Finally, we consider forecasts generated by bagging (see section 3.2). We use bagging to generate forecasts from a quarterly-frequency multiple-indicator ADL model where we allow the information set to consist of all the LIs or just the best three. We wish to check whether the gains to combination by the MIDAS approach hold up against this alternative method of generating forecasts using quarterly-frequency indicators. Although Inoue and Kilian (2005) show that bagging works well in selecting out of 26 regressors for forecasting inflation, we have ten LIs, and found that allowing eight lags of each (to match the two years of monthly data ($K = 24$) of the M-MIDAS model reported in Table 4) resulted in near-singularity of the second-moment matrix of the explanatory variables in the general unrestricted model. Instead we report comparisons for the number of months of data in the M-MIDAS $K = 6, 12$ and 15 , against a comparable quarterly maximum lag order for the bagged forecasts (2, 4 and 5 quarters, respectively). The need to estimate the unrestricted model in the bagging approach would appear to restrict the generality of that approach. We report results for bagging with a critical value of the t -test of the exclusion decision of $c = 2.575$. This gave the most favourable results, compared to testing at the 5 and 10% levels, and corresponds to testing at the 1% level. As we assume that there is always one autoregressive term of output growth, bagging is employed to select lags of the LIs.

The results in Table 5 indicate that the bagging method is worse than M-MIDAS at the shorter horizons, but better at the longer horizons, when we consider the information set consisting of all ten indicators. The better accuracy of M-MIDAS holds across different values of K , and is in excess of 20% at $h = 1$. The relative performance of bagging at the shorter horizons improves when the subset of three indicators is considered. Bagging is less accurate than equal-weighted combinations of single-indicator model forecasts for either the set of three or ten LIs. Because bagging is viewed as a competitive forecasting approach when there are multiple potential predictors, we view these results as underscoring the value of the MIDAS approach as a useful way of combining information in modelling.

5.2 Forecasting output growth using real-time coincident indicators

In this section we test our conjecture that MIDAS models may be better suited to exploit the more timely information in coincident indicators (CIs) than in LIs. We consider three coincident indicators which are available in the real-time monthly data sets of the Philadelphia Fed (see Table 1). Using the real-time monthly data series for these series and the real-time quarterly output

growth series we are able to carry out a proper real-time forecast comparison of the MIDAS and quarterly-frequency models using only the information that was available at each of the forecast origins.

For this forecasting exercise, we compute forecasts recursively using quarterly data vintages for output growth of 1985:Q1 to 2005:Q1 inclusive. The sample period starts in 1959Q1 in each instance. We generate 1 to 4-step ahead forecasts of 1985:Q1 to 1985:Q4 using the 1985:Q1 vintage, up to 1 to 4-step ahead forecasts of 2004:Q1 to 2004:Q4 using the 2004:Q1 vintage, giving a sample of 77 forecasts at each horizon. For the quarterly-frequency models the indicators are obtained by aggregating the monthly data from the monthly vintages data sets (1985:M1 until 2005:M3).

Table 6 describes the construction of the MIDAS forecasts with horizons from $1/3$ to 2, making explicit the way in which the monthly and quarterly real-time data sets are used. The table shows that the forecasts of 2000:Q1 use monthly information up to 1999M12, 2000M1 and 2000M2 to generate forecasts of length $h = 1$, $2/3$ and $1/3$ respectively, by making use of the monthly vintages 2000M1, 2000M2 and 2000M3. In each case, the quarterly output data run to 1999Q4, as given in the 2000:Q1 real-time data set. The table makes clear an important feature of the real-time generation of forecasts: the model is estimated at each horizon on the same vintage of monthly indicator data as is used in generating the forecasts. Hence the model estimates may change between the $2/3$ and $1/3$ horizon forecasts of 2000:Q1 because of the additional monthly data observation (the value of the indicator in 2000:M2) as well as because of vintage effects, namely revisions to the entire history of the indicator series.

In addition to the real-time exercise ‘RT’ (as described in Table 6), we compute forecasts for all the models using latest vintage data. That is, the ‘No RT’ columns in Table 7 refer to forecasts generated using the 2005:Q1 vintage of data for output growth and the 2003:M3 vintage for the CIs, as well as using this latest vintage data to construct forecast errors. A comparison of the MIDAS RMSFEs for ‘RT’ and ‘No RT’ in Table 7 indicates markedly more accurate forecasts using real-time data.

In the real-time exercise, the MIDAS model forecasts with a horizon of one month ($h = 1/3$) are significantly more accurate than the DL model $h = 1$ forecasts for both industrial production and employment, and for capacity utilisation the DL forecasts do not encompass the MIDAS. These results indicate marked improvements in forecast accuracy from MIDAS models that utilise within-quarter monthly data compared to quarterly-frequency models that do not use this information. For industrial production and employment the MIDAS model forecasts are significantly more accurate at the quarterly horizon ($h = 1$), as well as at longer horizons. There is also some indication that the relative performance of the MIDAS models is enhanced using real-time, as opposed to latest-vintage data. Although not shown to save space, the MIDAS model with industrial production is

also appreciably more accurate than an autoregressive model: nearly 20% for 1/3 : 1 and 10% for the 4/3 : 2 horizon comparisons.

5.3 Forecasting inflation

Our final empirical exercise is an investigation of the usefulness of MIDAS models for forecasting inflation, and the extent to which the findings depend on the use of real-time versus latest-vintage data. Inflation is measured by the change in the GDP deflator, and the forecast period and way in which the exercise is carried out parallels that of section 5.2. The monthly indicators we consider are the unemployment rate, capacity utilization, and M2. These indicators are evaluated using quarterly vintages (1985:Q1-2005:Q1) and are described in Table 1. The unemployment rate Phillips Curve has been a mainstay of inflation forecasting since its inception; capacity utilization offers a measure of the output gap; while monetary theories suggest variables such as M2 should have predictive ability for inflation. Stock and Watson (1999) consider the predictive ability of these and many other variables for inflation at a 12-month horizon, using latest-vintage data.

As GDP-deflator inflation is only available quarterly, whilst we have monthly observations on the indicator variables, it is of interest to compare MIDAS and quarterly-frequency model forecasts of inflation. As in section 5.2 we will explicitly address data vintage effects by carrying out the comparisons within a real-time framework, and also using latest-vintage data. There is some evidence that output gap measures are useful predictors using latest-vintage data but not in real-time (Orphanides and van Norden (2005)). We investigate whether the same is true of MIDAS models of inflation. Finally, given the more marked autoregressive nature of inflation, compared to output growth, our comparison will focus on MIDAS-AR versus ADL models.

The results are recorded in Table 8. The RMSFEs for the MIDAS model are markedly smaller for the latest-vintage exercise than in real-time, indicating that exercises such as that undertaken by Stock and Watson (1999) are likely to understate the uncertainty associated with forecasting inflation. For example, the RMSFE for $h = 1$ using the unemployment rate is some 20% lower using latest-vintage data. The comparison given by the ratio of the ADL to the AR model for both unemployment and capacity utilization supports Orphanides and van Norden (2005), as we find that both variables are useful for predicting inflation a year ahead ($h = 4$) using latest-vintage data (the null that the AR is as accurate as the ADL is rejected for capacity utilization, and the null of forecast encompassing is rejected for unemployment), but the evidence is weaker for the real-time exercise (especially for unemployment). M2 is useful irrespective of data vintage effects, with clear reductions in RMSFE at the longer horizons.

If instead we test for predictive ability using the MIDAS model, there is by way of contrast

clear evidence using the unemployment rate (see the RMSFE ratio of MIDAS-AR to AR, and the rejections of equal accuracy and forecast encompassing) at the longer horizons, again irrespective of the treatment of data vintage effects. Similarly, we find that the MIDAS-AR model outperforms the ADL for the unemployment rate. These results taken together suggest that the unemployment rate has little predictive ability within the quarterly-frequency ADL model, but that its predictive ability is clearly apparent using the MIDAS framework. The predictive ability of the other two indicators is apparent using the (quarterly-frequency) ADL (compared to the AR), and the temporal disaggregation afforded by the MIDAS-AR adds relatively little (compare MIDAS-AR to ADL).

6 Conclusions

The MIDAS approach has been shown to be a useful way of obtaining forecasts of quarterly output growth and inflation from monthly indicator information. Specifically, we have shown that MIDAS is useful for forecasting output growth using monthly coincident indicators, and for forecasting inflation based on the unemployment rate Phillips Curve relationship. The MIDAS approach is also an effective way of combining information from multiple predictors. We have also investigated data vintage effects, and the extent to which forecast performance depends on whether the exercises are run in real time or using latest vintage data. We also checked that the numeric differences in forecast accuracy measures between the MIDAS and quarterly-frequency models are significant by testing for equal predictive accuracy, and also for forecast encompassing.

In some respects the MIDAS approach is a simpler way of incorporating monthly information to improve quarterly forecasts of macroeconomic variables than other methods proposed in the literature, such as that of Miller and Chin (1996), for example. These methods require the estimation of separate models for quarterly series, and for monthly series. The forecasts of the two are then combined in some way, possibly using weights estimated over some prior period. By mixing monthly and quarterly-frequency data series in a single regression, forecasts from MIDAS models can be obtained in a single step. In common with Miller and Chin (1996), we also find that the inclusion of monthly information can significantly improve forecasts of the current quarter.

Another advantage of the MIDAS approach for macroeconomic forecasting arises from the restrictions imposed in the dynamic structure of the model. This simplifies the decision about which lags of the regressors to include, as only a maximum lag need be set. Choosing the specification generally requires the application of either pre-tests or information criteria, which can affect forecast performance. Against this are the possible costs to assuming that dynamic responses can be restricted to follow simple functions, such as an exponential function of the time lag. This aspect of the MIDAS approach also facilitates the combination of relatively large sets of predictors in a

regression equation, by requiring only a small number of parameters to model the effects of each on the dependent variable. We compared forecasts from a MIDAS model based on ten monthly leading indicators to forecasts generated from a ‘bagging procedure’, noting that bagging worked well in generating forecasts of US inflation (Inoue and Kilian (2005)). The comparison favoured the MIDAS approach. MIDAS appears to offer a promising approach to short-term macroeconomic forecasting.

7 Appendix I

Tests for predictive ability

Letting $MSE_{h,M}$ denote the MSFE for model M at horizon h , and $\hat{e}_{h,M,t}$ the h -step ahead forecast error for model M , then the test of equal accuracy (on MSFE) of the DL and MIDAS models at horizon h is calculated as:

$$MSE-F_h = n \left(\frac{MSE_{h,DL} - MSE_{h,MIDAS}}{MSE_{h,MIDAS}} \right).$$

The test statistic for the null that DL forecast encompasses MIDAS:

$$ENC-F_h = n \left(\frac{n^{-1} \sum_{t=1}^n (\hat{e}_{h,DL,t}^2 - \hat{e}_{h,DL,t} \hat{e}_{h,MIDAS,t})}{MSE_{h,MIDAS}} \right).$$

See McCracken (2004) and Clark and McCracken (2001), and Clark and McCracken (2005).

Calculating p-values for MSE-F and ENC-F tests by bootstrapping.

The bootstrap procedure is as follows. Firstly, estimate the quarterly-frequency DL model using the full-sample:

$$y_t = \beta_0 + \beta_1 x_{t-1} + \dots + \beta_p x_{t-p} + \varepsilon_t.$$

Here x_t is the monthly indicator series aggregated to the quarterly frequency, so that the null hypothesis is imposed - there is no useful information for predicting y_t in the monthly series beyond that contained in the indicator series aggregated to the quarterly frequency. The lag length is selected by SIC from a maximum lag of $K/3$. K depends on the empirical forecast comparison exercise. $p = 1$ is the most common selection.

Secondly, from the estimated residuals generate a vector of errors $\{\varepsilon_t^*\}$ by sampling with replacement, and from:

$$y_t^* = \hat{\beta}_0 + \hat{\beta}_1 x_{t-1} + \dots + \hat{\beta}_p x_{t-p} + \varepsilon_t^*$$

we obtain a bootstrap sample $\{y_t^*\}$ of size T ($T = 172$ for tables 2 to 5, and $T = 181$ for tables 7 and 8). Unlike Clark and McCracken (2005), we condition on the predictor series, here the observed

monthly indicator series $\{x_t^{(3)}\}$, when we generate bootstrap samples, so there is no need to specify an equation for $\{x_t^{(3)}\}$ (where $x_t = \frac{1}{3} \sum_{i=0}^2 x_{t-i/3}^{(3)}$). This is because it seems reasonable to assume that the monthly indicator series is not Granger-caused by y_t .

Thirdly, the sample is divided into estimation and forecasting samples containing, respectively, R and P observations (R and P are 116 and 56 for tables 2 to 5, and 104 and 77 for tables 7 and 8). The DL and MIDAS models are then estimated recursively beginning with the initial model estimation period R . The MIDAS parameter K is as defined in the first step, while p for the DL model is selected by SIC on each new window. The forecast errors for each forecast horizon are calculated for both models, as are the MSE-F and ENC-F test statistics.

Fourthly, the above procedure is repeated 199 times, and the p -values are computed by comparing the actual values of the test statistics to the estimates of their empirical distributions (under the null).

The adaptations to test ADL against MIDAS-AR models are straightforward.

Evaluation of Bootstrap procedure

Clark and McCracken (2005) establish the validity of the bootstrap as a means of conducting inference using the MSE-F and ENC-F test statistics in their setup. We carry out a simple Monte Carlo evaluation of the bootstrap procedure (sketched above) for comparing MIDAS and quarterly-frequency DL models' forecasts. Using samples sizes similar to the empirical part ($R = 133$ and $P = 77$), we use two data generating processes (DGPs), and estimate empirical sizes and powers.

The MIDAS DGP is:

$$y_t = 1[b(1; 1, -4)x_{t-1}^{(3)} + b(2; 1, -4)x_{t-4/3}^{(3)} + b(3; 1, -4)x_{t-5/3}^{(3)}] + \varepsilon_t.$$

This process implies that only data from one quarter is used to predict y_t . The parameters of the weight function indicate that the information of the last month ($x_{t-1}^{(3)}$) has weight equal to one, and the DGP can equivalently be written as:

$$y_t = x_{t-1}^{(3)} + \varepsilon_t.$$

For the DL DGP, we set:

$$y_t = x_{t-1} + \varepsilon_t,$$

where $x_{t-1} = \frac{1}{3} \sum_{i=0}^2 x_{t-1-i}^{(3)}$. For both DGPs, the monthly predictor series $x_t^{(3)}$ is an AR(1) with a three-month lag but shocks $v_t, v_{t-1/3}^{(3)}$ and $v_{t-2/3}^{(3)}$ have pairwise correlations of 0.8:

$$\begin{aligned} x_t^{(3)} &= 0.9x_{t-1}^{(3)} + v_t^{(3)} \\ x_{t-1/3}^{(3)} &= 0.9x_{t-1/3-1}^{(3)} + v_{t-1/3}^{(3)} \\ x_{t-2/3}^{(3)} &= 0.9x_{t-2/3-1}^{(3)} + v_{t-2/3}^{(3)} \end{aligned} .$$

For a given DGP, both MIDAS and DL models are estimated (with $K = 3$) on each replication with the simulated data (all disturbances are $N(0, 1)$), and the test statistics are calculated. The test statistics are then bootstrapped, and we estimate the p -values.

The rejection frequencies are presented in Table A1 for nominal sizes of 10%. For the DL DGP, the values in the table are estimates of size, and for the MIDAS DGP, they are estimates of power, as the null is that DL is as accurate as MIDAS (for MSE-F) and that the DL forecasts encompass those from MIDAS (for ENC-F). In accordance with the results of Clark and McCracken (2005), the ENC-F test has slightly better properties of the two, and the power decreases with the horizon. The small-sample performance of both tests appears reasonable, and indicates that they can be used in empirical forecast comparisons of MIDAS against quarterly-frequency models.

References

- Bhansali, R. J. (2002). Multi-step forecasting. In Clements, M. P., and Hendry, D. F. (eds.), *A Companion to Economic Forecasting*, pp. 206–221: Oxford: Blackwells.
- Birchenhall, C. R., Jessen, H., Osborn, D., and Simpson, P. (1999). Predicting U.S. business cycle regimes. *Journal of Business & Economic statistics*, **17**, No. 3, (July), 313–323.
- Breiman, L. (1996). Bagging predictors. *Machine Learning*, **36**, 105–139.
- Campos, J., Hendry, D. F., and Krolzig, H.-M. (2003). Consistent model selection by an automatic Gets approach. *Oxford Bulletin of Economics and Statistics*, **65**, 803–819.
- Chong, Y. Y., and Hendry, D. F. (1986). Econometric evaluation of linear macro-economic models. *Review of Economic Studies*, **53**, 671–690. Reprinted in Granger, C. W. J. (ed.) (1990), *Modelling Economic Series*. Oxford: Clarendon Press.
- Chow, G., and Lin, A. (1971). Best linear unbiased interpolation, distribution and extrapolation of time series by related time series. *Review of Economics and Statistics*, **53**, 372–375.
- Clark, T. E., and McCracken, M. W. (2001). Tests of equal forecast accuracy and encompassing for nested models. *Journal of Econometrics*, **105**, 85–110.
- Clark, T. E., and McCracken, M. W. (2005). Evaluating direct multi-step forecasts. *Econometric Reviews*, **24**, 369–404.
- Clements, M. P., and Galvão, A. B. (2006). Combining predictors and combining information in modelling: Forecasting US recession probabilities and output growth. In Milas, C., Rothman, P., and van Dijk, D. (eds.), *Nonlinear Time Series Analysis of Business Cycles. Contributions to Economic Analysis Series*, pp. 55–73: Elsevier.
- Clements, M. P., and Hendry, D. F. (1993). On the limitations of comparing mean squared forecast errors. *Journal of Forecasting*, **12**, 617–637. With discussion. Reprinted in Mills, T. C. (ed.) (1999), *Economic Forecasting. The International Library of Critical Writings in Economics*. Cheltenham: Edward Elgar.
- Croushore, D. (2006). Forecasting with real-time macroeconomic data. In Elliott, G., Granger, C., and Timmermann, A. (eds.), *Handbook of Economic Forecasting, Volume 1. Handbook of Economics 24*, pp. 961–982: Elsevier, North-Holland.
- Croushore, D., and Stark, T. (2001). A real-time data set for macroeconomists. *Journal of Econometrics*, **105**(November), 111–130.
- Croushore, D., and Stark, T. (2003). A real-time data set for macroeconomists: Does the data vintage matter?. *The Review of Economics and Statistics*, **85**(3), 605–617.

- Dhrymes, P. J. (1971). *Distributed Lags: Problems of Estimation and Formulation*. San Francisco: Holden Day.
- Diebold, F. X., and Lopez, J. A. (1996). Forecast evaluation and combination. In Maddala, G. S., and Rao, C. R. (eds.), *Handbook of Statistics*, Vol. 14, pp. 241–268: Amsterdam: North-Holland.
- Diebold, F. X., and Rudebusch, G. D. (1991). Forecasting output with the composite leading index: A real-time analysis. *Journal of the American Statistical Association*, **86**, 603–610.
- Faust, J., Rogers, J. H., and Wright, J. H. (2003). Exchange rate forecasting: The errors we’ve really made. *Journal of International Economic Review*, **60**, 35–39.
- Faust, J., Rogers, J. H., and Wright, J. H. (2005). News and noise in G-7 GDP announcements. *Journal of Money, Credit and Banking*, **37** (3), 403–417.
- Forni, M., Hallin, M., Lippi, M., and Reichlin, L. (2000). The generalized factor model: Identification and estimation. *Review of Economics and Statistics*, **82**, 540–554.
- Ghysels, E., Santa-Clara, P., and Valkanov, R. (2004a). The MIDAS touch: Mixed Data Sampling regression models. mimeo, Chapel Hill, N.C.
- Ghysels, E., Santa-Clara, P., and Valkanov, R. (2004b). Predicting volatility: Getting the most out of return data sampled at different frequencies. *Journal of Econometrics*. Forthcoming.
- Ghysels, E., Sinko, A., and Valkanov, R. (2004c). MIDAS regressions: Further results and new directions. *Econometric Reviews*. Forthcoming.
- Ghysels, E., and Wright, J. (2005). Forecasting Professional Forecasters. mimeo, Chapel Hill, N.C.
- Hendry, D. F., and Krolzig, H.-M. (2005). The properties of automatic Gets modelling. *Economic Journal*, **115**, C32–C61.
- Hendry, D. F., and Mizon, G. E. (1978). Serial correlation as a convenient simplification, not a nuisance: A comment on a study of the demand for money by the Bank of England. *Economic Journal*, **88**, 549–563. Reprinted in Hendry, D. F., *Econometrics: Alchemy or Science?* Oxford: Blackwell Publishers, 1993, and Oxford University Press, 2000.
- Inoue, A., and Kilian, L. (2005). How useful is bagging in forecasting economic time series? a case study of U.S. CPI inflation. mimeo, North Carolina State University.
- Joutz, F., and Stekler, H. O. (2000). An evaluation of the predictions of the Federal Reserve. *International Journal of Forecasting*, **16**, 17–38.
- Karamouzis, N., and Lombra, N. (1989). Federal Reserve policymaking: An overview and analysis of the policy process. *Carnegie Rochester Conference Series on Public Policy*, **30**, 7–62.
- Kilian, L. (1999). Exchange rates and monetary fundamentals: What do we learn from long-horizon

- regressions. *Journal of Applied Econometrics*, **14**, 491–510.
- Koenig, E. F., Dolmas, S., and Piger, J. (2003). The use and abuse of real-time data in economic forecasting. *The Review of Economics and Statistics*, **85**(3), 618–628.
- McCracken, M. W. (2004). Asymptotics for out-of-sample tests of causality. Discussion paper. Manuscript, University of Missouri.
- Miller, P. J., and Chin, D. M. (1996). Using monthly data to improve quarterly model forecasts. *Federal Reserve Bank of Minneapolis Quarterly Review*, **20**, 16–33. (2).
- Montgomery, A. L., Zarnowitz, V., Tsay, R. S., and Tiao, G. C. (1998). Forecasting the U.S. unemployment rate. *Journal of the American Statistical Association*, **93**, 478–493.
- Newbold, P., and Harvey, D. I. (2002). Forecasting combination and encompassing. In Clements, M. P., and Hendry, D. F. (eds.), *A Companion to Economic Forecasting*, pp. 268–283: Oxford: Blackwells.
- Orphanides, A. (2001). Monetary policy rules based on real-time data. *American Economic Review*, **91**(4), 964–985.
- Orphanides, A., and van Norden, S. (2005). The reliability of inflation forecasts based on output gaps in real time. *Journal of Money, Credit and Banking*, **37**, 583–601.
- Patterson, K. D. (1995). An integrated model of the data measurement and data generation processes with an application to consumers’ expenditure. *Economic Journal*, **105**, 54–76.
- Patterson, K. D. (2003). Exploiting information in vintages of time-series data. *International Journal of Forecasting*, **19**, 177–197.
- Robertson, J. C., and Tallman, E. W. (1998). Data vintages and measuring forecast model performance. *Federal Reserve Bank of Atlanta Economic Review, Fourth Quarter*, 4–20.
- Sims, C. A. (1974). Distributed lags. In Intriligator, M. D., and Kendrick, D. A. (eds.), *Frontiers of Quantitative Economics*, Vol. 2, Ch. 5. Amsterdam: North-Holland Publishing Company.
- Stock, J. H., and Watson, M. W. (1999). Forecasting inflation. *Journal of Monetary Economics*, **44**, 293–335.
- Stock, J. H., and Watson, M. W. (2003). How Did Leading Indicator Forecasts Perform During the 2001 Recession. *Federal Reserve Bank of Richmond, Economic Quarterly*, **89/3, Summer**, 71–90.
- Timmermann, A. (2006). Forecast combinations. In Elliott, G., Granger, C., and Timmermann, A. (eds.), *Handbook of Economic Forecasting, Volume 1. Handbook of Economics 24*, pp. 135–196: Elsevier, North-Holland.
- West, K. D. (1996). Asymptotic inference about predictive ability. *Econometrica*, **64**, 1067–1084.

West, K. D. (2006). Forecasting evaluation. In Elliott, G., Granger, C., and Timmermann, A. (eds.), *Handbook of Economic Forecasting, Volume 1. Handbook of Economics 24*, pp. 99–134: Elsevier, North-Holland.

Table 1: Description of Predictors/Indicators

Name	Description	Transformation
Leading Indicators		
spread	term spread (10year-Federal Funds)	level
stock	stock price index (500 common stocks)	$\Delta \ln$
hours	average weekly hours in manufacturing	$\ln$
claims	news claims for unemployment insurance	$\Delta \ln$
building	building permits	$\ln$
vendor	vendor performance diffusion index	$\ln$
ordersc	orders - consumer goods and materials	$\Delta \ln$
ordersn	orders - nondefense capital goods	$\Delta \ln$
expect	consumer confidence index (Michigan)	$\ln$
M2	real money supply M2	$\Delta \ln$
Coincident Indicators (available in real-time)		
Industrial Production	Total Industrial Production	$\Delta \ln$
Employment	Nonfarm payroll employment	$\Delta_{12} \ln$
Capacity Utilization	Manufacturing Utilization Capacity	$\Delta \ln$
Other Indicators (available in real-time)		
Unemployment	Civilian unemployment rate	level
M2-rt	M2 measure of the money stock	$\Delta_{12} \ln$

Notes: Leading Indicators of the Conference Board's Composite Leading Indicator (available from <http://www.wws.princeton.edu/~mwatson/publi.html>). The real-time data from the Philadelphia Fed (<http://www.phil.frb.org/econ/forecast/readow.html>). Two of the CI series are components of the Conference Board Coincident Index, namely employment and industrial production.

Table 2: Forecasting Output Growth: MIDAS versus DL models

The table presents RMSFEs for MIDAS relative to DL using each of the monthly LIs, based on recursive forecasts for the period 1989:Q1-2002:Q2 (and an estimation period starting in 1959:Q1). A forecast horizon of $h=1/3$ indicates that LI information for the second month of the quarter being forecast is used, while for $h=1$ only information up to the last quarter is employed. DL models use indicator series aggregated to the quarterly frequency, and only produce forecasts at quarterly horizons. Columns headed MSE and ENC report the bootstrap estimates of the p -values of the null of the DL forecasts being of equal accuracy (on MSE) as the MIDAS, and of the DL forecasts encompassing the MIDAS forecasts. Entries in bold denote rejections at the 10% significance level.

Using First-Released data to compute forecast errors																				
LIs	spread		stock		hours		claims		building		vendor		ordersc		ordersn		expect		M2	
MIDAS: RMSFE																				
h = 1/3	0.594		0.509		0.620		0.490		0.424		0.422		0.526		0.507		0.485		0.557	
h = 2/3	0.596		0.488		0.593		0.465		0.426		0.434		0.482		0.593		0.484		0.552	
h = 1	0.594		0.545		0.529		0.455		0.419		0.490		0.425		0.537		0.488		0.552	
h = 4/3	0.585		0.546		0.530		0.416		0.446		0.503		0.462		0.484		0.496		0.544	
h = 5/3	0.594		0.516		0.529		0.485		0.454		0.503		0.466		0.496		0.494		0.591	
h = 2	0.594		0.554		0.495		0.468		0.453		0.501		0.477		0.483		0.493		0.556	
h = 7/3	0.610		0.514		0.533		0.476		0.470		0.512		0.487		0.535		0.495		0.521	
h = 8/3	0.558		0.503		0.524		0.467		0.472		0.512		0.482		0.516		0.508		0.538	
h = 3	0.556		0.474		0.515		0.468		0.494		0.510		0.488		0.492		0.514		0.516	
h = 10/3	0.557		0.477		0.548		0.483		0.494		0.515		0.480		0.543		0.518		0.524	
h = 11/3	0.530		0.474		0.549		0.484		0.495		0.515		0.488		0.549		0.524		0.497	
h = 4	0.527		0.481		0.550		0.487		0.475		0.515		0.494		0.555		0.516		0.487	
Ratios RMSFE: MIDAS/DL																				
h=1/3;1	0.993		1.022		1.247		1.160		0.975		0.899		1.171		1.008		0.938		1.015	
h = 1;1	0.993		1.094		1.064		1.076		0.962		1.043		0.946		1.066		0.944		1.006	
h=4/3;2	0.993		1.067		1.078		0.943		0.941		1.013		0.997		1.042		1.005		1.047	
h = 2; 2	1.008		1.083		1.006		1.061		0.957		1.009		1.030		1.040		1.001		1.069	
h=7/3;3	1.113		1.087		1.060		0.995		1.007		1.000		1.035		1.062		0.962		1.035	
h = 3; 3	1.015		1.002		1.025		0.978		1.057		0.996		1.037		0.977		0.998		1.025	
h=10/3;4	1.060		0.989		1.054		1.006		1.009		0.997		0.990		1.035		0.924		1.067	
h =4, 4	1.002		0.997		1.058		1.015		0.969		0.997		1.020		1.058		0.920		0.992	
	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC
h=1/3;1	0.31	0.33	0.76	0.10	1.00	0.00	1.00	0.00	0.04	0.00	0.00	0.00	1.00	0.00	0.59	0.00	0.01	0.00	0.75	0.26
h = 1	0.06	0.09	1.00	0.36	1.00	0.00	1.00	0.00	0.01	0.00	1.00	1.00	0.00	0.00	0.97	0.04	0.02	0.02	0.48	0.03
h=4/3;2	0.27	0.12	0.96	0.30	1.00	0.98	0.00	0.00	0.00	0.00	0.96	0.93	0.36	0.00	0.94	0.02	0.80	0.03	0.99	0.04
h = 2	0.55	0.44	1.00	0.95	0.62	0.15	1.00	0.53	0.00	0.00	0.89	0.64	0.90	0.32	0.89	0.18	0.54	0.18	1.00	0.68
h=7/3;3	1.00	1.00	1.00	0.60	1.00	1.00	0.29	0.06	0.67	0.01	0.44	0.32	0.95	0.94	0.96	0.92	0.01	0.01	0.99	0.00
h = 3	0.86	0.90	0.50	0.02	0.97	0.94	0.02	0.02	1.00	0.03	0.12	0.09	0.94	0.95	0.08	0.03	0.27	0.18	0.96	0.06
h=10/3;4	1.00	1.00	0.20	0.08	1.00	0.99	0.62	0.56	0.74	0.02	0.21	0.21	0.18	0.04	0.86	0.02	0.00	0.00	1.00	0.00
h =4	0.47	0.53	0.38	0.34	1.00	1.00	0.81	0.37	0.02	0.00	0.14	0.13	0.79	0.77	0.97	0.82	0.00	0.00	0.29	0.04

Continuation of Table 2 ...

Using Last-Vintage data (2002:Q3 vintage) to compute forecast errors																				
	spread		stock		hours		claims		building		vendor		ordersc		ordersn		expect		M2	
MIDAS: RMSFE																				
h = 1/3	0.627		0.571		0.660		0.521		0.544		0.538		0.540		0.615		0.587		0.659	
h = 2/3	0.628		0.563		0.649		0.543		0.545		0.543		0.546		0.723		0.599		0.649	
h = 1	0.626		0.627		0.600		0.556		0.550		0.585		0.566		0.624		0.594		0.648	
h = 4/3	0.609		0.598		0.607		0.476		0.560		0.582		0.530		0.530		0.600		0.634	
h = 5/3	0.610		0.577		0.611		0.577		0.569		0.581		0.541		0.597		0.602		0.681	
h = 2	0.610		0.620		0.582		0.548		0.556		0.580		0.560		0.584		0.601		0.641	
h = 7/3	0.661		0.587		0.662		0.568		0.584		0.605		0.590		0.637		0.625		0.614	
h = 8/3	0.627		0.585		0.636		0.569		0.597		0.604		0.576		0.601		0.628		0.639	
h = 3	0.622		0.590		0.630		0.579		0.616		0.603		0.584		0.614		0.623		0.616	
h = 10/3	0.608		0.583		0.627		0.558		0.571		0.598		0.564		0.627		0.623		0.613	
h = 11/3	0.591		0.557		0.627		0.568		0.581		0.598		0.564		0.641		0.623		0.609	
h = 4	0.583		0.583		0.628		0.580		0.556		0.598		0.572		0.635		0.616		0.608	
Ratios RMSFE: MIDAS/DL																				
h=1/3;1	0.990		0.977		1.115		0.968		1.026		0.946		0.986		1.096		0.955		1.013	
h = 1; 1	0.987		1.073		1.014		1.034		1.037		1.028		1.032		1.112		0.966		0.996	
h=4/3;2	0.993		1.028		1.053		0.911		1.026		1.016		0.978		0.956		0.992		1.034	
h = 2; 2	0.994		1.066		1.008		1.050		1.020		1.012		1.032		1.053		0.994		1.045	
h=7/3;3	1.071		1.052		1.077		0.977		1.062		1.006		1.023		1.066		0.999		1.015	
h = 3; 3	1.008		1.057		1.026		0.994		1.120		1.003		1.014		1.027		0.995		1.019	
h=10/3;4	1.048		0.989		1.027		0.988		1.018		0.998		0.991		1.011		0.940		0.986	
h = 4; 4	1.005		0.990		1.029		1.025		0.990		0.999		1.005		1.023		0.928		0.978	
	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC
h=1/3;1	0.21	0.25	0.12	0.03	1.00	0.00	0.01	0.00	0.93	0.00	0.00	0.00	0.15	0.00	1.00	0.00	0.03	0.00	0.70	0.37
h = 1	0.04	0.06	0.98	0.53	0.82	0.00	0.97	0.00	0.99	0.02	1.00	0.98	0.92	0.02	1.00	0.71	0.10	0.06	0.25	0.04
h=4/3;2	0.25	0.12	0.85	0.12	0.99	0.90	0.00	0.00	0.95	0.01	0.97	0.97	0.08	0.01	0.02	0.00	0.21	0.06	0.98	0.07
h = 2	0.05	0.07	0.99	0.94	0.72	0.41	0.99	0.69	0.94	0.01	0.94	0.86	0.92	0.69	0.96	0.66	0.30	0.16	0.99	0.50
h=7/3;3	1.00	1.00	0.94	0.37	1.00	1.00	0.02	0.02	1.00	0.56	0.85	0.85	0.89	0.85	0.97	0.97	0.47	0.06	0.89	0.00
h = 3	0.70	0.76	0.99	0.85	0.98	0.96	0.19	0.14	1.00	1.00	0.72	0.54	0.76	0.78	0.80	0.56	0.14	0.13	0.90	0.13
h=10/3;4	1.00	0.99	0.20	0.13	0.95	0.82	0.12	0.14	0.92	0.07	0.29	0.30	0.18	0.06	0.63	0.02	0.00	0.00	0.22	0.00
h = 4	0.66	0.73	0.23	0.30	0.97	0.87	0.94	0.84	0.19	0.03	0.27	0.24	0.55	0.48	0.79	0.51	0.00	0.00	0.07	0.04

Table 3: Forecasting Output Growth: MIDAS-AR versus ADL models

The table presents RMSFEs for MIDAS relative to DL using the LIs and allowing for autoregressive dynamics using output growth vintages from 1989:Q1 to 2002:Q2 (sample starts at 1959:Q1). The RMSFEs are calculated using first-release data as actuals. p -values of the tests of equal forecast accuracy (MSE) and forecasting encompassing (ENC) are reported, where the null is either the ADL or the AR (depending on the denominator of the RMSFE ratios).

	spread		stock		hours		claims		building		vendor		ordersc		ordersn		expect		M2	
MIDAS-AR: RMSFE																				
h=1	0.568		0.432		0.477		0.468		0.395		0.434		0.421		0.511		0.473		0.533	
h=2	0.556		0.534		0.492		0.463		0.449		0.473		0.490		0.463		0.486		0.536	
h=3	0.586		0.475		0.506		0.473		0.489		0.496		0.472		0.490		0.529		0.540	
h=4	0.524		0.479		0.536		0.496		0.499		0.488		0.477		0.552		0.522		0.494	
Ratios RMSFE: MIDAS-AR/MIDAS																				
h=1	0.955		0.792		0.903		1.029		0.943		0.887		0.991		0.953		0.970		0.965	
h=2	0.936		0.964		0.994		0.989		0.990		0.943		1.027		0.959		0.985		0.965	
h=3	1.055		1.003		0.982		1.010		0.990		0.974		0.968		0.996		1.028		1.046	
h=4	0.994		0.996		0.975		1.019		1.050		0.947		0.965		0.994		1.011		1.014	
Ratios RMSFE: MIDAS-AR/ADL																				
h=1	1.007		0.992		1.023		1.109		0.935		0.988		0.934		1.138		0.933		1.043	
h=2	0.964		1.095		1.006		1.057		0.952		0.976		1.068		0.950		0.996		1.058	
h=3	1.022		1.015		1.012		0.993		1.049		1.003		1.007		0.993		0.987		1.050	
h=4	0.973		1.009		1.027		1.037		1.041		0.981		1.000		1.083		0.896		0.995	
	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC
h=1	0.90	0.49	0.19	0.00	0.97	0.00	1.00	0.00	0.00	0.00	0.05	0.02	0.00	0.00	1.00	0.32	0.00	0.00	1.00	0.31
h=2	0.03	0.03	1.00	0.44	0.67	0.59	1.00	0.69	0.00	0.00	0.01	0.02	1.00	0.80	0.01	0.00	0.27	0.02	0.99	0.20
h=3	0.94	0.97	0.72	0.01	0.85	0.87	0.17	0.10	1.00	0.01	0.79	0.63	0.48	0.37	0.23	0.03	0.05	0.05	0.99	0.21
h=4	0.05	0.07	0.57	0.16	0.96	0.94	1.00	0.64	1.00	0.11	0.01	0.01	0.35	0.22	1.00	0.99	0.00	0.00	0.23	0.05
Ratios RMSFE: ADL/AR																				
h=1	1.346		1.039		1.114		1.008		1.009		1.049		1.075		1.073		1.211		1.219	
h=2	1.251		1.058		1.061		0.951		1.023		1.051		0.996		1.058		1.059		1.100	
h=3	1.205		0.983		1.049		1.000		0.978		1.039		0.985		1.036		1.125		1.080	
h=4	1.129		0.996		1.095		1.003		1.005		1.044		1.000		1.068		1.222		1.041	
	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC
h=1	1.00	0.98	0.98	0.01	1.00	1.00	0.73	0.00	0.68	0.00	0.98	0.98	1.00	0.01	1.00	0.88	1.00	0.01	1.00	0.03
h=2	1.00	0.04	0.99	0.04	1.00	0.34	0.00	0.01	0.87	0.03	0.95	0.97	0.13	0.21	1.00	0.96	0.97	0.04	1.00	0.00
h=3	1.00	0.08	0.06	0.02	0.98	0.22	0.25	0.37	0.05	0.03	0.92	0.95	0.06	0.10	0.97	0.43	1.00	0.94	0.99	0.01
h=4	1.00	0.10	0.16	0.04	1.00	0.20	0.43	0.48	0.43	0.11	0.94	0.94	0.28	0.38	1.00	0.13	1.00	0.61	0.96	0.00
Ratios RMSFE: MIDAS-AR/AR																				
h=1	1.355		1.031		1.140		1.118		0.943		1.037		1.005		1.221		1.130		1.272	
h=2	1.207		1.159		1.067		1.005		0.974		1.026		1.064		1.006		1.055		1.163	
h=3	1.231		0.998		1.061		0.993		1.026		1.042		0.992		1.029		1.110		1.134	
h=4	1.098		1.005		1.124		1.040		1.046		1.024		1.000		1.157		1.095		1.036	
	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC
h=1	1.00	0.50	0.89	0.00	1.00	1.00	1.00	0.00	0.00	0.00	0.96	0.96	0.40	0.00	1.00	1.00	1.00	0.00	1.00	0.09
h=2	1.00	0.03	1.00	0.34	0.96	0.18	0.36	0.05	0.06	0.03	0.78	0.79	0.98	0.68	0.34	0.24	0.93	0.02	1.00	0.00
h=3	1.00	0.08	0.22	0.05	0.94	0.37	0.15	0.18	0.77	0.06	0.89	0.91	0.10	0.14	0.75	0.32	1.00	0.84	1.00	0.02
h=4	0.98	0.09	0.36	0.04	1.00	0.21	0.86	0.61	0.87	0.06	0.71	0.67	0.26	0.19	1.00	0.72	0.99	0.94	0.86	0.01

Table 4: Forecasting Output Growth using Multiple LI models and Forecast Combination

This table presents RMSFEs for the multiple indicator MIDAS and DL models (M-MIDAS and M-DL, respectively) using output growth vintages from 1989:Q1 to 2002:Q2 (sample starts at 1959:Q1). RMSFEs are also shown for equal-weighted combinations of the forecasts from the single-indicator models of Table 2 (denoted by the prefix 'Mean'). The models in the first panel use all 10 LIs, while in the second panel only a subset of the indicators are included. The forecasts errors are computed using the first-release data as actual values. The p -values of the tests of equal accuracy (MSE) and forecasting encompassing (ENC) have the M-DL model as the null.

	RMSFEs:				Ratios RMSFE:		H ₀ : M-DL X M-MIDAS	
	M-MIDAS	M-DL	Mean MIDAS	Mean DL	M-Midas/ M-DL	M-Midas/ Mean MIDAS	MSE	ENC
Using information of all leading indicators								
h=1	0.478	0.594	0.420	0.418	0.804	1.138	0.00	0.00
h=2	0.455	0.651	0.451	0.450	0.699	1.008	0.05	0.05
h=3	0.559	0.641	0.470	0.470	0.872	1.191	0.90	0.50
h=4	0.630	0.685	0.475	0.479	0.920	1.327	0.90	0.70
Using information of stock, building and expectations								
h=1	0.508	0.572	0.442	0.477	0.888	1.149	0.00	0.00
h=2	0.560	0.671	0.471	0.480	0.834	1.188	0.33	0.00
h=3	0.498	0.620	0.475	0.475	0.803	1.048	0.11	0.01
h=4	0.488	0.576	0.472	0.479	0.848	1.035	0.20	0.11

**Table 5: Forecasting Output Growth Combining Leading Indicators:
M-MIDAS versus Bagging**

This table compares RMSFEs for M-MIDAS (as in Table 4, same sample) with 'bagging' forecasts. Bagging forecasts are computed with a critical value of 2.57 for the t -test of the exclusion decision. The M-MIDAS models are estimated using $K=15$, 12 and 6 as the maximum number of monthly lags. The bagging forecasts are computed with a matching maximum quarterly lag length of 5, 4 or 2. Also shown is the ratio of the bagging MSFE to the RMSFE for the equal-weighted (mean) combination of single-regressor ADL model forecasts. The models in the first panel utilize information on all the 10 leading indicators, while the results in the second panel use only a subset. The forecasts errors are computed using first-release data as actuals.

	RMSFEs:	Ratios RMSFE:			
	Bagging (K =15)	Bagging/ M-MIDAS (K=15)	Bagging/ M-MIDAS (K=12)	Bagging/ M-MIDAS (K=6)	Bagging/ Mean ADL (K=15)
Using information of all leading indicators					
h=1	0.533	1.193	1.371	1.352	1.305
h=2	0.526	1.074	1.021	1.079	1.179
h=3	0.530	0.970	0.919	0.995	1.132
h=4	0.562	0.914	0.947	0.995	1.191
Using information of stock, building and expectations					
h=1	0.518	1.019	0.956	1.029	1.157
h=2	0.571	1.021	0.916	0.926	1.252
h=3	0.574	1.154	1.075	1.005	1.217
h=4	0.560	1.146	1.000	1.115	1.174

Table 6. Outline of the use of the monthly real-time data sets in model estimation and forecast generation using the coincident indicators.

d.s. is shorthand for ‘real-time data set’. The evaluation uses the first-release actuals, so that growth in 00:Q1 is evaluated using the first-release figure for 00Q1 in the 00Q2 d.s., and growth in 00:Q2 is evaluated using the first-release figure in the 00Q3 d.s.

Period being forecast	Horizon h	Model estimated on data up to:	Latest x data used
00Q1	1/3	y_{99Q4} on $x_{99:M11}$ (from 00M3 d.s.)	x_{00M2} (from 00M3 d.s.)
00Q1	2/3	y_{99Q4} on $x_{99:M10}$ (from 00M2 d.s.)	x_{00M1} (from 00M2 d.s.)
00Q1	1	y_{99Q4} on $x_{99:M9}$ (from 00M1 d.s.)	x_{99M12} (from 00M1 d.s.)
00Q2	4/3	y_{99Q4} on $x_{99:M8}$ (from 00M3 d.s.)	x_{00M2} (from 00M3 d.s.)
00Q2	5/3	y_{99Q4} on $x_{99:M7}$ (from 00M2 d.s.)	x_{00M1} (from 00M2 d.s.)
00Q2	2	y_{99Q4} on $x_{99:M6}$ (from 00M1 d.s.)	x_{99M12} (from 00M1 d.s.)

Table 7: Forecasting Output Growth using Real-Time Coincident Indicators: MIDAS versus DL models

The table presents RMSFE comparisons for MIDAS versus DL models for three indicators with monthly vintages from 1985:M1 to 2005:M2 (sample starts at 1959:M1). In the columns labeled “RT” real-time data is employed for output growth (quarterly vintages from 1985:Q1 to 2005:Q1) and the indicators, and the actuals are the first-release data. The columns “No RT” use last vintage data for output growth and the indicators (2005:Q2), and for computing forecast errors. For $h=1/3$, for example, information up to the second month of the current quarter is employed for forecasting the current quarter, while for $h=1$ only information up to the previous quarter is employed. The last panel presents p -values for tests of equal forecast accuracy (MSE) and forecasting encompassing (ENC), where rejections (emboldened at 10% signif. level.) are in favour of MIDAS.

	Industrial Production				Employment				Capacity utilization			
	RT		No RT		RT		No RT		RT		No RT	
h = 1/3	0.351		0.405		0.427		0.490		0.435		0.501	
h = 2/3	0.382		0.461		0.430		0.494		0.453		0.515	
h = 1	0.431		0.543		0.431		0.508		0.458		0.520	
h = 4/3	0.393		0.483		0.472		0.548		0.458		0.522	
h = 5/3	0.439		0.521		0.461		0.529		0.455		0.523	
h = 2	0.445		0.521		0.464		0.536		0.463		0.513	
h = 7/3	0.429		0.523		0.466		0.538		0.462		0.527	
h = 8/3	0.436		0.508		0.461		0.538		0.470		0.529	
h = 3	0.466		0.521		0.461		0.531		0.471		0.532	
h = 10/3	0.451		0.520		0.468		0.518		0.473		0.525	
h = 11/3	0.451		0.520		0.466		0.522		0.473		0.525	
h = 4	0.471		0.529		0.464		0.529		0.473		0.525	
Ratios RMSFE: MIDAS/DL												
h=1/3;1	0.781		0.772		0.977		0.988		0.990		0.964	
h=1;1	0.958		1.036		0.986		1.025		1.044		1.000	
h=4/3;2	0.882		0.939		1.037		1.044		0.973		0.983	
h=2;2	0.998		1.013		1.018		1.021		0.984		0.966	
h=7/3;3	0.946		1.003		0.985		1.000		0.976		0.991	
h=3;3	1.027		1.000		0.977		0.987		0.994		1.000	
h=10/3;4	0.985		1.002		0.987		0.980		0.997		0.995	
h=4;4	1.029		1.019		0.977		1.000		0.996		0.994	
	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC
h=1/3;1	0.00	0.00	0.00	0.00	0.01	0.00	0.04	0.00	0.53	0.01	0.16	0.02
h=1;1	0.00	0.00	0.98	0.00	0.05	0.00	0.99	0.01	0.98	0.13	0.54	0.04
h=4/3;2	0.00	0.00	0.00	0.00	0.99	0.58	1.00	0.98	0.13	0.19	0.32	0.39
h=2;2	0.33	0.12	0.81	0.38	0.93	0.02	0.94	0.20	0.06	0.03	0.00	0.00
h=7/3;3	0.00	0.00	0.53	0.06	0.03	0.00	0.39	0.03	0.03	0.03	0.16	0.16
h=3;3	0.97	0.96	0.43	0.53	0.00	0.00	0.02	0.00	0.10	0.05	0.32	0.24
h=10/3;4	0.05	0.03	0.51	0.03	0.02	0.00	0.00	0.00	0.22	0.31	0.17	0.29
h=4;4	0.97	0.97	0.90	0.87	0.01	0.01	0.40	0.23	0.15	0.18	0.11	0.14

Table 8: Forecasting Inflation in Real-Time: MIDAS-AR versus ADL models

The table presents RMSFE comparisons of MIDAS and ADL models using both real-time data and latest vintage data ('RT' and 'No RT'). The real-time data are quarterly vintages of the inflation and indicators from 1985:Q1 to 2005:Q1 (starts at 1959:Q1), with indicators available monthly. In the first part of the table the tests of equal forecast accuracy (MSE) and forecasting encompassing (ENC) are computed with ADL under the null. In the second part, the model under the null is the AR. Rejection implies the MIDAS is more accurate / not forecast encompassed in the first case, and favours either the ADL (indicator has predictive ability) or the MIDAS (indicator with temporal disaggregation) in the second case.

	Capacity utilization				Unemployment				M2-rt			
MIDAS-AR: RMSFE												
	RT		No RT		RT		No RT		RT		No RT	
h = 1	0.280		0.215		0.275		0.216		0.268		0.219	
h = 2	0.322		0.249		0.316		0.242		0.318		0.242	
h = 3	0.316		0.255		0.309		0.246		0.318		0.259	
h = 4	0.314		0.236		0.308		0.233		0.312		0.247	
Ratios RMSFE: MIDAS-AR/ADL												
h = 1	1.013		1.001		0.974		0.989		1.011		1.021	
h = 2	1.019		1.035		0.989		0.981		1.030		0.991	
h = 3	1.018		1.030		0.970		0.952		1.087		1.092	
h = 4	0.980		1.004		0.924		0.946		1.151		1.091	
	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC
h = 1	0.81	0.83	0.55	0.34	0.02	0.04	0.12	0.17	0.80	0.90	0.91	0.96
h = 2	0.83	0.85	0.97	0.93	0.25	0.28	0.13	0.13	0.93	0.58	0.42	0.17
h = 3	0.82	0.69	0.90	0.60	0.10	0.12	0.03	0.05	1.00	0.18	1.00	0.04
h = 4	0.20	0.12	0.60	0.12	0.01	0.03	0.03	0.03	1.00	0.43	1.00	0.03
Ratios RMSFE: ADL/AR												
h = 1	0.989		0.983		1.022		1.006		0.959		0.989	
h = 2	0.981		0.967		1.006		1.003		0.975		0.994	
h = 3	0.975		0.956		1.017		1.016		0.931		0.932	
h = 4	0.969		0.923		1.030		0.983		0.838		0.905	
	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC
h = 1	0.10	0.11	0.06	0.05	0.94	0.91	0.58	0.43	0.00	0.01	0.08	0.10
h = 2	0.09	0.09	0.03	0.04	0.34	0.36	0.26	0.20	0.04	0.04	0.16	0.07
h = 3	0.10	0.09	0.05	0.03	0.51	0.40	0.50	0.23	0.01	0.00	0.01	0.00
h = 4	0.10	0.05	0.02	0.00	0.60	0.27	0.12	0.03	0.00	0.00	0.01	0.00
Ratios RMSFE: MIDAS-AR/AR												
h = 1	1.001		0.984		0.995		0.996		0.970		1.011	
h = 2	1.000		1.001		0.995		0.985		1.004		0.985	
h = 3	0.992		0.985		0.986		0.967		1.012		1.018	
h = 4	0.950		0.927		0.951		0.930		0.964		0.988	
	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC	MSE	ENC
h = 1	0.34	0.25	0.02	0.02	0.12	0.13	0.14	0.12	0.00	0.00	0.85	0.87
h = 2	0.23	0.06	0.28	0.03	0.17	0.21	0.08	0.07	0.43	0.41	0.07	0.07
h = 3	0.12	0.02	0.05	0.01	0.11	0.15	0.05	0.04	0.59	0.47	0.73	0.65
h = 4	0.02	0.01	0.01	0.00	0.04	0.05	0.02	0.03	0.04	0.06	0.19	0.13

Table A1: Rejection Frequencies of the MSE and ENC tests using Bootstrapped p -values.

The DL DGP is $y_t = x_{t-1} + \varepsilon_t$, $\varepsilon_t \sim N(0,1)$. The MIDAS DGP is $y_t = x_{t-1}^{(3)} + \varepsilon_t$. The DGP of the monthly predictor x_t is an AR(1) with a three-month lag but with correlated disturbances. The quarterly predictor is the average over the ‘months’ in the quarter. The number of bootstrap replications is 150, and the number of Monte Carlo replications is 200. The total number of observations is 200 and the out-of-sample period has 77 observations. Recursive forecasts are used to compute MSE and ENC statistics for each forecast horizon with the DL model under the null. The p -values of the statistics are computed by a bootstrap on each replication and rejections are decided with a 10% significance level.

DGPs:	DL		MIDAS	
	MSE	ENC	MSE	ENC
h = 1	0.100	0.125	0.970	0.990
h = 2	0.165	0.130	0.825	0.870
h = 3	0.170	0.120	0.655	0.715
h = 4	0.145	0.110	0.545	0.555