

UNIVERSITY OF HAWAII
1907
MAU KE EA O KA 'ĀINA I KA PŪ
UHERO

THE ECONOMIC RESEARCH ORGANIZATION
AT THE UNIVERSITY OF HAWAII

NOWCASTING TOURISM INDUSTRY
PERFORMANCE USING HIGH FREQUENCY
COVARIATES

BY

ASHLEY HIRASHIMA, JAMES JONES, CARL S. BONHAM,
AND PETER FULEKY

Working Paper No. 2015-3

September 28, 2015

UNIVERSITY OF HAWAII AT MANOA
2424 MAILE WAY, ROOM 540 • HONOLULU, HAWAII 96822
WWW.UHERO.HAWAII.EDU

WORKING PAPERS ARE PRELIMINARY MATERIALS CIRCULATED TO STIMULATE
DISCUSSION AND CRITICAL COMMENT. THE VIEWS EXPRESSED ARE THOSE OF
THE INDIVIDUAL AUTHORS.

Nowcasting Tourism Industry Performance Using High Frequency Covariates

Ashley Hirashima, James Jones, Carl S. Bonham, and Peter Fuleky *

September 28, 2015

Abstract

We evaluate the short term forecasting performance of methods that systematically incorporate high frequency information via covariates. Our results indicate that including timely intra-period data into the forecasting process results in significant gains in predictive accuracy compared to relying exclusively on low frequency aggregates. Anticipating growing popularity of these tools among empirical analysts, we offer practical implementation guidelines to facilitate their adoption.

Keywords: Nowcast; Ragged edge; Mixed frequency models.

JEL classifications: C22, C82, L83, Z32

*University of Hawaii Economic Research Organization and Department of Economics, University of Hawaii, 540 Saunders Hall, 2424 Maile Way, Honolulu, HI 96822. Email: fuleky@hawaii.edu

1 Introduction

Where tourism is a key component of the local economy, firms, government agencies, and other organizations often update their predictions of tourism activity within a quarter. Inevitably, several stages of the forecasting process involve frequency conversions. The underlying dataset usually contains economic indicators released with different lags and sampled at various frequencies. Tourism agencies use short term forecasts to plan their operations to satisfy multi-period performance targets. Government organizations feed high frequency tourism forecasts into macroeconomic models that evolve at a lower frequency. Such practical issues arising in the forecasting process lead to two questions: (1) how can data released with different lags and frequencies be combined in the generation of multi-period forecasts, and (2) what benefits can be derived from such combinations.

Forecasters often turn to classical methods to predict industry performance (see for example Athanasopoulos et al., 2011; Dwyer et al., 2012). Methods, like Exponential Smoothing (ES) and Autoregressive Integrated Moving Average (ARIMA) models, only use the history of the variable of interest for prediction. Others, such as Autoregressive Distributed Lag (ADL) models, go a step further by incorporating a set of explanatory variables. A common feature of these models is that they all operate at a single frequency. For example, a typical ADL model may predict quarterly tourist arrivals using its own lags and a set of explanatory variables, all indexed at

the quarterly frequency. If some of the explanatory variables are available at a higher frequency, such as monthly or daily, they must be aggregated to the quarterly frequency before they can be included in the model. The disadvantage of such aggregation is that it may discard valuable information, especially at the end of the sample when only a partial quarter, say one or two months, of data is available.

Some forecasters make adjustments to their near-term forecasts to incorporate information from incomplete periods. This process can be cumbersome and, by definition, subjective. Several approaches have been recently developed that avoid such limitations by *directly* using high frequency regressors to predict a low frequency variable of interest. This is a rapidly growing area of research with over fifty studies published on the topic in the last decade (see for example Camacho et al., 2013). While our paper is the first one to introduce some of the new methods to the tourism literature, macroeconomic forecasting studies have already shown that they often improve predictions for the current period (nowcasting). An additional benefit of a model-based approach to incorporate high frequency information is the streamlining of the forecasting process, with near-term predictions generated through a mechanical process, as opposed to a series of subjective adjustments. It is no surprise then that these tools are quickly gaining popularity in applied research: among others, the Federal Reserve Banks of Atlanta and Philadelphia publish nowcasts of economic activity on their websites.

Direct transfer of information across frequencies can be accomplished in a

number of ways. The Mixed Data Sampling (MIDAS) regression of Ghysels et al. (2004, 2007) uses high frequency regressors to predict a low frequency variable of interest. The coefficients in a MIDAS model are determined nonlinearly by a small set of hyperparameters. This solves the parameter proliferation problem that can occur in the case of a large frequency mismatch between the predictors and the variable of interest, but it means that the regression can no longer be estimated by Ordinary Least Squares (OLS).

Rather than estimating the hyperparameters nonlinearly, Foroni et al. (2015) introduced an Unrestricted Mixed Data Sampling (U-MIDAS) regression that only requires OLS estimation. They show that when the frequency mismatch is small, as in the case of using monthly data to predict a quarterly series, U-MIDAS models tend to outperform MIDAS models. Factor models based on the Kalman filter have also been used in prediction (see for example Fuleky and Bonham, 2015), but are beyond the scope of our study. Their main purpose is the estimation of an unobserved overall business cycle variable and its fluctuations. While, Bai et al. (2013) found that the Kalman filter based methods and the MIDAS models analyzed in our paper have similar forecasting performances, the implementation of the former is more complex.

There is a large body of literature that evaluates the performance of mixed-frequency methods in real-world applications. Most studies show that high frequency variables improve predictive accuracy. They include the survey by Camacho et al. (2013) investigating a variety of short-term forecasting

methods, and the study by Jansen et al. (2012) evaluating eleven different models to forecast real GDP for several European countries. Nonetheless, these authors point out that the performance differences across mixed frequency specifications are marginal, and while monthly predictors increase the accuracy of nowcasts when the dependent series is measured at the quarterly frequency, the gains in performance appear to be muted for forecasts one or two quarters ahead. We know of only one tourism related study dealing with a similar topic. Bangwayo-Skeete and Skeete (2015), found that weekly Google search data help to predict monthly tourist arrivals to certain Caribbean destinations. Finally, some studies do not find a significant improvement in predictive accuracy from the use of high frequency covariates. Baumeister et al. (2015) note that forecast precision depends on whether the high frequency data provides a useful signal or simply introduces additional noise.

We evaluate several classical and modern forecasting methods in a tourism setting to assess their practical advantages and disadvantages. Specifically, we use monthly regressors to produce nowcasts and one-quarter ahead forecasts of tourist arrivals to Hawaii, and, in a separate exercise, we predict quarterly earnings in the accommodation and food services industry in Hawaii. Our results are largely in line with the existing literature. We find that, relative to a low frequency baseline model, incorporating high frequency information results in an overall improvement in predictive accuracy, both for nowcasting and one-period ahead forecasting. However, differences in accu-

racy across mixed frequency models tend to be small. Therefore, while practitioners should not ignore high frequency information, they should use those methods that can be applied to the particular problem at the least cost. Our results also indicate that gains in predictive accuracy are the greatest when the high frequency information is contemporaneous. Consequently, mixed-frequency methods are most valuable when the high frequency regressors are available with relatively short publication lags.

The layout for the rest of the paper is as follows. In Section 2, we introduce the mixed frequency models used in our analysis. In Section 3, we describe our dataset, apply the models to Hawaii’s tourism sector, and discuss our results. Our concluding remarks are in Section 4.

2 Methods

Autoregressive distributed lag (ADL) models benefit from multiple sources of information: contemporaneous and lagged observations on the variable of interest and its predictors. In contrast, pure time series approaches—such as exponential smoothing and ARIMA models—only consider lags of the variable of interest and neglect the information in associated variables, while contemporaneous multiple regressions ignore any dynamics. An ADL model nests both of these building blocks and is therefore a more general and

versatile tool for forecasting. A typical ADL model can be written as

$$y_t = \alpha + \phi(L)y_{t-1} + \beta(L)x_t + \epsilon_t, \quad \text{for } t = \delta + 1, \delta + 2 \dots, \quad (1)$$

where y_t and x_t are the variable of interest and a predictor, respectively. The lag polynomials, $\phi(L) = \sum_{i=0}^{ymax} \phi_i L^i$ and $\beta(L) = \sum_{j=0}^{xmax} \beta_j L^j$, have lengths $ymax$ and $xmax$ whose optimal values can be determined using standard model selection criteria. Lags of the predictors are accommodated in the model by applying a $\delta = \lceil \max(ymax + 1, xmax) \rceil$ offset to the initial period, where $\lceil z \rceil$ is the *ceiling* operator producing the smallest integer not less than z . Additional regressors can be included at the cost of further notational complexity.

A forecast produced on forecast date T for a horizon h is based on the estimated relationship between the response variable and the h^{th} -and-greater lags of the predictors in equation (1). Reporting delays can result in missing observations for recent periods in a vintage T dataset, a phenomenon sometimes called the “ragged edge” problem. This problem can be solved by accounting for the lag of the most recent observation of each variable relative to time T . Denoting this lag by Δ_y and Δ_x for y and x , respectively, we obtain a forecast for horizon h by evaluating

$$\hat{y}_{T+h} = \hat{\alpha} + \hat{\phi}(L)y_{T-\Delta_y} + \hat{\beta}(L)x_{T-\Delta_x}, \quad (2)$$

with coefficients previously estimated in a regression that maintains equiva-

lent lags of the predictors *relative to the response variable*

$$y_t = \alpha + \phi(L)y_{t-h-\Delta_y} + \beta(L)x_{t-h-\Delta_x} + \epsilon_t, \quad \text{for } t = \delta + 1, \delta + 2 \dots, \quad (3)$$

where the δ offset applied to the initial period is determined by the maximum lag of the regressors. Note, the time subscript denotes the reference period for a particular observation and not the release date.

Equation (2) uses the relationship captured between the left and right hand sides of equation (3) to map predictor information available at time T into a forecast for horizon h . The model can also accept forward looking indicators by letting $-\Delta_x$ represent the lead time. For example, a predictor containing information about the period ahead of time t can be incorporated into equations (2) and (3) as x_{T+1} and x_{t-h+1} , respectively.

ADL models can be extended to map high frequency information into forecasts of low frequency variables. In the following we describe forecasting models that combine data sampled at different frequencies.

2.1 Mixed Frequency Models

Mixed frequency models are typically used when the variable of interest evolves at a low frequency while the predictors are observed at a high frequency. To illustrate, let y_t and x_t denote two economic indicators sampled at the quarterly and monthly frequency, respectively. Let us initially assume that there is no delay in the reporting of new information, with the data re-

lease occurring immediately at the end of the given period. This assumption helps us simplify the exposition, but we will relax it later.

The time index t refers to the end of a particular period. Without loss of generality (see Fuleky, 2012), we set the unit of time to a quarter, so that it matches the frequency of the response variable. Consequently, the monthly observations are indexed with $t = \frac{1}{3}, \frac{2}{3}, 1, 1\frac{1}{3}, \dots$, while the quarterly ones with $t = 1, 2, \dots$, and $L^{1/3}$ and L represent a monthly and a quarterly lag, respectively.¹ The relationship between the quarterly response variable, its own lags, and the lags of the monthly predictors can be estimated using a simple mixed frequency regression

$$y_t = \alpha + \phi(L)y_{t-1} + \beta(L^{1/3})x_t + \epsilon_t, \quad \text{for } t = \delta + 1, \delta + 2, \dots, \quad (4)$$

where $\phi(L) = \sum_{i=0}^{y_{max}} \phi_i L^i$ and $\beta(L^{1/3}) = \sum_{j=0}^{x_{max}} \beta_j L^{j/3}$. The time increments of unit length indicate that the rows of the data set are a quarter, or three months, apart. Equation (4) illustrates that the quarterly regressand can be directly related to its own lags and the lags of the monthly regressors. Therefore it is unnecessary to aggregate the monthly regressors beforehand to match the quarterly frequency of the regressand.

The forecast date T can fall at the end of any month. Hence, forecasts can be produced for horizons $h = \{0, \frac{1}{3}, \frac{2}{3}, 1, 1\frac{1}{3}, \dots\}$, where the first three

¹The fractional lag operator, $L^{1/3}$, is only applied to monthly indicators. For example, while x_2 is the value of the monthly indicator in the last month of quarter 2, $L^{1/3}x_2 = x_{1\frac{2}{3}}$ and $L^{2/3}x_2 = x_{1\frac{1}{3}}$ are the values of x in the second and first months of quarter 2, respectively.

specify predictions for the current quarter and are usually called “nowcasts”. A forecast for a certain horizon h necessitates a relationship between the response variable and the h^{th} -and-greater lags of the predictors in equation (4). Using the coefficients estimated in such a regression and the data available at the forecast date T , we obtain a forecast for horizon h , \hat{y}_{T+h} , by evaluating

$$\hat{y}_{T+h} = \hat{\alpha} + \hat{\phi}(L)y_{T-\Delta} + \hat{\beta}(L^{1/3})x_T, \quad (5)$$

where Δ denotes the lag of the most recent quarterly observation that is available as of time T . For example, if the forecast date is at the end of the first or second month of the quarter, then Δ is $\frac{1}{3}$ or $\frac{2}{3}$, respectively. If the forecast date coincides with the end of the quarter, then $\Delta = 0$.

Let us now relax the assumption that data is released immediately at the end of a particular period. Denoting the release lag by Δ_y and Δ_x for y and x , respectively, we obtain a forecast for horizon h by evaluating

$$\hat{y}_{T+h} = \hat{\alpha} + \hat{\phi}(L)y_{T-\Delta_y} + \hat{\beta}(L^{1/3})x_{T-\Delta_x}, \quad (6)$$

with coefficients previously estimated in a regression that maintains equivalent lags of the predictors *relative to the response variable*

$$y_t = \alpha + \phi(L)y_{t-h-\Delta_y} + \beta(L^{1/3})x_{t-h-\Delta_x} + \epsilon_t, \quad \text{for } t = \delta + 1, \delta + 2, \dots \quad (7)$$

As in the case of the single-frequency ADL model, this framework can also

accept forward looking predictors by letting $-\Delta_x$ represent the lead time.

Foroni et al. (2015) called the model described above an unrestricted MIDAS, or UMIDAS, model. Because the lag-structure of equation (7) is unconstrained, it potentially requires the estimation of a large number of parameters. To avoid parameter proliferation, we consider various constraints on the lag-polynomials $\phi(L)$ and $\beta(L^{1/3})$. Specifically, we examine the performance of mixed frequency models under the following restrictions:

Autometrics-based model selection relies on the automatic model selection features of the PcGive software (see Hendry and Krolzig, 2004) to identify an optimal set of predictors and their lags. Autometrics uses a wide variety of diagnostic tools to simplify a general unconstrained model.

Non-overlapping predictors are obtained by separating highly correlated regressors with similar information content based on their availability at time T . The regressor with the most recent observation is incorporated into the model with lags up until the period for which an observation for another regressor is available. This second regressor is incorporated into the model with lags up until the period for which an observation for yet another regressor is available, and so on. Such lag structure implied by data availability is more parsimonious than using all series in parallel.

MIDAS of Ghysels et al. (2007) eliminates parameter proliferation by

defining $\beta(L^{1/3})$ as an exponential Almon lag polynomial

$$\beta(L^{1/3}, \theta) = \beta_0(\theta)L^{0/3} + \beta_1(\theta)L^{1/3} + \dots + \beta_{y_{max}}(\theta)L^{x_{max}/3}, \quad (8)$$

where

$$\beta_j(\theta) = \frac{e^{\theta_1 j + \theta_2 j^2}}{\sum_{j=0}^{x_{max}} e^{\theta_1 j + \theta_2 j^2}} \quad (9)$$

so that the estimated values of only two hyper-parameters, θ_1 and θ_2 , determine the distribution of weights along the lag polynomial. Because the hyper-parameters enter the model nonlinearly, they can not be estimated by ordinary least squares, and we have to rely on other estimation techniques, such as the maximum likelihood or generalized method and moments estimators.

2.2 Forecasting Methods Based on Aggregates

The various flavors of MIDAS models described above take into account high frequency information contained in the explanatory variables. In contrast, conventional models tend to lack the flexibility to efficiently incorporate this information into forecasts. We can gauge the impact of high frequency information on forecast precision by comparing the two types of models. The simplest way to generate a quarterly forecast is to apply a fourth order autoregressive, or AR(4), model to quarterly data. Some variant of this model usually serves as a benchmark in the ranking of various forecasting methods. However, a quarterly AR(4) model does not take into account either

explanatory variables or monthly information available within a quarter.

A partial solution to the limitations of a quarterly AR(4) model is afforded by a *bridge*, consisting of two steps (see also Schumacher, 2014). In the first step, an autoregressive model is used to iteratively forecast the values of the monthly regressors for the remainder of the current quarter, and the monthly values within each quarter are aggregated. In the second step, a quarterly explanatory regression is estimated from historical data, and then evaluated with projected quarterly values of the regressors.

If, in addition to the predictors, the variable of interest is also available at the monthly frequency, then predictions for the current quarter and beyond can be generated by high frequency equivalents of the AR and MIDAS models. In particular, single frequency AR and ADL models can be applied to monthly data, and subsequently the monthly forecasts can be aggregated to the quarterly frequency. In our empirical illustration, we will compare the forecasting performance of all methods described above that are feasible.

3 Empirical Examples

Our goal is to demonstrate the impact of high frequency information on the accuracy of nowcasts and forecasts. We accomplish this by comparing the mixed and single frequency models described in Section 2 in two separate forecasting exercises. We also address two empirical issues associated with data availability at time t : the ragged edge problem (unbalanced data set)

and regressions with real time data (vintages).

3.1 Data

Our first application illustrates how to obtain forecasts of quarterly tourist arrivals to the state of Hawaii using monthly tourist arrivals, monthly passenger counts, and monthly airline passenger seats outlook. Although historical monthly values of tourist arrivals are available, a quarterly prediction corresponds to a multi-period, or three-month, forecast. The quarterly predictions are used, among others, to evaluate the industry's expected performance against a quarterly target.

Our second application illustrates how to obtain forecasts of quarterly earnings for the accommodation and food services industry for the state of Hawaii using the monthly consumer price index, monthly accommodation and food services jobs, and monthly tourist days. Quarterly earnings for the accommodation and food services industry, like tourist arrivals, is a useful indicator on its own, but also an important component of quarterly macroeconomic models for Hawaii given that the accommodations and food services industry accounts for more than 8% of Hawaii's state GDP.

3.1.1 Application 1 - Prediction of Quarterly Tourist Arrivals

Tourists are defined as persons on arriving airline flights excluding in-transit travelers and returning residents. The Hawaii Tourism Authority (HTA) estimates the number of in-transit travelers and residents by surveying pas-

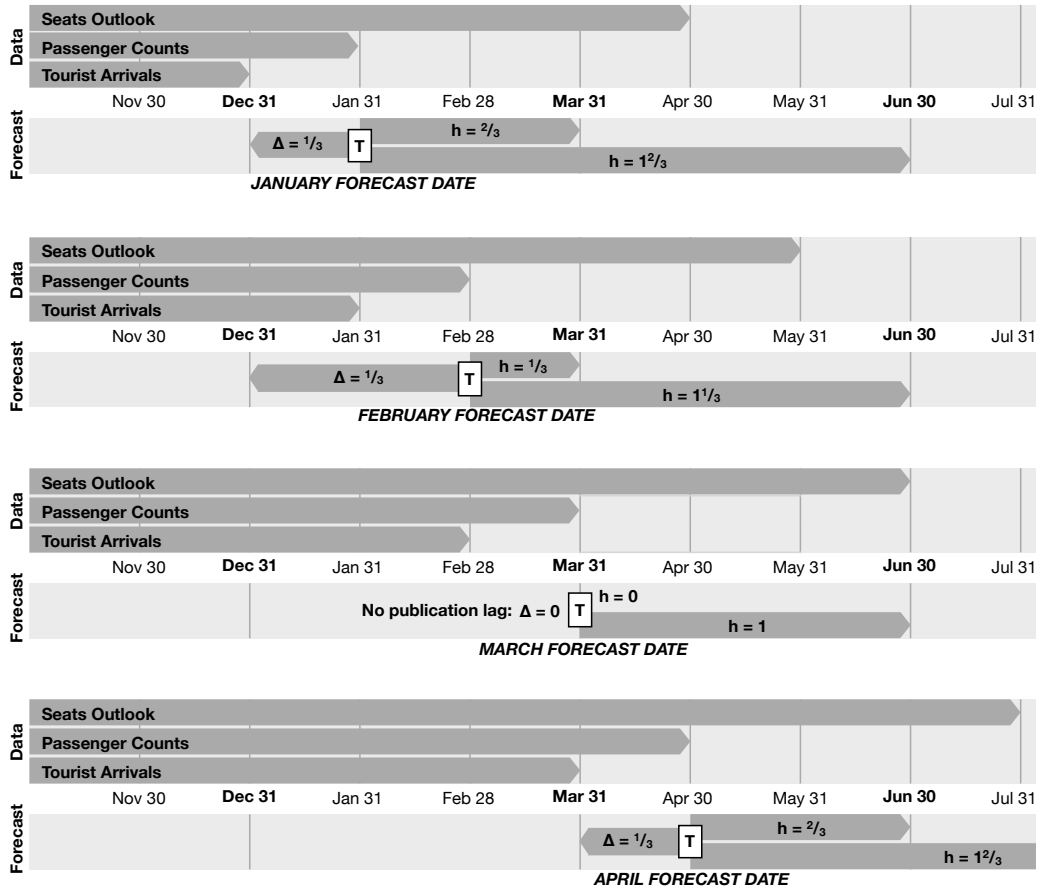


Figure 1: Increase of the information set and change of forecast horizon as the forecast date T progresses through a quarter in Application 1. Δ denotes the “release lag” of quarterly tourist arrivals.

sengers on domestic flights and analyzing US Customs Declarations Forms from international flights. HTA then calculates tourist arrivals by subtracting non-tourists from the total passenger counts reported by airlines. Monthly tourist arrivals estimates are released with a one month lag. Quarterly tourist arrivals are the sum of the monthly values within a quarter.

We obtained airline passenger counts from the Hawaii Department of

Business, Economic Development, and Tourism (DBEDT). The monthly value of this indicator, available with a two-day lag, captures the total number of airline passengers within a month. It includes passengers that arrive on both international and domestic flights with the exception of flights originating in Canada. Since this indicator is available almost contemporaneously, we include it in our model to inform us about current changes in traveler volumes.

The airline seats outlook captures the total number of scheduled seats expected to be flown on future direct flights to Hawaii excluding charter flights. This indicator is prepared by HTA based on data from Diio Mi flight schedules and is available for three months ahead. Due to its forward looking nature, the seats outlook is subject to greater uncertainty than historical data. It tends to undergo significant revisions from one release to the next, especially during rapid changes in airlift. For example, the January vintage of the March seats outlook may be quite different from the February one.

Figure 1 illustrates the increasing amount of information available as the forecast date T progresses through a quarter. At the end of January, we have tourist arrivals for December and consequently for the fourth quarter of the previous year, passenger counts for January, and seats outlook through April. The forecast horizons for tourist arrivals in the first and second quarters are $h = \frac{2}{3}$ and $h = 1\frac{2}{3}$, respectively. To illustrate the construction of the non-overlapping model, first consider the nowcast at horizon $h = \frac{2}{3}$; it is based on four lags of quarterly tourist arrivals in the previous year, passenger counts

for January, and seats outlook for February and March. For the $h = 1\frac{2}{3}$ forecast horizon, the non-overlapping model also uses seats outlook for April in addition to the information used for the nowcast.

At the end of February, we have tourist arrivals for January, passenger counts for February, and seats outlook through May. The forecast horizons for tourist arrivals in the first and second quarters are $h = \frac{1}{3}$ and $h = 1\frac{1}{3}$, respectively. For the nowcast at horizon $h = \frac{1}{3}$, the non-overlapping model uses four lags of quarterly tourist arrivals in the previous year, tourist arrivals for January, passenger counts for February, and seats outlook for March. For the $h = 1\frac{1}{3}$ forecast horizon, the non-overlapping model also uses seats outlook for April and May in addition to the information used for the nowcast.

At the end of March, we have tourist arrivals for February, passenger counts for March, and seats outlook through June. The forecast horizons for tourist arrivals in the first and second quarters are $h = 0$ and $h = 1$, respectively. For the nowcast at horizon $h = 0$, the non-overlapping model uses four lags of quarterly tourist arrivals in the previous year, tourist arrivals for January, passenger counts for February and March. For the $h = 1$ forecast horizon, the non-overlapping model also uses seats outlook for April, May, and June in addition to the information used for the nowcast. By the end of April the information set has shifted forward by a full quarter relative to January, and our focus turns to predictions for the second and third quarter, or horizons $h = \frac{2}{3}$ and $h = 1\frac{2}{3}$, respectively. The analysis therefore covers

forecast horizons between $h = 0$ and $h = 1\frac{2}{3}$, in $\frac{1}{3}$, or monthly, increments.

We construct a real time data set that contains each vintage of data. This means that for all variables we collect unrevised historical values and subsequent revisions. The goal of constructing a real time data set is to replicate the actual data that would have been available to produce a forecast at a given time. This is especially important because of the frequent and sizable revisions that the seats outlook series undergoes. To avoid issues related to unit roots and seasonality, we convert levels to year-over-year growth rates.

Our sample starts in January of 2001, and we produce quasi out-of-sample forecasts between January of 2008 and June of 2014. We estimate the model parameters from recursive samples where the starting period is held fixed and the ending period advances with the forecast date. The Autometrics based model, determined by diagnostic criteria, is respecified in each iteration of the forecasting exercise. We set the maximum lag length to 4 for quarterly tourist arrivals ($y_{max} = 4$) and to 12 for monthly tourist arrivals, passenger counts, and airline seats ($x_{max} = 12$). The MIDAS model also uses these lag limits.

3.1.2 Application 2 - Prediction of Quarterly Income

Industry earnings are defined as the labor income paid out to employees and proprietors' of a particular industry. The US Bureau of Economic Analysis (BEA) produces estimates of industry earnings based on a number of administrative data sources as well as surveys and census data. We focus on

labor income in the accommodation and food services industry, which—in contrast to tourist arrivals—is not available at the monthly frequency. Estimates are released quarterly with roughly a one quarter lag: earnings for the first quarter are released in June, earnings for the second quarter are released in September, and so on.




Figure 2: Increase of the information set and change of forecast horizon as the forecast date T progresses through a quarter in Application 2. Δ denotes the release lag of quarterly labor income.

We use several predictors of labor income. Figure 2 illustrates the increasing amount of information available as the forecast date T progresses through a quarter. Payroll jobs for the accommodation and food services industry in Hawaii are estimated jointly by the BLS and the Hawaii Department of Labor and Industrial Relations (DLIR) as part of the Current Employment Statistics program. Since the vast majority of industry earnings consist of payments to employees, payroll jobs should provide information on changes in earnings associated with changes in the total number of jobs. Payroll jobs at the state level are available with a half-month publication lag.

The headline Consumer Price Index for All Urban Consumers, CPI-U, is a US city average for all items from the US Bureau of Labor Statistics (BLS). Industry earnings are only released in nominal dollars so it follows that the CPI could have considerable predictive power at least for changes in earnings associated with changes in the overall price level. While there is a consumer price index for Honolulu, HI, this index is only available semi-annually and with a lengthy publication lag, limiting its usefulness for producing quarterly nowcasts. The national CPI, in contrast, is available monthly, and similarly to payroll jobs with a short, roughly two-week, publication lag.

The accommodation and food services industry in Hawaii is heavily influenced by tourism activity, which can be captured by tourist days. Tourist days are defined as the total number of days spent in the state by tourists who arrive by air. Tourist days are estimated by HTA from the same surveys and administrative sources used to estimate tourist arrivals and are released

with the same one-month lag.

Since the publication lag on industry earnings is almost a full quarter, industry earnings for the previous quarter are not available during the first two months of each quarter. For example in January and February, the last observation available for earnings is the third quarter of the previous year. Therefore, in addition to nowcasts for the current quarter and forecasts for the subsequent quarter, in this application we also produce backcasts for the previous quarter.

Although in this application we have a larger sample that begins in January 1990, the format of the forecasting exercise largely follows the first application. We produce predictions during the identical period ranging from January 2008 to June 2014. The data are transformed to year-over-year differences of log-levels, the maximum lag-length is set to four quarters and twelve months, and we use recursive estimation, the same as in the first application.

3.2 Results

We evaluate the forecasting performance of all methods by comparing their Root Mean Squared Forecast Errors (RMSFE). We expect the forecast accuracy to improve as the amount of useful information contained in the predictors increases.

Table 1: Comparison of Forecasting Performance for Application 1

Model	Forecast			Nowcast		
	$h = 1\frac{2}{3}$	$h = 1\frac{1}{3}$	$h = 1$	$h = \frac{2}{3}$	$h = \frac{1}{3}$	$h = 0$
Quarterly AR	161.3	161.3	161.3	100.7	100.7	100.7
Monthly AR	153.8	139.5	129.8	80.3	41.5	26.7
Monthly ADL	73.8	78.5	73.1	35.8	22.1	10.4
Bridge	115.6	111.7	111.5	43.1	30.4	27.1
Autometrics	107.1	100.4	82.0	41.3	26.4	13.2
Non-Overlapping	96.1	84.1	75.5	57.3	27.7	13.0
MIDAS	94.2	88.4	77.9	46.0	27.5	18.4

Note: Root mean squared forecast error for each model and forecast horizon. Numbers in bold font represent the lowest RMSFE for a particular forecast horizon, h .

3.2.1 Application 1 - Prediction of Quarterly Tourist Arrivals

Table 1 and Figure 3 report the results for our first application. For all models, RMSFE tends to decrease as the forecast horizon shortens. For most of the models, the largest reduction in RMSFE occurs when the forecast horizon shrinks from $h = 1$ to $h = \frac{2}{3}$ as tourist arrivals for the full previous quarter become available. In fact, the quarterly AR model benefits from new information only at this horizon. In contrast, the mixed frequency models take advantage of monthly data, and therefore their precision tends to continually improve as the forecast horizon shrinks.

Regardless of the method used to incorporate the high frequency information, all models show clear improvement in predictive accuracy relative to the quarterly AR model. Furthermore, a small-sample corrected Diebold-Mariano test (Diebold and Mariano, 1995; Harvey et al., 1997) indicates that

Application 1: Root Mean Squared Forecast Error by Forecast Horizon

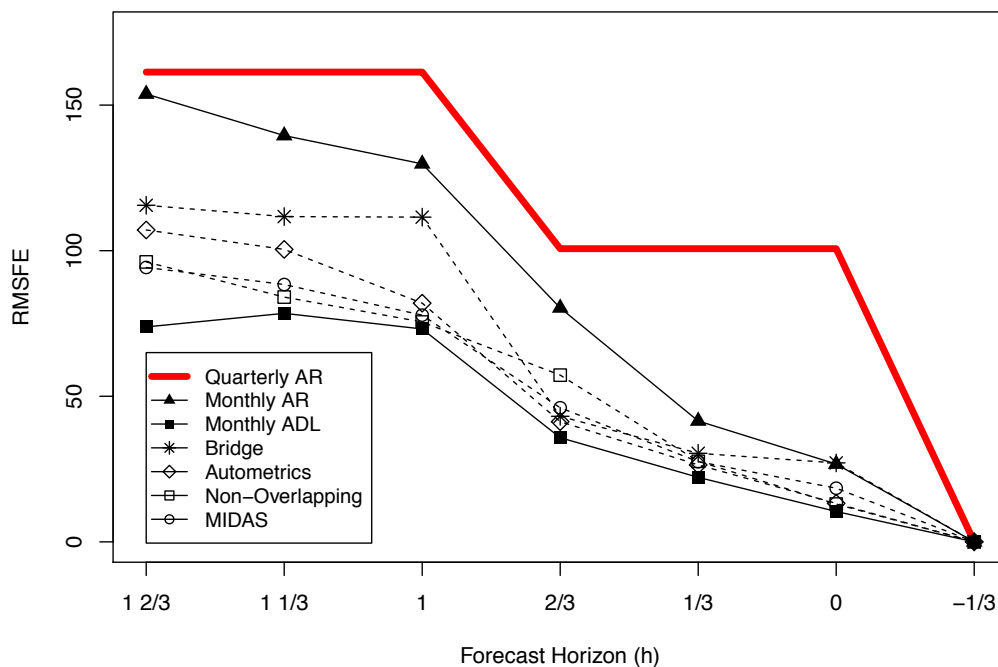


Figure 3: RMSFE for each model for forecast horizons between $h = 1\frac{2}{3}$ and $h = 0$.

the differences between the quarterly AR model and all other models except the monthly AR model are statistically significant. This demonstrates the contribution of the high frequency information.

The monthly AR model is dominated by the multivariate models in almost all cases. The only exception is $h = 0$, where the monthly AR model slightly outperforms the Bridge model, although this difference is not statistically significant. This result illustrates the value of multivariate methods. The monthly AR model incorporates high frequency information about the dependent variable and produces more accurate predictions relative to

the quarterly AR model. But the multivariate methods go a step further by also using high frequency information contained in a set of explanatory variables, and consequently yield more accurate predictions than either univariate model.

Among the multivariate models, the distributed lag model clearly performs the best. At every forecast horizon, h , predictions from the distributed lag model have the lowest RMSFE. This may be due to the fact that quarterly predictions from distributed lag models are aggregates of monthly values. The aggregation either uses an available actual monthly observation, or a prediction for a particular month. In contrast, the mixed frequency models use monthly variables as predictors of quarterly tourist arrivals in a regression and no aggregation takes place. Although the use of actuals and monthly predictions results in a slight advantage of the distributed lag model, it is important to remember that this approach is only feasible if the variable of interest is also available at the monthly frequency.

Across the rest of the multivariate models—essentially the mixed frequency ones—there is no clear ranking. For example the MIDAS model outperforms the Autometrics based model in all three forecasting periods, whereas the Autometrics model outperforms MIDAS in all three nowcasting periods. However, none of these differences are statistically significant. And, for each horizon, even the worst performing multivariate model still provides a meaningful increase in predictive accuracy relative to the quarterly AR model. The choice of model is relatively unimportant; incorporating

Table 2: Comparison of Forecasting Performance for Application 2

Model	Forecast			Nowcast			Backcast	
	$h = 1\frac{2}{3}$	$h = 1\frac{1}{3}$	$h = 1$	$h = \frac{2}{3}$	$h = \frac{1}{3}$	$h = 0$	$h = -\frac{1}{3}$	$h = -\frac{2}{3}$
Quarterly AR	233.4	233.4	169.5	169.5	169.5	114.4	114.4	114.4
Quarterly ADL	179.4	179.4	170.3	121.7	121.7	114.9	67.8	67.8
Bridge	198.6	184.7	129.7	96.8	82.7	74.8	67.8	67.8
Autometrics	173.3	168.4	140.2	112.8	106.7	74.9	62.2	62.2
MIDAS	164.5	136.7	129.6	110.3	91.3	73.3	64.4	64.4

Note: Root mean squared forecast error for each model and forecast horizon. Numbers in bold font represent the lowest RMSFE for a particular forecast horizon, h .

high frequency information into the forecasting process results in a statistically significant improvement in forecasting accuracy but moving from one method to another provides only modest gains.

3.2.2 Application 2 - Prediction of Quarterly Income

Table 2 and Figure 4 report the results for our second application. The results for this application are largely similar to those for the first application. Again, RMSFE declines as the forecast horizon, h , shrinks—increasing the amount of useful information in the model results in a more accurate prediction. The quarterly AR model ranks worst in terms of forecasting performance; all of the models incorporating explanatory variables, at either the monthly or quarterly frequency, outperform the quarterly AR model, except at $h = 1$ and $h = 0$ where the quarterly ADL model performance is very similarly.

In general, incorporating high frequency information through any of the proposed methods results in a reduction in RMSFE relative to the quarterly

Application 2: Root Mean Squared Forecast Error by Forecast Horizon

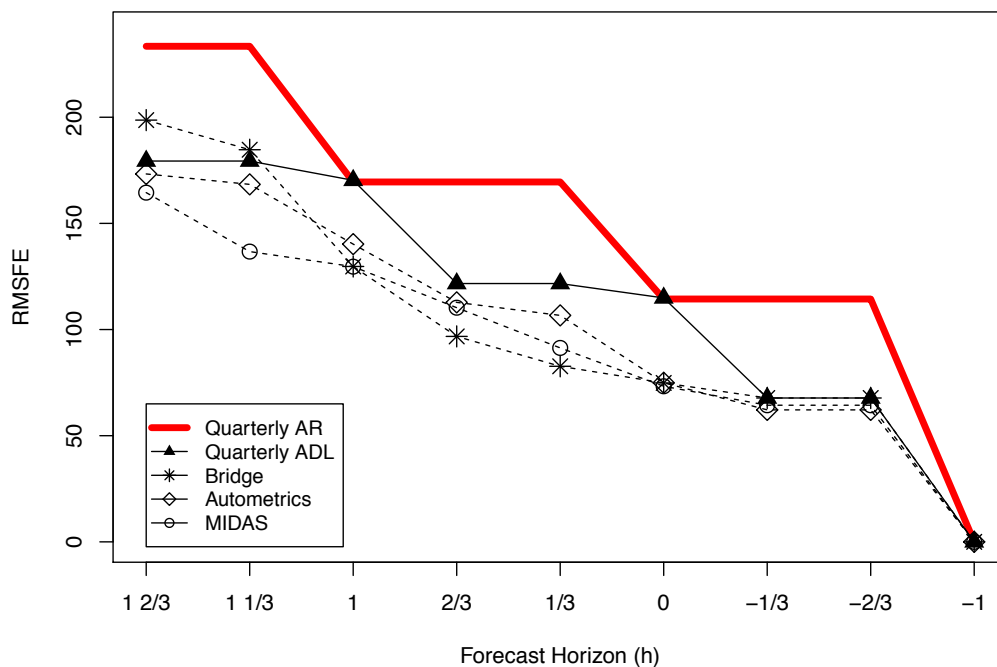


Figure 4: RMSFE for each model for forecast horizons between $h = 1\frac{2}{3}$ and $h = -\frac{2}{3}$.

models. However in the backcast periods, $h = -\frac{1}{3}$ and $-\frac{2}{3}$, the high frequency methods no longer provide any statistically significant improvement in precision over the quarterly ADL model. At $h = -\frac{1}{3}$ and $h = -\frac{2}{3}$, observations for all three months of the quarter are available for the explanatory variables, and the quarterly ADL model incorporates that information at the quarterly frequency. This highlights the value of examining the availability and relevance of high frequency predictors, when considering the use of mixed frequency forecasting methods. If the high frequency predictors contain useful information, then the greater accuracy of mixed frequency methods may

outweigh the complexity they introduce to the forecasting process; otherwise, working at the low frequency may be preferable.

In this application, there is no model that statistically dominates all others. At most forecast, backcast, and nowcast horizons, the MIDAS model performs relatively well. But the only horizon where the MIDAS forecasts statistically dominate the Bridge forecasts is at the first two forecast horizons, $h = 1\frac{2}{3}$ and $h = 1\frac{1}{3}$, suggesting that the latter may still be a worthy alternative for nowcasting and backcasting. The greater nowcasting precision of mixed frequency models relative to the quarterly ADL model illustrates the importance of intra-period information at the end of the sample. However, the similarity of mixed frequency model performances implies that the choice between the Autometrics and MIDAS models is less consequential.

4 Conclusion

We contribute to the existing literature in several important ways. First, we examine a number of econometric tools that can be employed for short term prediction of tourist arrivals. Specifically, we introduce to the tourism literature techniques that directly incorporate high frequency information into forecasting models. Second, for each tool, we highlight distinguishing features and limitations that practitioners need to be aware of. Similarly to their growing popularity in empirical macroeconomics, these techniques will also become standard forecasting tools in tourism research. To facilitate their

adoption, we provide practical guidelines for their implementation. Third, we illustrate the merits of these tools in a real life setting by evaluating their performance in forecasting tourist arrivals and labor income in a tourism related industry in Hawaii.

Our study confirms the hypothesis that using high frequency data contributes to an improvement in forecasting performance. The main benefit of high frequency data is that it contains more timely information than low frequency data released with a long publication lag. However, among the models that incorporate high frequency information, the differences tend to be small and often statistically insignificant. This implies that, while practitioners should take advantage of high frequency data, the particular method used to do so is relatively unimportant. Incorporating high frequency data into the forecasting process through any of the methods outlined is likely to result in a substantial improvement in accuracy, whereas moving from one method to another leads to marginal gains at best. Therefore, the optimal model incorporating high frequency information may be the one that is easiest to implement.

References

- Athanasopoulos, G., Hyndman, R. J., Song, H., and Wu, D. C. (2011). The tourism forecasting competition. *International Journal of Forecasting*, 27(3):822 – 844.
- Bai, J., Ghysels, E., and Wright, J. (2013). State space models and MIDAS regressions. *Econometric Reviews*, 32(7):779–813.
- Bangwayo-Skeete, P. and Skeete, R. (2015). Can google data improve the forecasting performance of tourist arrivals? Mixed-data sampling approach. *Tourism Management*, 46:454–464.
- Baumeister, C., Guérin, P., and Kilian, L. (2015). Do high-frequency financial data help forecast oil prices? The MIDAS touch at work. *International Journal of Forecasting*, 31(2):238–252.
- Camacho, M., Perez-Quiros, G., and Poncela, P. (2013). Short-term forecasting for empirical economists: A survey of the recently proposed algorithms. *Foundations and Trends (R) in Econometrics*, 6(2):101–161.
- Diebold, F. and Mariano, R. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3):253–265.
- Dwyer, L., Gill, A., and Seetaram, N. (2012). *Handbook of research methods in tourism: quantitative and qualitative approaches*. Edward Elgar Publishing.

- Foroni, C., Marcellino, M., and Schumacher, C. (2015). U-MIDAS: MIDAS regressions with unrestricted lag polynomials. *Journal of the Royal Statistical Society, Series A*, 178:57–82.
- Fuleky, P. (2012). On the choice of the unit period in time series models. *Applied Economics Letters*, 19(12):1179–1182.
- Fuleky, P. and Bonham, C. S. (2015). Forecasting with mixed frequency factor models in the presence of common trends. *Macroeconomic Dynamics*, 19(4):753–775.
- Ghysels, E., Santa-Clara, P., and Valkanov, R. (2004). The MIDAS touch: Mixed Data Sampling regression models. *Working paper*.
- Ghysels, E., Sinko, A., and Valkanov, R. (2007). MIDAS regressions: Further results and new directions. *Econometric Reviews*, 26(1):53–90.
- Harvey, D., Leybourne, S., and Newbold, P. (1997). Testing the equality of prediction mean squared errors. *International Journal of Forecasting*, 13(2):281–291.
- Hendry, D. F. and Krolzig, H.-M. (2004). Automatic model selection: a new instrument for social science. *Electoral Studies*, 23(3):525–544.
- Jansen, W. J., Jin, X., and de Winter, J. (2012). Forecasting and nowcasting real GDP: Comparing statistical models and subjective forecasts. *DNB Working Papers, Netherlands Central Bank, Research Department*.

Schumacher, C. (2014). Midas and bridge equations. *Discussion Paper, Deutsche Bundesbank*, (26/2014).