Forecasting in a Mixed Up World: Nowcasting Hawaii Tourism

BY

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

Working Paper No. 2015-3R

September 28, 2015

revised: August 2, 2016

Forecasting in a Mixed Up World: Nowcasting Hawaii Tourism

August 2, 2016

Abstract

We evaluate the short term forecasting performance of methods that systematically incorporate high frequency information via covariates. Our study provides a thorough introduction of these methods. We highlight the distinguishing features and limitations of each tool and evaluate their forecasting performance in two tourism-specific applications. The first uses monthly indicators to predict quarterly tourist arrivals to Hawaii; the second predicts quarterly labor income in the accommodations and food services sector. Our results indicate that compared to the exclusive use of low frequency aggregates, including timely intra-period data in the forecasting process results in significant gains in predictive accuracy. Anticipating growing popularity of these techniques among empirical analysts, we present practical implementation guidelines to facilitate their adoption.

Keywords: Nowcast; Ragged edge; Mixed frequency models.
JEL classifications: C22, C82, L83, Z32

1 Introduction

The importance of tourism across the world has led to an extensive literature which seeks to predict visitor volumes. In places where tourism is a key component of the local economy, a variety of organizations rely on predictions of tourism activity to plan their operations. Because many tourism services are perishable, firms have an incentive to use forecasts of tourism demand to efficiently manage supply and prices. Tourism agencies use forecasts to manage promotional strategies and set performance targets; other governmental organizations incorporate tourism forecasts into larger macroeconomic models. A wide variety of indicators can be used to predict tourism activity, but each indicator may be sampled at a different frequency and may only be available after a publication lag. In practice this issue is often solved by transforming the data to a single frequency. But the aggregation process eliminates valuable intra-period information that could be used to update the forecast. 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.

There exists a wide range of methods used to forecast tourism demand, but the method that offers the best forecasting performance varies by application. In two meta-studies, Li et al. (2005) and Song and Li (2008) review 22 and 55 published articles, respectively, that compare alternative forecasting methods in tourism applications. Both studies conclude that no single method dominates all others. The most popular methods used to forecast tourism demand can be classified into two groups: univariate time series methods and multivariate regressions. Univariate methods, like Exponential Smoothing (ES) and Autoregressive Integrated Moving Average (ARIMA) models, use only the history of the variable of interest for prediction. In contrast econometric methods, like Autoregressive Distributed Lag (ADL) and Error Correction (EC) models, incorporate information from a set of explanatory variables. We provide a brief outline of several of these methods and their use in the tourism literature in Section 2. Although several non-regression approaches, many using Artificial Intelligence models, have appeared in the tourism forecasting literature in recent years, these methods are beyond the scope of this study (for more information on these methods see Kon and Turner (2005)).

A common feature of both univariate and multivariate models is that they operate at a single frequency, with the variable of interest and any explanatory variables aggregated to the same frequency before estimation. Consider predicting quarterly tourist arrivals using the monthly number of inbound airline passengers as an explanatory variable. If passenger counts are highly correlated with tourist arrivals, then the monthly passenger counts may contain substantial predictive power. For example, if the forecaster observes that passenger counts were much higher than expected after the first month of the quarter, she may want to incorporate this information to update the tourist arrivals forecast for the quarter. However with a single frequency model the passenger count series must be aggregated to the quarterly frequency, and the intra-quarter information cannot be used for prediction. Some forecasters make ad-hoc adjustments to their short term forecasts to incorporate informa-

tion from incomplete periods at the end of sample. But this process can be cumbersome and, by definition, subjective. Using a model-based approach to incorporate high frequency information streamlines the forecasting process and may improve forecasting accuracy.

Several approaches have been recently developed to *directly* use high frequency regressors to predict a low frequency variable of interest. This is a rapidly growing area of research with over 50 studies in the last decade (see for example Camacho et al., 2013). The simple example above could be formulated as an Unrestricted Mixed Data Sampling (U-MIDAS) model developed by Foroni et al. (2015). U-MIDAS uses high frequency regressors to predict a low frequency variable, so that monthly passenger counts can be directly used to predict quarterly tourist arrivals without any aggregation. Consequently, all of the monthly information can be used as soon as it is available. However the U-MIDAS method results in a parameter proliferation problem when there is a large frequency mismatch between the variable of interest and the high frequency regressors. The Mixed Data Sampling (MIDAS) method of Ghysels et al. (2004, 2007) solves this problem by imposing non-linear restrictions on the model, but it sacrifices the simplicity of ordinary least squares parameter estimation. Foroni et al. (2015) 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.

Despite their appeal, few tourism studies have applied mixed frequency forecasting methods. In fact, we are aware of only one study; Bangwayo-Skeete and Skeete (2015) use a MIDAS model to predict monthly tourist arrivals to several Caribbean destinations with weekly Google search data as an explanatory variable. They find that the mixed frequency model outperforms several univariate methods. We use monthly tourist arrivals and other explanatory variables to obtain a forecast of quarterly tourist arrivals and find similar results. In contrast, mixed frequency forecasting methods have been used extensively in macroeconomic applications. 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 find that incorporating high frequency information improves predictive accuracy for the current period (nowcasting), but gains in forecasting one or two periods ahead appear to be muted. Finally, some studies find that the use of high frequency covariates does not significantly improve predictive accuracy. Baumeister et al. (2015) note that forecast precision depends on whether the high frequency data provides a useful signal or simply introduces additional noise.

Factor models based on the Kalman filter have also been used for prediction in mixed-frequency environments (see for example Fuleky and Bonham, 2015). Factor models are beyond the scope of our study because their main purpose is the estimation of an unobserved overall business cycle variable and its fluctuations (Stock and Watson, 1989, 1991). While, Bai et al. (2013) found that factor models and the MIDAS models analyzed in our paper have similar forecasting performances, the implementation of the former is more complex.

In the rest of the paper, we first discuss several classical and modern forecasting methods used in the tourism literature and highlight their practical advantages and disadvantages. We then evaluate their performance in a mixed frequency environment. Specifically, we use monthly regressors to produce nowcasts and one-quarter-ahead forecasts of tourist arrivals,

and, in a separate exercise, we predict quarterly earnings in the accommodation and food services industry.[1] 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 accuracy across mixed frequency models tend to be small. Therefore, while practitioners should use high frequency information, they should use the method 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.

2   Methods

Tourism forecasters often turn to classical methods to predict industry performance (see for example Athanasopoulos et al., 2010; Dwyer et al., 2012). The most popular of these methods include autoregressive integrated moving average (ARIMA), autoregressive distributed lag (ADL), error correction (ECM), and vector autoregression (VAR) models and their variants. Of the 121 tourism demand modeling and forecasting studies identified by Song and Li (2008) over 2001-2006, nearly 60% used univariate time series techniques. These are the least costly to implement in terms of data requirements and computational complexity because they only use the history of the variable of interest for prediction. More sophisticated econometric methods incorporate information from a set of explanatory variables. Multivariate regressions have greater data requirements and can be more technically challenging to implement, but they allow researchers to identify which components drive fluctuations in tourism demand, and the included predictors may improve tourism demand forecasts.

The empirical literature does not provide clear evidence about the superiority of one of these approaches over the other. A number of studies have found that multivariate regressions outperform univariate time-series approaches in tourism forecasting applications. Song et al. (2000) find that an EC model outperforms ARIMA and random-walk models, and Li et al. (2006) combine an Error Correction and Time Varying Parameter (TVP-EC) model and find that it outperforms TVP, EC, and several univariate methods. However, some studies find that univariate methods outperform multivariate models. Kulendran and Witt (2001) and Kulendran and Witt (2003) find that ARIMA models outperform random walk and EC models for one-quarter-ahead forecasts but for longer horizons the ARIMA models are dominated by simple random walk models. Athanasopoulos et al. (2010) find that univariate approaches outperform methods that include explanatory variables; they suggest that these results may be due to model misspecification and the possibility that forecasting the dependent variable directly may be easier than forecasting the explanatory variables. In contrast, we find clear evidence that multivariate methods incorporating timely intra-period data improve forecasting performance relative to a baseline univariate model.

---

[1]All data and R programs used are available on request.

Many of the aforementioned methods are special cases of the general ADL model (see Banerjee et al., 1993). The ADL model takes advantage of contemporaneous and lagged observations of the variable of interest and its predictors. A typical ADL model can be written as

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

where y_t is the variable of interest and x_t is the predictor, or a set of predictors. For example, consider predicting tourist arrivals, y_t, using lagged tourist arrivals, $y_{t-1}, y_{t-2}, \ldots$ and airline passenger counts, $x_t, x_{t-1}, \ldots$. The lag polynomials, $\phi(L) = \sum_{i=0}^{p} \phi_{i+1}L^i$ and $\beta(L) = \sum_{j=0}^{q} \beta_j L^j$, have lengths p and q whose optimal values can be determined using standard model selection criteria. Additional regressors can be included at the cost of further notational complexity. Univariate time series approaches only consider lags of tourist arrivals and neglect the information in the associated passenger count variable, while contemporaneous multiple regressions ignore any dynamics. An ADL model nests both approaches and is therefore a more general and versatile tool for forecasting.

A forecast produced on forecast date T for a horizon h is based on the estimated relationship between y_t and the h^{th}-and-greater lags of x_t in equation (1). Reporting delays can result in missing observations for recent periods in a vintage T dataset, a phenomenon sometimes called a "ragged edge." This problem can be addressed by allowing for a lag between time T and the most recent published observation. We denote this lag by Δ_y for the target series and Δ_x for the predictor variables and obtain a forecast for horizon h

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

with coefficients previously estimated in a regression that maintains equivalent lags of arrivals and passenger counts *relative to the response variable, y_t,*

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

Note, the time subscript denotes the reference period for a particular observation and not the release date. To illustrate, consider the case where quarterly tourist arrivals are released with a one quarter lag, $\Delta_y = 1$, but quarterly airline passenger counts are released immediately at the end of the quarter, $\Delta_x = 0$. So at the end of the second quarter of the year, tourist arrivals will only be available through the first quarter, but the passenger count series will be available through the end of the second quarter. To estimate an equation for a one-period ahead forecast, $h = 1$, equation (2) will estimate the relationship between tourist arrivals, y_t, and lagged arrivals beginning with y_{t-2} and passenger counts beginning with x_{t-1}.

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, consider the previous example of tourist arrivals and airline passenger counts but now with arrivals sampled quarterly and passenger counts sampled monthly.[2]

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 tourist arrivals, the response variable. Consequently, the monthly passenger count observations are indexed with $t = \frac{1}{3}, \frac{2}{3}, 1, 1\frac{1}{3}, \ldots$, and the quarterly tourist arrivals with $t = 1, 2, \ldots$.[3] The relationship between quarterly tourist arrivals, its own lags, and the lags of monthly passenger counts can be estimated using a simple mixed frequency regression

$$y_t = \alpha + \phi(L)y_{t-1-\Delta_y} + \beta(L^{1/3})x_{t-\Delta_x} + \epsilon_t , \quad \text{for} \quad t = 1, 2 \ldots , \tag{4}$$

where $\phi(L) = \sum_{i=0}^{p} \phi_{i+1}L^i$, $\beta(L^{1/3}) = \sum_{j=0}^{q} \beta_j L^{j/3}$, and Δ_y and Δ_x are the release lags for arrivals and passenger counts, respectively. The time increment of unit length indicates that the rows of the data set are a quarter, or three months, apart. Equation (4) illustrates that the tourist arrivals variable can be directly related to its own lags and the lags of monthly passenger counts without any aggregation beforehand. This is an advantage relative to the single frequency ADL, given by equation (1), where monthly passenger counts need to be aggregated to the quarterly frequency before estimation and prediction.

The forecast date T can fall at the end of any month. Hence, tourist arrival forecasts can be produced for horizons $h = \{0, \frac{1}{3}, \frac{2}{3}, 1, 1\frac{1}{3}, \ldots\}$, where the first three specify predictions for the current quarter and are usually called "nowcasts". For example, a forecast of first quarter arrivals made at the end of January has a forecast horizon $h = \frac{2}{3}$. At the end of February the first quarter forecast can be updated, and the updated forecast has a horizon $h = \frac{1}{3}$. A forecast of tourist arrivals at horizon h requires an estimated relationship between tourist arrivals and passenger counts with lags $h + \Delta_x$ and greater in equation (4). Using the coefficients estimated in such a regression and the data available at the forecast date T, we obtain a forecast, $\hat{y}_{T+h}$, by evaluating

$$\hat{y}_{T+h} = \hat{\alpha} + \hat{\phi}(L)y_{T-\Delta_y} + \hat{\beta}(L^{1/3})x_{T-\Delta_x} , \tag{5}$$

where Δ_y and Δ_x denote lags between the forecast period, T, and the most recent quarterly and monthly observation, respectively. For example, if the target variable y is published at the end of the last month of each quarter, and the forecast date is the end of the first month

---

[2]While this example uses high frequency (monthly) predictors to forecast a quarterly variable of interest, the methods described could use higher frequency predictors such as weekly or even daily data. For example, daily data on snowfall might be very useful when nowcasting monthly tourist arrivals or revenue in a region dominated by winter sports tourism.

[3]The 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.

of the quarter, then $\Delta_y = \frac{1}{3}$. The most recent quarterly observation available at the end of the first month of the quarter is the previous quarter's value, $y_{T-\frac{1}{3}}$. Similarly, for a forecast date at the end of the second month, $\Delta_y = \frac{2}{3}$, and a forecast date at the end of the third month of the quarter implies $\Delta_y = 0$ as the current quarter value has just been released.

Foroni et al. (2015) called the model described above an unrestricted mixed data sampling (U-MIDAS) regression. Because the lag-structure of equation (4) 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 OxMetrics 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.[4]

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 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. Refer to Section 3.1.1 for an illustration.

MIDAS of Ghysels et al. (2007) eliminates parameter proliferation by defining $\beta(L^{1/3})$ as an exponential Almon lag polynomial

$$\beta(L^{1/3}, \theta) = \beta_0(\theta)L^{0/3} + \beta_1(\theta)L^{1/3} + \ldots + \beta_q(\theta)L^{q/3}, \tag{6}$$

where

$$\beta_j(\theta) = \frac{e^{\theta_1 j + \theta_2 j^2}}{\sum_{j=0}^{q} e^{\theta_1 j + \theta_2 j^2}} \tag{7}$$

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 nonlinear estimation techniques.

---

[4]The Autometrics model selection algorithm begins from a General Unrestricted Model (GUM) and sequentially discards variables that are statistically insignificant, subject to the restriction that the reduced model continues to pass a number of specification tests such as tests of serially uncorrelated and homoscedastic errors. The goal is to obtain a specific model that is parsimonious, excludes only irrelevant variables and lags, and is an adequate representation of the data.

2.2 Forecasting Methods Based on Aggregates

The various flavors of mixed frequency models described above take into account high frequency information contained in the explanatory variables. In contrast, conventional single frequency models tend to lack the flexibility to efficiently incorporate such information. 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 use an autoregressive model with the quarterly data. Autoregressive models are typically used as benchmarks in the ranking of various forecasting methods (for a list of papers see Song and Li (2008) and Li et al. (2005)). A disadvantage of the quarterly AR model is its neglect of both explanatory variables and monthly information available within a quarter.

A partial solution to the limitations of a quarterly AR 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 explanatory variables for the remainder of the current quarter and then the monthly forecasts are appended to the available history and aggregated to the quarterly frequency. In the second step, a quarterly model similar to the single frequency ADL is estimated from historical data, and then evaluated using the projected quarterly values.

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 using monthly AR and ADL 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 one-quarter-ahead 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 any 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, or three-month forecast is useful for evaluating industry performance against a quarterly target or as inputs into planning, tax, and forecasting models of the macroeconomy.

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

In our application 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 passengers 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: the tourist arrival statistics for January are released at the end of February. 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. Each release includes a three month outlook, so the release at the end of January includes an outlook for February, March and April. The next release at the end of February will include an outlook for March, April, and May, and so on. The number of airline seats actually flown is published with a one month lag together with tourist arrivals. We combine the two seats indicators into a single series by using all available historical values of seats flown and appending the latest seats outlook to the end of the series. 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 outlook for March seats outlook published in the February release may be substantially different from the March seats outlook published in the January release.

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 first and second quarter tourist arrivals are $h = \frac{2}{3}$ and $h = 1\frac{2}{3}$, respectively. To illustrate, the construction of the non-
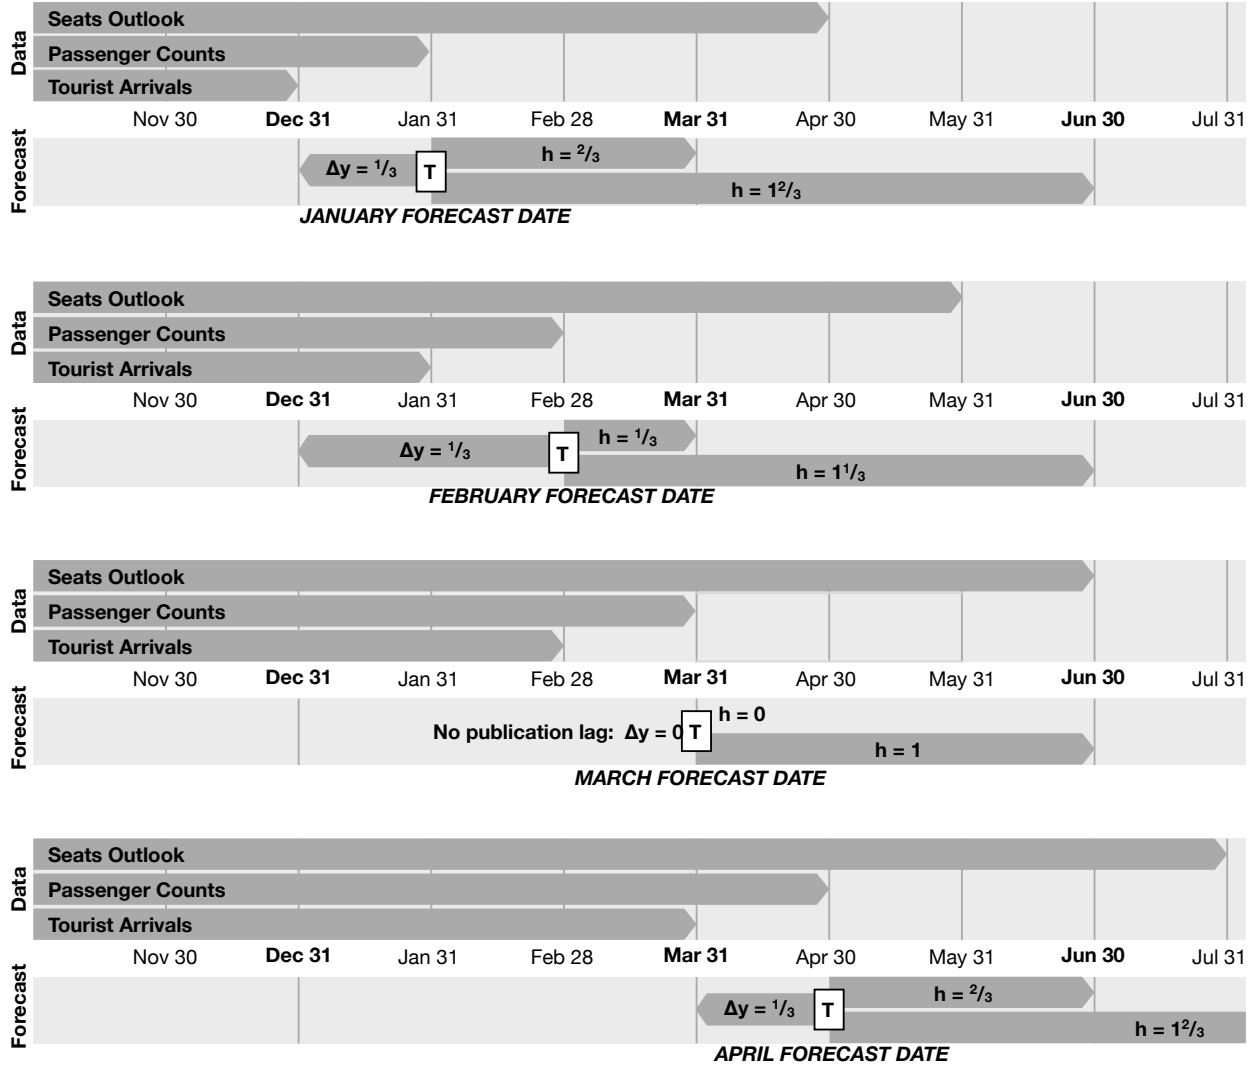
Figure 1: Increase of the information set and change of forecast horizon as the forecast date T progresses through a quarter in Application 1. Δ denotes the "release lag" of quarterly tourist arrivals.

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, in addition to the information used for the nowcast the non-overlapping model uses the seats outlook for April.

At the end of February, we have tourist arrivals for January, passenger counts for February, and seats outlook through May. The forecast horizons for first and second quarter tourist arrivals 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 uses the nowcast information set plus the seats outlook for April and May.

At the end of March, we have tourist arrivals for February, passenger counts for March, and seats outlook through June. The forecast horizons for first and second quarter tourist arrivals 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 and February, and passenger counts for 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 and to 12 for monthly tourist arrivals, passenger counts, and airline seats. 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 of employees and proprietors in 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.

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. 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.

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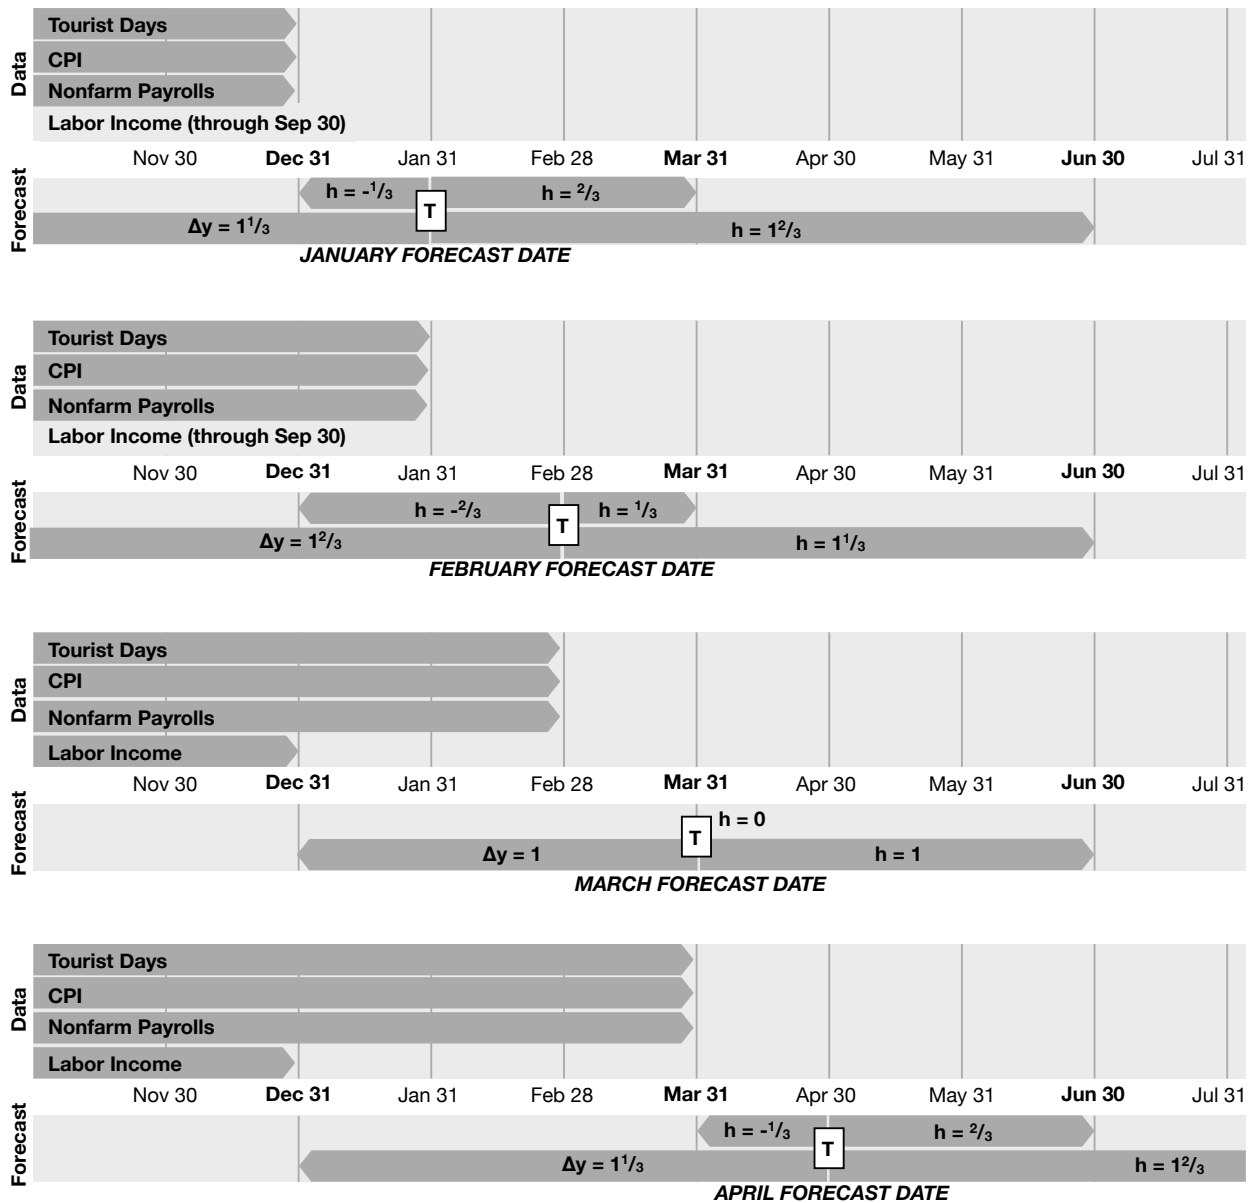
Figure 2: Increase of the information set and change of forecast horizon as the forecast date T progresses through a quarter in Application 2. Δ denotes the release lag of quarterly labor income.

as part of the Current Employment Statistics program. Since the vast majority of industry earnings consist of payments to employees, payroll jobs should provide useful information on changes in earnings due to 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 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.

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.

The format of the forecasting exercise largely follows the first application. They only key difference is that in this application our sample begins in January 1990. We produce predictions for the period 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, as in the first application.

3.2   Results

We evaluate the forecasting performance of all methods by comparing their Root Mean Squared Errors (RMSE) and Mean Absolute Percent Errors (MAPE). Because MAPE is based on a linear loss function, while RMSE is based on a quadratic loss function and will place a greater weight on large forecast errors, our results vary depending on the accuracy measure considered. While comparing forecasting methods using RMSE and MAPE is common in the tourism literature (Li et al., 2005; Song and Li, 2008); it is less common to see tests of the hypothesis that one forecast is more accurate than another. We use a small sample adjustment of the Diebold-Mariano Test (Diebold and Mariano, 1995; Harvey et al., 1997) to formally test whether the observed differences in forecast accuracy are statistically significant. We use two versions of the test, one with a quadratic loss function that corresponds to RMSE and a second with a linear loss function that corresponds to MAPE; the results from both tests are similar. An alternative to the modified DM test has been proposed by Ashley (1998). We choose the modified DM test because it provides a simpler and less computationally costly means of model comparison. We expect forecast accuracy to improve as the horizon shrinks and more contemporaneous intra-period information is used in our mixed frequency models.

3.2.1   Application 1 - Prediction of Quarterly Tourist Arrivals

For our first application, Table 1 and Figure 3 report the results based on RMSE, Table 2 and Figure 4 report the results based on MAPE, and Table 3 reports a summary of the DM test results. For all models, forecast accuracy improves as the forecast horizon shortens. For most of the models, the largest reduction in RMSE or MAPE occurs when the forecast horizon shrinks from $h = 1$ to $h = \frac{2}{3}$ as tourist arrivals for the full previous quarter
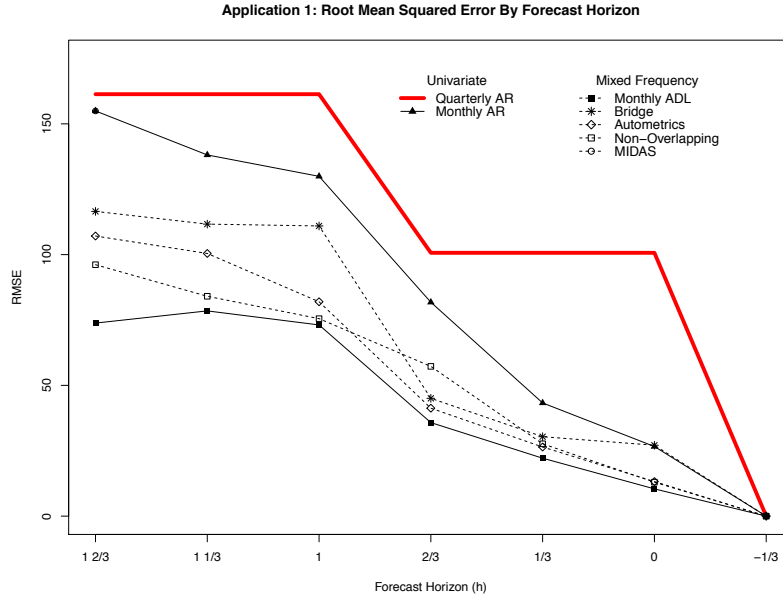
Figure 3: RMSE for each model for forecast horizons between $h = 1\frac{2}{3}$ and $h = 0$.

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 can be updated each month as new data becomes available resulting in a continuous improvement in forecasting accuracy at each horizon.

Figure 4: MAPE for each model for forecast horizons between $h = 1\frac{2}{3}$ and $h = 0$.

14

Table 1: Comparison of Forecasting Performance by RMSE for Application 1

	Forecast			Nowcast		
Model	$h = 1\frac{2}{3}$	$h = 1\frac{1}{3}$	$h = 1$	$h = \frac{2}{3}$	$h = \frac{1}{3}$	$h = 0$
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 error for each model and forecast horizon. Numbers in bold font represent the lowest RMSE for a particular forecast horizon, h.

Table 2: Comparison of Forecasting Performance by MAPE for Application 1

	Forecast			Nowcast		
Model	$h = 1\frac{2}{3}$	$h = 1\frac{1}{3}$	$h = 1$	$h = \frac{2}{3}$	$h = \frac{1}{3}$	$h = 0$
Quarterly AR	4.71	4.71	4.71	2.95	2.95	2.95
Monthly AR	5.13	5.33	5.45	2.05	1.90	1.05
Monthly ADL	**3.20**	**3.21**	**2.36**	1.84	**0.91**	0.56
Bridge	3.77	3.53	3.35	2.14	1.54	1.33
Autometrics	4.27	3.87	2.87	**1.46**	0.94	0.63
Non-Overlapping	4.22	3.58	2.66	2.57	1.07	**0.53**
MIDAS	3.49	3.76	3.00	2.05	1.00	0.66

Note: Mean absolute percent error for each model and forecast horizon. Numbers in bold font represent the lowest MAPE for a particular forecast horizon, h.

All models show clear improvement in predictive accuracy relative to the quarterly AR model regardless of the method used to incorporate the high frequency information. Additionally the methods that incorporate explanatory variables produce more accurate forecasts than both the monthly and quarterly AR models in almost all cases. The only exception is $h = 0$ where the monthly AR model slightly outperforms the Bridge model. These two results illustrate the value of multivariate methods and mixed frequency information. 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 resulting in more accurate predictions than either univariate model.

Model	Quadratic		Linear	
	RMSE	DM	MAPE	DM
Quarterly AR	0	0	3	0
Monthly AR	7	3	3	1
Monthly ADL	**36**	**20**	**28**	**13**
Bridge	13	5	10	1
Autometrics	23	12	17	6
Non-Overlapping	24	11	16	7
MIDAS	22	9	15	4

Note: Values for RMSE and MAPE indicate the total number of models dominated based on RMSE and MAPE across the six forecast horizons. Values for DM indicate the number of models statistically dominated at the 5% significance level based on a quadratic or linear loss function. In both cases a higher value indicates a higher ranking and superior forecasting performance. See the Appendix for full test results.

Among all the methods, the distributed lag model clearly performs the best based on RMSE. At every forecast horizon, h, predictions from the distributed lag model have the lowest RMSE. This may be due to the fact that quarterly predictions from distributed lag models are aggregates of monthly values, which include available monthly observations. In contrast, the mixed frequency models only use monthly variables as predictors of quarterly tourist arrivals in a regression, but no aggregation of actual monthly observations takes place. Although the use of published monthly values 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.

Comparing methods by MAPE, many of the same results hold. The monthly AR model preforms relatively poorly for horizons $h = 1\frac{2}{3}, 1\frac{1}{3}, 1$. All other models continue to offer substantial improvements in predictive accuracy over the quarterly AR model at all forecast horizons. However the monthly ADL model no longer dominates at all forecast horizons. The monthly ADL forecast has the lowest MAPE for four of the six horizons but for $h = \frac{2}{3}$ the Autometrics based model delivers more accurate predictions and for $h = 0$ the Non-Overlapping model forecasts are slightly more accurate.

The DM test results indicate that the mixed frequency methods deliver a statistically significant improvement in forecasting precision over the quarterly AR model, but among the mixed frequency models the differences in forecasting performance are generally not statistically significant. The full DM test results are included in the Appendix and Table 3 summarizes the test results across all six horizons. The monthly ADL model dominates

all other methods based on RMSE point values but less than half of those differences are statistically significant. Among the other mixed frequency methods, differences in accuracy are generally not statistically significant. For a practitioner, these results highlight the importance of incorporating high frequency information, but the method used to incorporate the high frequency information is of second-order importance.

3.2.2 Application 2 - Prediction of Quarterly Income

Table 4 and Figure 5 report the results based on RMSE for our second application, Table 5 and Figure 6 report the results based on MAPE, and Table 6 reports a summary of the DM test results. The results for this application are largely similar to those for the first application. Again, RMSE 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.

The mixed frequency methods tend to offer substantially more accurate forecasts than the quarterly ADL late in the quarter, at horizons $h = 1$ and $h = 0$. At the end of the quarter there are two months of high frequency data used in the mixed frequency model predictions, but these two months of data are unavailable in the quarterly ADL. In contrast, at $h = 1\frac{2}{3}, -\frac{1}{3}, -\frac{2}{3}$ the mixed frequency methods only provide at best a marginal improvement over the quarterly ADL model; at these horizons all three months of high frequency data are available for the previous quarter, so when forming predictions using the quarterly ADL the
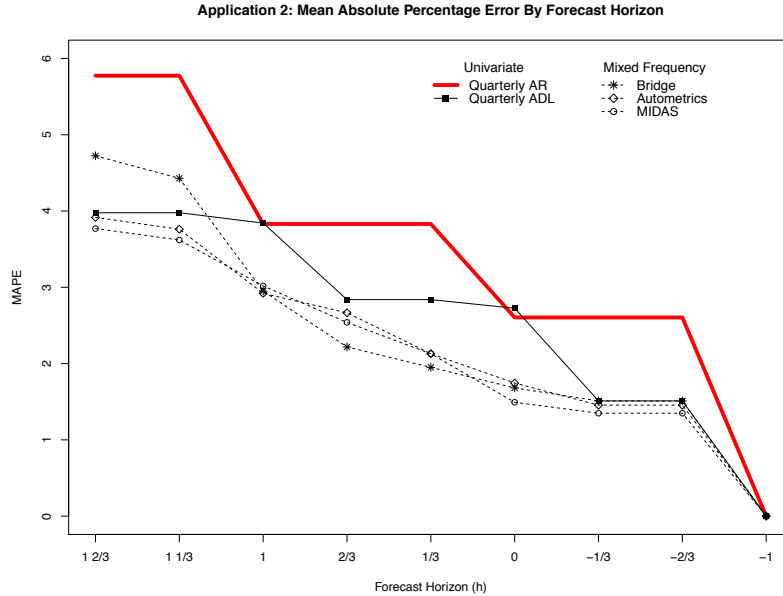

Figure 5: RMSE for each model for forecast horizons between $h = 1\frac{2}{3}$ and $h = -1$.

Figure 6: MAPE for each model for forecast horizons between $h = 1\frac{2}{3}$ and $h = -1$.

latest data on all regressors is available. 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 increased accuracy of mixed frequency methods may outweigh the complexity they introduce to the forecasting process; otherwise, working at the low frequency may be preferable.

Unlike in the first application, there is no clear ranking among the mixed frequency methods as no single model dominates at all horizons. The MIDAS model performs relatively well, dominating at four horizons based on RMSE and five horizons based on MAPE. However the other two mixed frequency methods also perform well; the Bridge model dominates at

Table 4: Comparison of Forecasting Performance by RMSE 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 error for each model and forecast horizon. Numbers in bold font represent the lowest RMSE for a particular forecast horizon, h.

Table 5: Comparison of Forecasting Performance by MAPE for Application 2

	Forecast			Nowcast			Backcast	
Model	$h = 1\frac{2}{3}$	$h = 1\frac{1}{3}$	$h = 1$	$h = \frac{2}{3}$	$h = \frac{1}{3}$	$h = 0$	$h = -\frac{1}{3}$	$h = -\frac{2}{3}$
Quarterly AR	5.77	5.77	3.83	3.83	3.83	2.60	2.60	2.60
Quarterly ADL	3.98	3.98	3.84	2.84	2.84	2.73	1.51	1.51
Bridge	4.72	4.43	2.95	**2.22**	**1.95**	1.68	1.51	1.51
Autometrics	3.92	3.76	**2.92**	2.67	2.13	1.74	1.46	1.46
MIDAS	**3.77**	**3.62**	3.02	2.54	2.13	**1.49**	**1.35**	**1.35**

Note: Mean absolute percentage error for each model and forecast horizon. Numbers in bold font represent the lowest MAPE for a particular forecast horizon, h.

two horizons based on RMSE and MAPE while the Autometrics based model dominates at two horizons based on RMSE and one horizon based on MAPE.

A summary of our statistical comparison of forecast accuracy is included in Table 6 with the full results included in the Appendix. The DM tests indicate that differences in forecasting accuracy between the Quarterly AR model and the other four methods are almost always statistically significant. However differences between the Quarterly ADL model and the mixed frequency methods are only statistically significant in a handful of cases. Differences between the mixed frequency methods are generally not statistically significant.

These results highlight the importance of multivariate methods when explanatory variables contain strong signals about the path of the target series. This is particularly true when intra-period information becomes available in mixed-frequency models at $h = \frac{1}{3}$ and $h = 0$. However when there is no intra-period information, for example at $h = 1\frac{2}{3}, -\frac{1}{3}, -\frac{2}{3},$

Table 6: Statistical Comparison of Forecasting Performance for Application 2

	Quadratic		Linear	
Model	RMSE	DM	MAPE	DM
Quarterly AR	0	0	1	0
Quarterly ADL	7	5	7	5
Bridge	16	8	17	8
Autometrics	8	8	19	9
MIDAS	**24**	**10**	**23**	**13**

Note: Values for RMSE and MAPE indicate the total number of models dominated based on RMSE and MAPE across the eight forecast horizons. Values for DM indicate the number of models statistically dominated at the 5% significance level based on a quadratic or linear loss function. In both cases a higher value indicates a higher ranking and superior forecasting performance. See the Appendix for full test results.

then mixed frequency methods may offer little benefit over a low frequency ADL and add complexity to the forecasting process. Finally, in contrast to the first application where the monthly ADL model generally dominated, this application indicates no clear winner or ranking among the mixed frequency methods. The most accurate mixed frequency method will likely vary from one application to another so practitioners should test a number of methods for their particular application and data availability. For some applications there may be no single method that dominates; the best mixed frequency method may be the one that can be incorporated into the forecasting process at least cost.

4   Conclusion

We contribute to the existing literature on tourism forecasting in several important ways. First, we examine a number of econometric methods that incorporate high frequency information in the forecasting process. These techniques are at the frontier of academic research and are gaining widespread adoption in empirical macroeconomic analysis. Our study provides a thorough introduction of these methods. We highlight each method's distinguishing features and limitations that practitioners need to be aware of. Second, to facilitate their adoption, we present practical guidelines for their implementation. Third, we illustrate the merits of these techniques by evaluating their performance in forecasting Hawaii tourist arrivals and labor income in the accommodations and food service industry.

Our study confirms the hypothesis that using high frequency data improves 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

Ashley, R. (1998). A new technique for postsample model selection and validation. *Journal of Economic Dynamics and Control*, 22(5):647–665.

Athanasopoulos, G., Hyndman, R. J., Song, H., and Wu, D. C. (2010). 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.

Banerjee, A., Dolado, J. J., Galbraith, J. W., and Hendry, D. F. (1993). *Co-integration, Error Correction and the Econometric Analysis of Non-Stationary Data*. Oxford: Oxford University Press.

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. *CIRANO, Centre Interuniversitaire de Reserche en Analyse des Organizations, 2004s-20*.

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*.

Kon, S. C. and Turner, L. W. (2005). Neural network forecasting of tourism demand. *Tourism Economics*, 11(3):301–328.

Kulendran, N. and Witt, S. F. (2001). Cointegration versus least squares regression. *Annals of Tourism Research*, 28(2):291–311.

Kulendran, N. and Witt, S. F. (2003). Forecasting the demand for international business tourism. *Journal of Travel Research*, 41:265–271.

Li, G., Song, H., and Witt, S. (2005). Recent developments in econometric modeling and forecasting. *Journal of Travel Research*, 44(1).

Li, G., Wong, K. K. F., Song, H., and Witt, S. F. (2006). Tourism demand forecasting: A time varying parameter error correction model. *Journal of Travel Research*, 45:175–185.

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

Song, H. and Li, G. (2008). Tourism demand modelling and forecasting–A review of recent research. *Tourism Management*, 29(2):203–220.

Song, H., Romilly, P., and Liu, X. (2000). An empirical study of outbound tourism demand in the uk. *Applied Economics*, 32:611–624.

Stock, J. and Watson, M. (1989). New indexes of coincident and leading economic indicators. *NBER macroeconomics annual*, pages 351–394.

Stock, J. and Watson, M. (1991). A probability model of the coincident economic indicators. *Leading Economic indicators: new approaches and forecasting records*, 66.

5 Appendix

Table 7: Comparison of Forecasting Performance for Application 1: $h = 1\frac{2}{3}$

Model	QAR	MAR	MADL	BR	Auto	N-O	MIDAS
Quarterly AR	1	1.04	2.19**	1.38	1.51*	1.68*	1.71*
Monthly AR	0.96	1	2.1*	1.33	1.45	1.62	1.65
Monthly ADL	0.46**	0.48*	1	0.63*	0.69**	0.77*	0.78
Bridge	0.72	0.75	1.58*	1	1.09	1.21	1.24
Autometrics	0.66*	0.69	1.45**	0.92	1	1.11	1.14
Non-Overlapping	0.6*	0.62	1.3*	0.82	0.9	1	1.02
MIDAS	0.58*	0.61	1.27	0.81	0.88	0.98	1

Note: Values represent the ratio of RMSE of the row model relative to the column model. A value less than one indicates that the row model has lower RMSE than the column model. Differences in forecast accuracy are tested with the DM test using a quadratic loss function: *Significant at 5% level. **Significant at the 1% level.

Table 8: Comparison of Forecasting Performance for Application 1: $h = 1\frac{1}{3}$

Model	QAR	MAR	MADL	BR	Auto	N-O	MIDAS
Quarterly AR	1	1.17	2.05*	1.44*	1.6*	1.91*	1.82*
Monthly AR	0.86	1	1.76*	1.24	1.37	1.64*	1.56
Monthly ADL	0.49*	0.57*	1	0.7*	0.78	0.93	0.89
Bridge	0.69*	0.81	1.42*	1	1.11	1.33	1.26
Autometrics	0.62*	0.73	1.28	0.9	1	1.2	1.14
Non-Overlapping	0.52*	0.61*	1.07	0.75	0.84	1	0.95
MIDAS	0.55*	0.64	1.13	0.79	0.88	1.05	1

Note: Values represent the ratio of RMSE of the row model relative to the column model. A value less than one indicates that the row model has lower RMSE than the column model. Differences in forecast accuracy are tested with the DM test using a quadratic loss function: *Significant at 5% level. **Significant at the 1% level.

Table 9: Comparison of Forecasting Performance for Application 1: $h = 1$

Model	QAR	MAR	MADL	BR	Auto	N-O	MIDAS
Quarterly AR	1	1.24	2.21**	1.45*	1.96*	2.14**	2.07**
Monthly AR	0.81	1	1.79**	1.17	1.59*	1.73**	1.67**
Monthly ADL	0.45**	0.56**	1	0.66*	0.89	0.97	0.94
Bridge	0.69*	0.85	1.52*	1	1.35	1.47	1.42
Autometrics	0.51*	0.63*	1.13	0.74	1	1.09	1.05
Non-Overlapping	0.47**	0.58**	1.03	0.68	0.92	1	0.97
MIDAS	0.48**	0.6**	1.07	0.7	0.95	1.04	1

Note: Values represent the ratio of RMSE of the row model relative to the column model. A value less than one indicates that the row model has lower RMSE than the column model. Differences in forecast accuracy are tested with the DM test using a quadratic loss function: *Significant at 5% level. **Significant at the 1% level.

Table 10: Comparison of Forecasting Performance for Application 1: $h = \frac{2}{3}$

Model	QAR	MAR	MADL	BR	Auto	N-O	MIDAS
Quarterly AR	1	1.23*	2.81*	2.24*	2.46*	1.76*	2.18*
Monthly AR	0.81*	1	2.29*	1.82	2*	1.44	1.78*
Monthly ADL	0.36*	0.44*	1	0.8	0.87	0.63**	0.78
Bridge	0.45*	0.55	1.26	1	1.1	0.79	0.97
Autometrics	0.41*	0.5*	1.15	0.91	1	0.72*	0.89
Non-Overlapping	0.57*	0.7	1.6**	1.27	1.39*	1	1.24
MIDAS	0.46*	0.56*	1.29	1.03	1.13	0.81	1

Note: Values represent the ratio of RMSE of the row model relative to the column model. A value less than one indicates that the row model has lower RMSE than the column model. Differences in forecast accuracy are tested with the DM test using a quadratic loss function: *Significant at 5% level. **Significant at the 1% level.

Table 11: Comparison of Forecasting Performance for Application 1: $h = \frac{1}{3}$

Model	QAR	MAR	MADL	BR	Auto	N-O	MIDAS
Quarterly AR	1	2.33*	4.56**	3.27*	3.8**	3.64**	3.66**
Monthly AR	0.43*	1	1.96**	1.4	1.63**	1.56*	1.57*
Monthly ADL	0.22**	0.51**	1	0.72	0.83	0.8	0.8
Bridge	0.31*	0.71	1.4	1	1.16	1.12	1.12
Autometrics	0.26**	0.61**	1.2	0.86	1	0.96	0.96
Non-Overlapping	0.27**	0.64*	1.25	0.9	1.04	1	1
MIDAS	0.27**	0.64*	1.25	0.89	1.04	1	1

Note: Values represent the ratio of RMSE of the row model relative to the column model. A value less than one indicates that the row model has lower RMSE than the column model. Differences in forecast accuracy are tested with the DM test using a quadratic loss function: *Significant at 5% level. **Significant at the 1% level.

Table 12: Comparison of Forecasting Performance for Application 1: $h = 0$

Model	QAR	MAR	MADL	BR	Auto	N-O	MIDAS
Quarterly AR	1	3.78**	9.62**	3.62**	7.67**	7.77**	5.45**
Monthly AR	0.26**	1	2.54**	0.96	2.03**	2.05**	1.44
Monthly ADL	0.1**	0.39**	1	0.38**	0.8	0.81	0.57*
Bridge	0.28**	1.04	2.66**	1	2.12**	2.14**	1.51
Autometrics	0.13**	0.49**	1.25	0.47**	1	1.01	0.71
Non-Overlapping	0.13**	0.49**	1.24	0.47**	0.99	1	0.7
MIDAS	0.18**	0.69	1.76*	0.66	1.41	1.42	1

Note: Values represent the ratio of RMSE of the row model relative to the column model. A value less than one indicates that the row model has lower RMSE than the column model. Differences in forecast accuracy are tested with the DM test using a quadratic loss function: *Significant at 5% level. **Significant at the 1% level.

Table 13: Comparison of Forecasting Performance for Application 2: $h = 1\frac{2}{3}$

Model	QAR	QADL	BR	Auto	MIDAS
Quarterly AR	1	1.3**	1.18	1.35*	1.46**
Quarterly ADL	0.77**	1	0.9	1.04	1.12
Bridge	0.85	1.11	1	1.15	1.24**
Autometrics	0.74*	0.97	0.87	1	1.08
MIDAS	0.69**	0.89	0.81**	0.93	1

Note: Values represent the ratio of RMSE of the row model relative to the column model. A value less than one indicates that the row model has lower RMSE than the column model. Differences in forecast accuracy are tested with the DM test using a quadratic loss function: *Significant at 5% level. **Significant at the 1% level.

Table 14: Comparison of Forecasting Performance for Application 2: $h = 1\frac{1}{3}$

Model	QAR	QADL	BR	Auto	MIDAS
Quarterly AR	1	1.3**	1.26*	1.39*	1.57**
Quarterly ADL	0.77**	1	0.97	1.07	1.21
Bridge	0.79*	1.03	1	1.1	1.25*
Autometrics	0.72*	0.94	0.91	1	1.14
MIDAS	0.64**	0.83	0.8*	0.88	1

Note: Values represent the ratio of RMSE of the row model relative to the column model. A value less than one indicates that the row model has lower RMSE than the column model. Differences in forecast accuracy are tested with the DM test using a quadratic loss function: *Significant at 5% level. **Significant at the 1% level.

Table 15: Comparison of Forecasting Performance for Application 2: $h = 1$

Model	QAR	QADL	BR	Auto	MIDAS
Quarterly AR	1	1	1.31*	1.21*	1.32
Quarterly ADL	1	1	1.31	1.21	1.33
Bridge	0.77*	0.76	1	0.92	1.01
Autometrics	0.83*	0.82	1.08	1	1.09
MIDAS	0.76	0.75	0.99	0.91	1

Note: Values represent the ratio of RMSE of the row model relative to the column model. A value less than one indicates that the row model has lower RMSE than the column model. Differences in forecast accuracy are tested with the DM test using a quadratic loss function: *Significant at 5% level. **Significant at the 1% level.

Table 16: Comparison of Forecasting Performance for Application 2: $h = \frac{2}{3}$

Model	QAR	QADL	BR	Auto	MIDAS
Quarterly AR	1	1.39*	1.75*	1.5*	1.49*
Quarterly ADL	0.72*	1	1.26	1.08	1.07
Bridge	0.57*	0.8	1	0.86	0.85
Autometrics	0.67*	0.93	1.17	1	0.99
MIDAS	0.67*	0.94	1.18	1.01	1

Note: Values represent the ratio of RMSE of the row model relative to the column model. A value less than one indicates that the row model has lower RMSE than the column model. Differences in forecast accuracy are tested with the DM test using a quadratic loss function: *Significant at 5% level. **Significant at the 1% level.

Table 17: Comparison of Forecasting Performance for Application 2: $h = \frac{1}{3}$

Model	QAR	QADL	BR	Auto	MIDAS
Quarterly AR	1	1.39*	2.05*	1.59*	1.8**
Quarterly ADL	0.72*	1	1.47*	1.14	1.3**
Bridge	0.49*	0.68*	1	0.78	0.88
Autometrics	0.63*	0.88	1.29	1	1.14
MIDAS	0.55**	0.77**	1.14	0.88	1

Note: Values represent the ratio of RMSE of the row model relative to the column model. A value less than one indicates that the row model has lower RMSE than the column model. Differences in forecast accuracy are tested with the DM test using a quadratic loss function: *Significant at 5% level. **Significant at the 1% level.

Table 18: Comparison of Forecasting Performance for Application 2: $h = 0$

Model	QAR	QADL	BR	Auto	MIDAS
Quarterly AR	1	1	1.53**	1.53**	1.66**
Quarterly ADL	1	1	1.54**	1.53**	1.67**
Bridge	0.65**	0.65**	1	1	1.09
Autometrics	0.66**	0.65**	1	1	1.09
MIDAS	0.6**	0.6**	0.92	0.92	1

Note: Values represent the ratio of RMSE of the row model relative to the column model. A value less than one indicates that the row model has lower RMSE than the column model. Differences in forecast accuracy are tested with the DM test using a quadratic loss function: *Significant at 5% level. **Significant at the 1% level.

Table 19: Comparison of Forecasting Performance for Application 2: $h = -\frac{1}{3}$

Model	QAR	QADL	BR	Auto	MIDAS
Quarterly AR	1	1.69**	1.69**	1.84**	1.78**
Quarterly ADL	0.59**	1	1	1.09	1.05
Bridge	0.59**	1	1	1.09	1.05
Autometrics	0.54**	0.92	0.92	1	0.97
MIDAS	0.56**	0.95	0.95	1.04	1

Note: Values represent the ratio of RMSE of the row model relative to the column model. A value less than one indicates that the row model has lower RMSE than the column model. Differences in forecast accuracy are tested with the DM test using a quadratic loss function: *Significant at 5% level. **Significant at the 1% level.

Table 20: Comparison of Forecasting Performance for Application 2: $h = -\frac{2}{3}$

Model	QAR	QADL	BR	Auto	MIDAS
Quarterly AR	1	1.69**	1.69**	1.84**	1.78**
Quarterly ADL	0.59**	1	1	1.09	1.05
Bridge	0.59**	1	1	1.09	1.05
Autometrics	0.54**	0.92	0.92	1	0.97
MIDAS	0.56**	0.95	0.95	1.04	1

Note: Values represent the ratio of RMSE of the row model relative to the column model. A value less than one indicates that the row model has lower RMSE than the column model. Differences in forecast accuracy are tested with the DM test using a quadratic loss function: *Significant at 5% level. **Significant at the 1% level.