Nowcasting Portuguese tourism exports

Sónia Cabral
Banco de Portugal

Cláudia Duarte
Banco de Portugal

October 2016

Abstract
Given the increasing importance of the continuous monitoring of economic activity, techniques that allow taking advantage of the timely releases of high-frequency data play a key role in short-term forecasting. This article compares two single-equation approaches, namely the traditional bridge models and the more recent Mixed Data Sampling (MIDAS) regressions, to nowcast Portuguese quarterly tourism exports. We consider different specifications of bridge and MIDAS models, as well as combinations of nowcasts, in a recursive pseudo real-time exercise. The evidence is in favour of using short-term indicators for nowcasting tourism exports. MIDAS regressions tend to outperform bridge equations, especially when less current-quarter information is available. The best results are always obtained from a combination of nowcasts from a MIDAS specification with autoregressive dynamics. (JEL: C53, F47, Z39)

Introduction

Travel is the most important sector in Portuguese international trade in services and it has been a major driver of the average surplus of the services account in the last two decades (Figure 1). Even if the importance of exports of other services has progressively risen over time, nominal travel exports still represented more than 45 per cent of total exports of services and more than 15 per cent of total Portuguese exports of goods and services in 2015. In addition, Portuguese exports of travel services have increased strongly in the last years, growing by around 50 per cent from 2010 to 2015. As a result, nominal travel exports represented 6.3 per cent of GDP in 2015 and the surplus of the travel account amounted to more than 4 per cent of GDP in 2015, the highest value of the last two decades.

Comparing with other European Union (EU) countries, the economic importance of the tourism sector for Portugal is also evident (Figure 2). The

Acknowledgements: The authors thank Statistics Portugal (INE - `http://ine.pt`) for kindly providing the data on quarterly tourism exports at constant prices used herein. We also thank António Rua for his helpful comments and suggestions. We appreciate the useful discussions about the data with Carla Ferreira and Ana Mouta. Any errors and omissions are the sole responsibility of the authors. The opinions expressed in the paper are those of the authors and do not necessarily coincide with those of Banco de Portugal or the Eurosystem.

E-mail: scabral@bportugal.pt; cfduarte@bportugal.pt

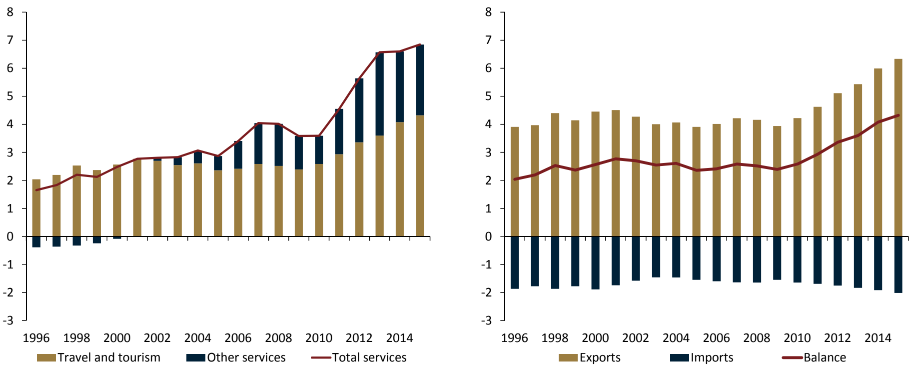

(A) Portuguese services account, balances (B) Portuguese travel account

FIGURE 1: Portuguese exports of tourism

Notes: Figures as a percentage of nominal GDP.
Sources: Statistics Portugal (INE) and Banco de Portugal.

ratio of Portuguese international tourism receipts to GDP increased from 4.8 per cent on average over the years 1995-2000 to 6.3 per cent in the period 2009-2014. This ratio of GDP is more than double the EU average and it is only surpassed by six other EU countries, most of them economies typically associated with significant tourism exports.

As tourism contributes significantly to the growth of the Portuguese economy, accurate forecasts of tourism demand are of particular importance. Calculating timely forecasts typically requires the identification of variables that not only bring useful information, but are also released early. The aim of this article is to use short-term monthly indicators to nowcast the real growth of tourism exports from Portuguese quarterly national accounts. The basic principle is to use information that is published early and at higher frequencies than the variable of interest in order to obtain projections before having observed data.

Considering that we are interested in projecting a quarterly variable on a monthly basis, nowcasting usually refers to the monthly projections of the current quarter and, hence, for each quarter, there are at least 3 different projections, one done in each month of the quarter. In this article, we define "nowcasting" as the projections of a quarter since the first month of that quarter until the official figures are released. Given that Portuguese quarterly national accounts are typically available 60 days after the end of the reference quarter, we produce 5 distinct nowcasts for each quarter. We have data from October 2000 to March 2016 and make use of reduced-form models in a pure time-series approach to nowcast inbound tourism in a recursive pseudo real-time exercise, which mimics the release pattern of the high-frequency

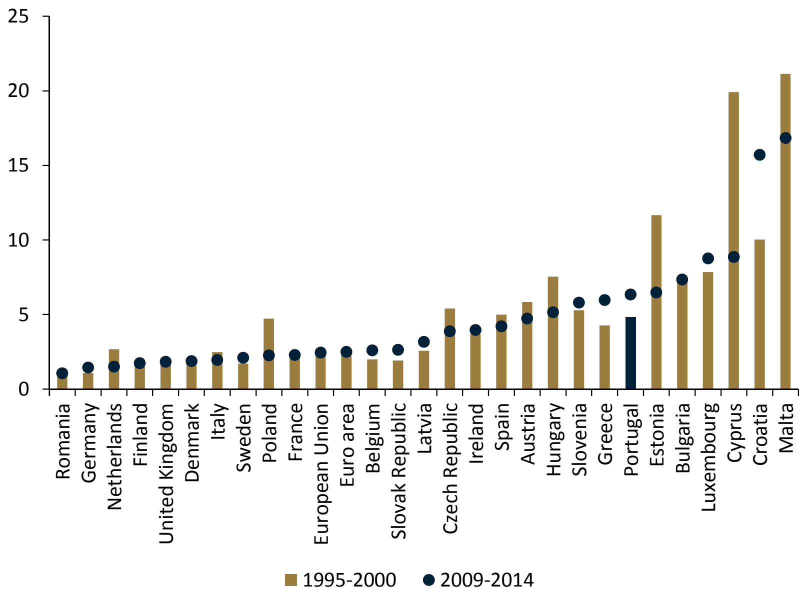

FIGURE 2: International tourism receipts

Notes: International tourism receipts, as percentage of GDP (current U.S. dollars). International tourism receipts are expenditures by international inbound visitors, including payments to national carriers for international transport. These receipts include any other prepayment made for goods or services received in the destination country.
Source: The World Bank - World Development Indicators (WDI).

indicators in real-time situations. The article goes beyond the traditional bridge equations and applies Mixed Data Sampling (MIDAS) regressions as proposed by Ghysels *et al.* (2007). Compared with other mixed-frequency models, the MIDAS approach is appealing because it is a simple, flexible and parsimonious single-regression framework. As far as we know, this is the first application of MIDAS regressions to short-term forecasting quarterly tourism exports.

The results obtained show that, in general, using short-term indicators to nowcast tourism exports is useful, as it delivers more accurate projections than those of a univariate benchmark throughout the whole period. MIDAS models tend to outperform the traditional bridge equations, especially when less current-quarter data on the indicators is available. Pooling nowcasts is a winning strategy for all mixed-frequency models considered: it allows improving on both univariate and single-indicator models. Considering all models and combinations of nowcasts, the best performing result in any period is always from a combination of nowcasts of a MIDAS model with autoregressive dynamics. Overall, despite the relatively short evaluation sample, we provide robust evidence on the improvement in the nowcast accuracy by pooling projections of MIDAS models.

The article is organised as follows. Section 2 discusses some of the related research that frames this study and highlights its main contributions to the literature. Section 3 briefly presents the bridge equations and the MIDAS regressions. Section 4 describes the variables and the design of the empirical exercise of nowcasting Portuguese exports of tourism. Section 5 discusses the results of the exercise conducted. Finally, section 6 presents some final remarks.

Related literature

Tourism demand modelling and forecasting has been an important area of research over the last decades and a number of new methods and techniques have emerged in the literature (see Song and Li (2008) for a survey of studies published after 2000). A special issue of the *International Journal of Forecasting* provides a useful and detailed description of recent developments in tourism forecasting (see Song and Hyndman (2011) for an introduction to this special issue and the papers therein). The review of the vast empirical literature on tourism modelling and forecasting is beyond the scope of this article. Instead, this section offers a non-exhaustive list of references in different strands of the literature that are related to our study and provide a framework for our analysis, with a special focus on the Portuguese economy.

In general, the tourism forecasting literature is still dominated by two main methods: non-causal time-series models and causal econometric approaches. Our study fits in the former broad category. Athanasopoulos *et al.* (2011) conclude that pure time-series approaches forecast tourism demand data more accurately than alternative methods. However, this finding is not unanimous in the literature, as there are numerous conflicting results, especially when using more sophisticated causal models (see, for instance, Song *et al.* (2011)).

Within the time-series models, our work contributes to the general study of tourism activities in Portugal. Given the relevance of the tourism sector, there are some research and policy oriented studies in this area focusing on the Portuguese economy. However, the literature on Portuguese international trade in tourism services is still limited when compared with the large number of studies on Portuguese international trade in goods. Some notable exceptions are Daniel and Ramos (2002) that perform an econometric analysis of the number of tourists arriving from five different origins using cointegration and error correction methods and Teixeira and Fernandes (2012, 2014) that use artificial neural networks models to forecast Portuguese tourism revenues and overnights. Using monthly data on tourist overnight stays in hotel accommodations, Gouveia and Rodrigues (2005) apply a nonparametric method to identify tourism growth cycles, concluding that there is a time lag between tourism demand cycles and economic cycles and Rodrigues and Gouveia (2004) use a parsimonious periodic autoregressive model and demonstrate its superiority in forecasting performance compared to other models. Andraz *et al.* (2009) use a diffusion index model for forecasting tourism demand in Algarve from the UK and confirm its better forecasting performance. More recently, Serra *et al.* (2014) use dynamic panel data models to model international tourism demand in seven different Portuguese tourist regions, finding a heterogeneous behaviour by region.

In our case, we focus on nowcasting developments in quarterly exports of tourism, in real terms. Hence, our article is also related to a large empirical literature on short-term forecasting of economic variables, and in particular, on forecasting Portuguese demand-side components of GDP. In this context, bridge equations are one of the most commonly used techniques to deal with mixed-frequency datasets. Typically, they link monthly and quarterly variables that show a significant correlation and the choice of the regressors tends to take into account their timeliness (see, for instance, Baffigi *et al.* (2004)). Esteves and Rua (2012) provide a general description of the methodology of the short-term forecasting exercise of the Banco de Portugal, where bridge models are the preferred modelling tool. Other applications of bridge models in the short-term forecasting exercises of the Portuguese economy include Cardoso and Duarte (2006) for exports of goods, Maria and Serra (2008) for investment and Esteves (2009) for private consumption.

In addition to the traditional bridge model approach, we consider MIDAS regressions. By using this technique, our article also contributes to a recent stream of empirical literature that uses MIDAS models for handling different sampling frequencies and asynchronous releases of information. Inspired in the distributed lag models, MIDAS regressions are very flexible, being able to account for different frequencies, different aggregation polynomials and different forecast horizons (for a brief overview of the main topics related with MIDAS modelling, see Andreou *et al.* 2011). Recently, Duarte *et al.* (2016) use MIDAS models for nowcasting and forecasting quarterly private consumption in Portugal. As far as we know, the MIDAS approach has not been applied to short-term forecast quarterly tourism exports yet and this article aims at filling that gap.

Bridge equations and MIDAS regressions

Early information on the state of the economy is crucial for policy-making. However, important official statistics, such as those of national accounts, are only available on a quarterly basis and with relevant publication delays. For example, the flash estimate for Portuguese GDP is available 45 days after the end of the quarter, while the main aggregates on the expenditure side are available 60 days after the end of the reference quarter. In this context, techniques for dealing with mixed-frequency data are useful tools to take advantage of the large number of relevant short-term indicators, allowing a timely evaluation of the current economic situation. Several types of econometric tools to combine data with different frequencies and exploit early releases of high-frequency data for improving forecast accuracy have been proposed in the literature; Foroni and Marcellino (2014) briefly describe the main approaches. This section focuses on two specific econometric approaches, which deal with mixed-frequency data in a simple and appealing way: bridge equations and MIDAS regressions (see, for instance, Schumacher (2016) for a recent comparison of these models).

Bridge equations

Bridge equations are one of the most commonly used techniques to link data with different time frequencies. Typically, the series with higher time frequency are, first, aggregated to the (lower) frequency of the dependent variable and, then, included in traditional forecasting models. These models have been widely considered in the literature, especially to forecast GDP growth in national and international institutions (e.g., Baffigi *et al.* 2004, Diron 2008, Barhoumi *et al.* 2012 and Bulligan *et al.* 2015).

Considering y_t sampled at a quarterly frequency (interval of reference) as the dependent variable, the specification of a simple bridge equation with a single indicator and autoregressive terms is given by:

$$y_{t+h} = \beta_0 + \beta(L)x_t^Q + \gamma(L)y_t + \varepsilon_{t+h}, \tag{1}$$

where the predictor x_t^Q is a quarterly variable obtained by aggregating its high-frequency counterpart $x_t^{(m)}$ sampled m times faster (for example, for monthly data m equals 3), h is the quarterly horizon, and ε_{t+h} is a standard i.i.d. error term. The quarterly lag polynomial $\beta(L)$ of order k is defined as $\beta(L) = \sum_{i=0}^{k} \beta_{i+1}L^i$, with $Lx_t^Q = x_{t-1}^Q$. Similarly, $\gamma(L)$ is a p-order polynomial in the lag operator defined as $\gamma(L) = \sum_{i=1}^{p} \gamma_i L^i$, where p is the number of autoregressive terms and $Ly_t = y_{t-1}$. Equation 1 can be easily extended to a multivariate format simply by including additional regressors and each one can have a distinct $\beta(L)$ polynomial.

Depending on the data release lags, the high-frequency indicators may need to be extended with estimates, before being temporally aggregated and included in the bridge model. Considering quarterly and monthly data, estimates for the missing monthly observations, obtained from simple univariate models, are plugged in the monthly data, which are transformed into quarterly series and, then, used for forecasting in the quarterly bridge model.

MIDAS regressions

This section gives a brief overview of the MIDAS regressions used in this article. Armesto *et al.* (2010) provide a simple and a intuitive introduction to the subject and comprehensive discussions of MIDAS regression for short-term forecasting can be found in Andreou *et al.* (2011), Foroni and Marcellino (2014), Schumacher (2016) and references therein. Finally, a recent annals issue of the *Journal of Econometrics* (Ghysels and Marcellino 2016) discusses in detail several econometric methods designed to handle mixed-frequency data.

The MIDAS regressions, introduced by Ghysels *et al.* (2004), are a direct multi-step forecasting tool inspired in the distributed lag models. In addition, as discussed in Duarte (2014), an autoregressive term can simply be added to the MIDAS equation. Consider again y_t sampled at a quarterly frequency and $x_t^{(m)}$ sampled m times faster. A simple MIDAS regression with autoregressive terms is:

$$y_{t+h} = \beta_0 + \beta_1 B(L^{1/m};\theta)x_t^{(m)} + \gamma(L)y_t + \varepsilon_{t+h}, \tag{2}$$

where h is the quarterly horizon, $B(L^{1/m};\theta) = \sum_{j=0}^{jmax} B(j;\theta)L^{j/m}$ is a polynomial of length $jmax$ in the $L^{1/m}$ operator, $B(j;\theta)$ represents the weighting scheme used for the aggregation, which is assumed to be normalised to 1, $L^{j/m}x_t^{(m)} = x_{t-j/m}^{(m)}$, and ε_{t+h} is a standard i.i.d. error term.

Although the order of the polynomial $B(L^{1/m};\theta)$ is potentially infinite, some restrictions must be imposed for the sake of tractability. In a MIDAS regression, the coefficients of $B(L^{1/m};\theta)$ are captured by a known weighting function $B(j;\theta)$, which depends on a few parameters summarized in vector θ. MIDAS models are, thus, tightly parameterised, which is one of the key features of this technique.

Some alternatives for the weighting function have been suggested in the literature; see, namely, Ghysels *et al.* (2007). The most commonly used polynomial is the exponential Almon lag polynomial:

$$B(k;\theta_1,\theta_2) := \frac{e^{(\theta_1 k + \theta_2 k^2)}}{\sum_{k=1}^{K} e^{(\theta_1 k + \theta_2 k^2)}}, \tag{3}$$

where $f(q,\theta_1,\theta_2) = (q^{\theta_1 - 1}(1-q)^{\theta_2 - 1}\Gamma(\theta_1+\theta_2))/(\Gamma(\theta_1)\Gamma(\theta_2))$ and $\Gamma(\theta) = \int_0^\infty e^{-k}k^{\theta-1}dk$. Since the exponential Almon polynomial has a nonlinear functional specification, MIDAS regressions have to be estimated using nonlinear methods, namely nonlinear least squares.

A MIDAS variant discussed by Chen and Ghysels (2011) is the multiplicative MIDAS (M-MIDAS), which is closer to traditional aggregation schemes. Instead of aggregating all lags in the high-frequency variable to a single aggregate, multiplicative MIDAS models include m aggregates of high-frequency data and their lags, i.e.,

$$y_{t+h} = \beta_0 + \sum_{i=1}^{p} \beta_i x_{t-i+1}^{mult} + \gamma(L) y_t + \varepsilon_{t+h}, \tag{4}$$

where $x_t^{mult} = \sum_{j=0}^{m-1} B(j;\theta) L^{j/m} x_t^{(m)}$.

A different MIDAS approach is the unrestricted MIDAS (U-MIDAS) regression proposed by Foroni *et al.* (2015):

$$\begin{aligned} y_{t+h} &= \beta_0 + B_u(L^{1/m}) x_t^{(m)} + \gamma(L) y_t + \varepsilon_{t+h} \\ &= \beta_0 + \sum_{j=0}^{J} \beta_{j+1} L^{j/m} x_t^{(m)} + \gamma(L) y_t + \varepsilon_{t+h} \\ &= \beta_0 + \beta_1 x_t^{(m)} + \beta_2 x_{t-1/m}^{(m)} + \ldots + \beta_{J+1} x_{t-J/m}^{(m)} + \gamma(L) y_t + \varepsilon_{t+h}. \end{aligned} \tag{5}$$

The U-MIDAS regression does not resort to functionals of distributed lag polynomials and, hence, has the advantage that it can be estimated by OLS. However, given the parameter proliferation, the U-MIDAS models are better able to deal with monthly data, than weekly or daily data, as large differences in sampling frequencies between the variables considered are very penalised in terms of parsimony.

Finally, Clements and Galvão (2008) suggested an alternative way of introducing autoregressive dynamics in MIDAS regressions. The authors proposed interpreting the dynamics on y_t as a common factor, resting on the hypothesis that y_{t+h} and $x_t^{(m)}$ share the same autoregressive dynamics. Consider a simple MIDAS regression where the error term can be represented by an autoregressive model of order 1. The common factor MIDAS (CF-MIDAS) model can be written as:

$$(1-\gamma L) y_t = \beta_0(1-\gamma) + \beta_1 (1-\gamma L) B(L^{1/m};\theta) x_t^{(m)} + \varepsilon_t. \tag{6}$$

Although the initial work by Clements and Galvão (2008) only considers a single autoregressive term, it is possible to extend this technique to allow for more autoregressive terms.

In summary, MIDAS models have a more flexible weighting structure than traditional low-frequency models and tend to be more parsimonious. The MIDAS framework can also easily accommodate the timely releases of high-frequency data. In equation 2, it is assumed that all high-frequency observations of $x_t^{(m)}$ over the low-frequency period of reference are known. Considering quarterly and monthly data, this means that the three months of information on the quarter of interest are already available for the short-term indicator. If instead of a full-quarter of data, only, say, the first month is available, then the MIDAS regression can be written as:

$$y_{t+h} = \beta_0 + \beta_1 B(L^{1/3};\theta)x^{(3)}_{t-2/3} + \gamma(L)y_t + \varepsilon_{t+h}. \tag{7}$$

Furthermore, MIDAS regressions can be extended to accommodate additional high-frequency indicators, and, in some cases, without requiring many more parameters to be estimated. Moreover, different polynomials $B(L^{1/m};\theta)$ for each regressor can also be considered.

Data and design of the exercise

Data

The dependent variable is tourism exports from the Portuguese quarterly national accounts at constant prices and seasonally and calendar effects adjusted. Throughout this article, tourism exports refer to the System of National Accounts concept of household final consumption expenditure of tourism of non-resident visitors in Portugal, and does not include the intermediate tourism consumption associated with business travels of non-residents.[1]

Four types of short-term variables related to tourism exports are published monthly and, hence, were the basis of the four individual indicators included in the exercise to nowcast quarterly tourism exports.

Firstly, we use the nominal exports (credits) from the travel account of the Portuguese Balance of Payments (BoP) deflated with the total Harmonised Index of Consumer Prices (HICP).[2]

Secondly, we consider the transactions with cards issued abroad in terminals located in Portugal (ATM/POS). These transactions include both Automated Teller Machines (ATM) cash withdrawals and Points of Sale (POS) transactions and are available since September 2000. The values of the monthly ATM/POS transactions were deflated using the total HICP.[3]

Thirdly, another indicator is the number of non-resident overnight stays in hotel establishments in Portugal. To account for potential quality effects,

1. The detailed data on tourism exports was kindly provided by Statistics Portugal (INE - `http://ine.pt/`).

2. Two other deflators were also tested to price-adjust BoP data. First, the HICP for the services aggregate was used. Second, a composite deflator was built by weighting several price components by their share in the expenditure of tourists in Portugal. We opted for using total HICP, which had the best performance, but the results do not qualitatively change with the two alternative deflators.

3. Similarly to nominal tourism exports from the travel account, we also considered two alternative deflators for the ATM/POS transactions (see footnote 2 for details) and the results remained broadly unchanged.

the number of overnight stays in each type of accommodation establishment was weighted by the respective average total income in the previous year. Five different individual types of hotel establishments (hotels, lodging houses, apartment hotels, tourist villages, tourist apartments) and a residual category (including boarding houses, inns and motels) were considered.[4]

Finally, we calculate a composite index of consumer sentiment in some of the main origin countries of tourists - Spain, the United Kingdom, France, Germany, Italy and the Netherlands. Surveys are particularly valuable because of their timeliness: they are the first monthly releases relating to the current quarter. The monthly consumer confidence indicator of each country published by the Directorate General for Economic and Financial Affairs (DG ECFIN) of the European Commission was weighted by the its importance as an origin of non-resident overnight stays in Portugal in the previous year.[5]

When needed, monthly series were seasonally and calendar effects adjusted. We applied the same procedure used by Statistics Portugal for seasonally adjusting monthly official statistics, namely the X-13 ARIMA with calendar effects adjustment resorting to JDemetra+ software provided by Eurostat. The sample period starts in the October 2000, which corresponds to the first month of the first quarter for which ATM/POS transactions are available, and ends in March 2016. With the exception of the confidence indicator, the original series were transformed to their year-on-year rate of change. In the case of the confidence indicator, absolute differences relative to the same period in the previous year were used.

Design of the exercise

The aim of this article is to nowcast the quarterly growth of Portuguese real tourism exports using four different monthly indicators. For that, we implement a pseudo real-time recursive and direct multi-step exercise with the following features.

All bridge and MIDAS models were recursively estimated with an expanding window and selected using the Bayesian Information Criterion (BIC). Starting from the initial in-sample period (from $2000Q4$ to $2007Q4$) that was used to specify the models, the estimation sample is expanded by adding a new observation in each round. [6] As a new observation is added to the sample, all models are re-estimated and, thus, the coefficients are allowed

4. We also experimented with the raw data on total non-resident overnight stays, but the nowcasting performance was not better.

5. We also used the standard consumer confidence indicators for both the EU and the euro area published monthly by the DG ECFIN and the results were qualitatively similar.

6. We also tested a rolling window and the main results regarding the differences between bridge and MIDAS models do not differ much.

to change over time. Regarding the out-of-sample nowcasting exercise, the evaluation sample covers the period from $2008Q1$ to $2016Q1$.

Different lags were used (up to 3 quarters), also for the autoregressive terms. MIDAS models were estimated using the exponential Almon polynomial defined in equation 3.[7] Bridge equations and the different MIDAS models described in section 3 were estimated with and without autoregressive terms.

An adequate selection of the predictors is crucial for obtaining the best forecast results over the periods considered. Given that in our case the information set comprises a small number of variables, we considered both single- and multi-variable models. In addition, we also tried a different strategy that can improve forecasting accuracy: pooling forecasts. Different pooling techniques are available in the literature, ranging from simple equal (and constant) weights to performance based weights. As simple combination schemes often show good performances, in this article two different pooling techniques are used: the equal-weight mean and the discounted mean squared forecast error (MSFE) combination proposed by Stock and Watson (2004). The Stock and Watson (2004) weights are as follows:

$$w_{it} = \frac{m_{it}^{-1}}{\sum_{i=1}^{n} m_{it}^{-1}} \qquad m_{it} = \sum_{s=t_0}^{T} \delta^{T-s}(y_s - \hat{y}_s^i)^2, \tag{8}$$

where $\hat{y}^i$ are the forecasts from model i and δ is the discount coefficient. The weights of this pooling technique depend inversely on the historical forecasting performance of each model. So, the greater the MSFE of an individual forecast, the smaller the associated weight.[8]

The dataset is a final vintage dataset, meaning that it refers to the latest release available when the database was built. In the case of the consumer confidence indicators and ATM/POS transactions final data equal real-time data, as these series are typically not revised. The revisions to BoP exports, overnight stays and quarterly tourism exports are not taken into account in this analysis but they are usually relatively small in Portugal, so the impact should be minor.

The existence of asynchronous release schedules of high-frequency series implies unbalanced panels with different patterns of missing values in the end of the sample (the so-called "ragged-edge" problem). There is evidence in the literature that accounting for this ragged-edge structure of the dataset can have a considerable impact in nowcast accuracy (see, for instance, Giannone *et al.* (2008)). Hence, we take into account this important characteristic of

7. The traditional Almon lag polynomial was also tested as an alternative for the weighting function. However, it did not improve the performance of the models.

8. Regarding the discount parameter, different values were considered and the non-discounting option ($\delta = 1$) showed the best results.

macroeconomic data in real-time. Following Foroni and Marcellino (2014) and Schumacher (2016), our pseudo real-time exercise mimics the release pattern of the indicators as they become available in real-time situations. More specifically, we replicate the unbalanced structure of the dataset in each of the recursive sub-samples, following a stylised publication calendar: for each series, we observe the number of missing values at the end and impose the same number of missing observations at each recursion.

As discussed in Banbura *et al.* (2011), one important feature of a nowcasting exercise is that one rarely performs a single projection for a given quarter but rather a sequence of nowcasts that are updated as new data arrive. Hence, considering forecasts of quarterly variables on a monthly basis, typically nowcasting refers to the monthly projections of the current quarter and there are at least 3 different projections for that quarter (one in each month of the quarter). However, by taking into account the publication delays of the variable of interest, it is possible to increase the number of projections before there is observed data. For example, Banbura *et al.* (2013) produce nowcasts of US GDP starting in the first month of the current quarter up to the first month of the following quarter, when the official data is published.

In our exercise, we also share this broader perspective about nowcasting. Hence, from the end of the first month of quarter t to the end of the second month of quarter $t+1$, when the official data is observed, it is possible to have up to 5 different nowcasts for quarter t, depending on the information set and the amount of within-quarter data available for each predictor. Given that all monthly indicators are typically observed before the release day of Portuguese quarterly national accounts (recall that the publication delay of expenditure-side aggregates is 60 days after the end of the reference quarter), in the months of their publication, i.e., February, May, August and November, we can obtain an early estimate for tourism exports before the official figure becomes available.

A simple example can help clarifying the structure of the dataset in our pseudo real-time exercise. Assume that one is interested in obtaining a projection of the real growth of Portuguese tourism exports in the first quarter of 2016. In the end of January 2016 (1st m Qt), the consumer confidence indicator is available for January but there is no current quarter information for the other variables: the ATM/POS transactions is available for December 2015 and both BoP exports and overnight stays are available for November 2015. A month later, in the end of February (2nd m Qt), there are two months of current quarter information for the consumer confidence indicator, data for the ATM/POS transactions is available for January, and there is still no current quarter data for the other two variables: both BoP exports and overnight stays are available for December 2015. Again, a month later, in the end of March (3rd m Qt) there are three months of current quarter data for the consumer confidence indicator, two months of data for the ATM/POS transactions and information for both BoP exports and overnight stays is available for January. In addition, from this date onwards, data on quarterly tourism exports for the fourth quarter of 2015 can also be included. In the end of April 2016 (1st m Qt+1), both the consumer confidence indicator and the ATM/POS transactions have three months of data of the quarter of interest and information for BoP exports and overnight stays is available until February. Finally, in the end of May (2nd m Qt+1), full-quarter information for all variables is observed.

In this example, the last two projections are performed in April and May 2016 and refer to the previous quarter. Note that, in contrast with our broad perspective on the term "nowcasting", which allows us to simplify the wording, in some applications, current and previous quarter forecasts are labelled as "nowcasts" and "backcasts", respectively (see Banbura *et al.* (2011)).

As using full-quarter data for all indicators allows having nowcasts for the growth of Portuguese tourism exports in a given quarter only a couple days before the release of the official GDP figures, it is essential to have projections that exploit partial within-quarter information much earlier than that. In the bridge model framework, when not all months of the quarter are available for the predictors, estimates for the missing monthly observations obtained from simple univariate models are used, as described in section 3. All nowcasts are computed directly, i.e., no projections of the dependent variable are used in order to obtain the nowcasts, which implies different bridge models for each quarterly horizon. In the MIDAS framework, the different nowcasts for the quarter of interest are computed using distinct models for each within-quarter information set of the variables, i.e., a new regression is used as new (monthly and quarterly) information is included.

Finally, to evaluate the nowcasting performance of the different bridge and MIDAS models in the out-of-sample period, we used the root mean squared forecast error (RMSE). Relative RMSE are computed to compare the performance of these two approaches with a quarterly benchmark model. The benchmark model is a univariate autoregressive (AR) model, which is estimated recursively, and the lag length (from 0 to 3 lags) is chosen according to the BIC.

Main results

This section presents the results of the pseudo real-time nowcast exercise. As, on average, MIDAS models with AR dynamics outperformed MIDAS regressions without them throughout the whole evaluation periods, in what follows we focus only in the former MIDAS specifications. This finding is in line with other studies that showed that the MIDAS models without an AR component generally perform worse than the MIDAS specifications that include it (see, for instance, Kuzin *et al.* (2011) and Duarte (2014)). In addition, CF-MIDAS regressions were the worst models in terms of nowcast accuracy, so we also excluded them from the analysis.[9]

Regarding the results for single-variable regressions, Figure 3 provides evidence on the performances of the different classes of mixed-frequency models. The figures show the relative RMSE performances, at the different nowcast periods, against an AR benchmark. A ratio lower than 1 denotes a forecasting gain by the bridge and/or MIDAS approaches, whereas a value higher than 1 means that the univariate model outperforms the alternative models. Figure 3 shows heuristically one of the stylised facts of this literature: forecasting accuracy of this type of models tends to increase as time goes by and more information becomes available.

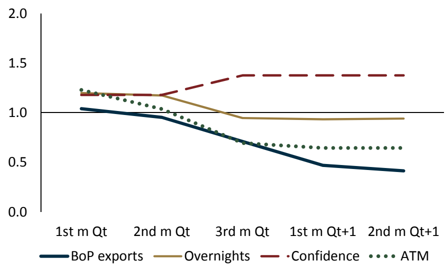

(A) Bridge equations

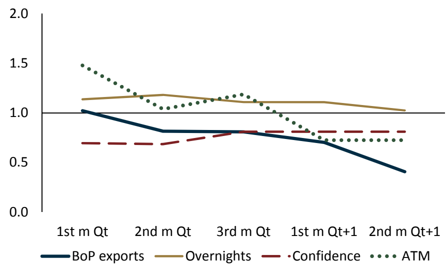

(B) AR-MIDAS regressions

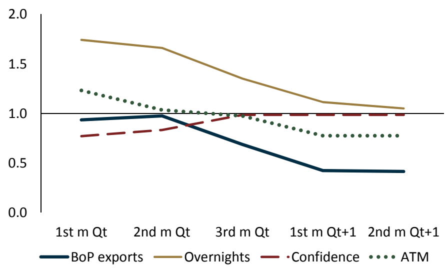

(C) AR-U-MIDAS regressions

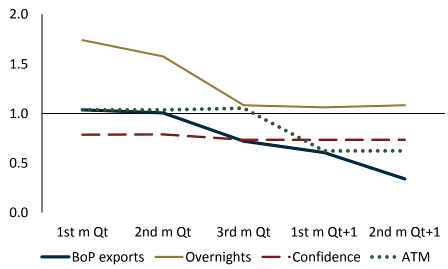

(D) AR-M-MIDAS regressions

FIGURE 3: Relative RMSE of single-variable models (benchmark = AR)

Notes: See Section 4 for a detailed description of the variables and the information used for each nowcast. Ratios of the RMSE with respect to an AR model. A ratio lower than 1 denotes a forecasting gain by the bridge and/or MIDAS models, whereas a value higher than 1 means that the univariate benchmark model outperforms the alternative models.

Starting with bridge equations, in the first two months of the reference quarter, no indicator outperforms the AR benchmark, but, as more data on the quarter is observed, taking into account exports from the BoP travel account,

9. All results are available from the authors upon request.

ATM/POS transactions and, to a lesser extent, overnight stays leads to a RMSE lower than the univariate benchmark in next three evaluation periods. Exports from the BoP travel account are the best indicator in all cases but one: the exception are ATM/POS transactions in the third month of the reference quarter, when there is only one month of BoP exports data, but two months information on ATM/POS transactions.

In contrast, the most accurate MIDAS regressions always outperform the AR benchmark throughout the evaluation periods. Focusing on the short-term indicators, there is a common pattern across the different MIDAS variants: in the first two months of the reference quarter, the best performing indicator is the consumer confidence index; henceforth, BoP exports have the best performance, as more data on this indicator for the reference quarter gradually becomes available. Moreover, the overnight stays variable tends to perform badly in MIDAS models, being worse than the AR benchmark in all cases.

In order to better investigate their properties and capture their differences and similarities over the whole set of individual indicators, Figure 4 provides evidence on the minima and average relative RMSE performances (against an AR benchmark) of the different classes of mixed-frequency models considered. Overall, the best performing model is always MIDAS, i.e., the MIDAS variant with the lowest RMSE always outperforms bridge models and this is true for both minima and average performances. However, in both cases, the best performing MIDAS model is not always the same variant.

Focusing on the minimum relative RMSE, the best nowcasting performance of a MIDAS model is always better than the AR benchmark and allows for gains from around 30 to 65 per cent throughout the whole period. In fact, compared to bridge equations, MIDAS regressions seem to work particularly well for short-term horizons, i.e., when less current-quarter information is available. In contrast, the lack of current-quarter data on BoP exports in the first two evaluation periods is critical for the performance of bridge equations, which never do better than the AR benchmark. In the last three evaluation periods, there are only mild differences between the MIDAS regressions with the lowest relative RMSE and the bridge model. In both cases, the nowcasting gains relatively to the univariate benchmark increase from around 30 per cent to about 60 percent in the last period.

Regarding the average nowcasting performances, it is difficult to outperform the AR benchmark in the first two months of the quarter (the only exception is the AR-MIDAS models in the second period). Moreover, there are no substantial differences between the average performances of the single-variable approaches over the whole period, even if the best MIDAS models perform (slightly) better, on average, than bridge equations in all periods.

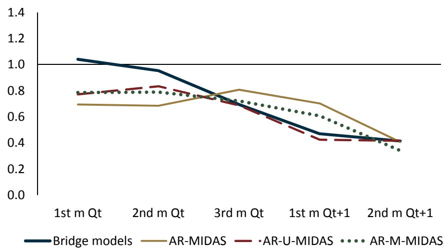

(A) Minima performances

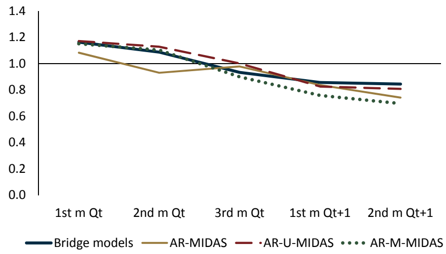

(B) Average performance

FIGURE 4: Minima and average relative RMSE of single-variable models (benchmark = AR)

Notes: See Section 4 for a detailed description of the information used for each nowcast. Minima and average of the relative RMSE ratios with respect to an AR benchmark within a model class across all indicators. A ratio lower than 1 denotes a forecasting gain by the bridge and/or MIDAS models, whereas a value higher than 1 means that the univariate benchmark model outperforms the alternative models.

Nowcast pooling

This section examines the results of the nowcast pooling exercise within each class of models. The combination of individual projections described in equation 8 has a better overall performance than multi-variable models and that the simple equal-weight mean of nowcasts. Hence, we will only analyse the results of the combination of nowcasts using the Stock and Watson (2004) weights.

Figure 5 depicts the minima and average relative RMSE performances of the different models against an AR benchmark, considering all possible combinations of the four single indicators within each model. [10]

Comparing the results included in Figures 4 and 5, it is clear that nowcast pooling is a winning strategy that tends to outperform single-variable models for every period and type of model considered. The finding that pooling of nowcasts is more stable than nowcasting with single models is in line with other studies in the MIDAS literature. Kuzin *et al.* (2013) concluded that pooling outperforms single-variable models for nowcasting quarterly GDP growth and Ghysels and Ozkan (2015) showed that forecast combinations of MIDAS regression models provide gains over traditional models for forecasting the US annual federal budget. Moreover, Clements and Galvão

10. Appendix A includes the detailed results of the nowcast accuracy of the eleven possible combinations for all mixed-frequency models considered: the first table reports the relative RMSE performances of each model against the AR benchmark and the second table includes the RMSE performances of the different MIDAS variants relative to the RMSE of the bridge equations for each combination of predictors.

(2008) found that combinations of MIDAS forecasts are at least as good as combinations of forecasts from bridge models and other mixed-frequency models.

The relative RMSE of pooled nowcasts are always lower than 1 in all cases depicted in Figure 5, implying that not only the best model in each class performs better than the AR benchmark but also that, on average, it is possible to improve nowcasting accuracy by using mixed-frequency models. As in the single-variable models, the best performing model is always a MIDAS regression, both in terms of minima and average performances. Even if the best results are not always obtained from the same type of MIDAS model, the AR-M-MIDAS model delivers good nowcasting results throughout the whole period.

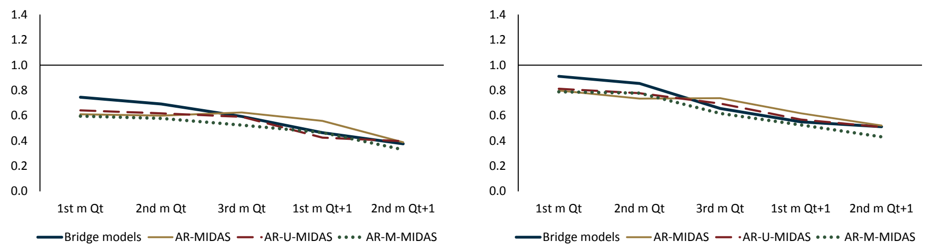

(A) Minima performances

(B) Average performance

FIGURE 5: Minima and average relative RMSE of nowcast pooling (benchmark = AR)

Notes: See Section 4 for a detailed description of the information used for each nowcast. Minima and average of the relative RMSE ratios with respect to an AR benchmark within a model class across all possible combinations of indicators. A ratio lower than 1 denotes a forecasting gain by the bridge and/or MIDAS models, whereas a value higher than 1 means that the univariate benchmark model outperforms the alternative models.

To examine in more detail the performance of nowcast pooling of the different models, Table 1 compares their relative RMSE performances against an AR benchmark. From the eleven possible combinations presented in Appendix A, this table shows the best performing ones for any given bridge/MIDAS model in each evaluation period. Following Foroni and Marcellino (2014), we test the hypothesis of equal accuracy in forecast performance using the Diebold and Mariano (1995) test modified for short samples by Harvey *et al.* (1997). The cases in which the hypothesis of equal forecast accuracy is rejected according to this test are indicated by one of more $*$ in the table, depending on the significance level.

The results show that nowcast pooling of both bridge and MIDAS models performs fairly well: it outperforms the AR benchmark for most of the combinations in each period and the differences in terms of RMSE are, in the vast majority of cases, statistically significant. For instance, in the last

	1st m Qt	2nd m Qt	3rd m Qt	1st m Qt+1	2nd m Qt+1
Bridge models					
Overnights + Confidence	0.861	0.845	0.755 *	0.769 *	0.786
ATM + Confidence	0.754 *	0.727	0.623 ***	0.599 ***	0.599 ***
BoP exports + Confidence	**0.745** *	**0.691** **	0.694 **	0.493 ***	0.450 ***
BoP exports + Overnights	1.046	0.981	0.672 ***	**0.464** ***	**0.376** ***
BoP exports + Overnights + Confidence	0.805 *	0.762 **	0.633 **	0.473 ***	0.406 ***
BoP exports + Confidence + ATM	0.794 *	0.747 **	**0.592** ***	0.483 ***	0.435 ***
BoP exports + Overnights + Confidence + ATM	0.848 *	0.802 **	0.604 ***	0.488 ***	0.427 ***
AR-MIDAS					
Overnights + Confidence	0.611 **	**0.599** **	0.656 **	0.655 **	0.738 *
ATM + Confidence	**0.609** **	0.604 **	0.657 **	0.614 **	0.614 **
BoP exports + Confidence	0.683 **	0.631 **	0.662 **	0.662 **	0.433 ***
BoP exports + Overnights	0.995	0.822 **	0.770 **	0.574 ***	**0.388** ***
BoP exports + Overnights + Confidence	0.668 **	0.613 ***	**0.624** ***	0.569 ***	0.422 ***
BoP exports + Confidence + ATM	0.657 **	0.624 ***	0.637 ***	0.561 ***	0.434 ***
BoP exports + Overnights + Confidence + ATM	0.679 ***	0.637 ***	0.653 ***	**0.557** ***	0.439 ***
AR-U-MIDAS					
Overnights + Confidence	0.823	0.814	0.922	0.754 *	0.810
ATM + Confidence	0.715 *	0.666 *	0.662 ***	0.588 ***	0.588 ***
BoP exports + Confidence	0.655 **	0.668 **	0.631 ***	0.495 ***	0.407 ***
BoP exports + Overnights	0.880	0.849	0.658 ***	**0.425** ***	0.407 ***
BoP exports + Overnights + Confidence	0.679 **	0.676 **	0.629 ***	0.518 ***	0.407 ***
BoP exports + Confidence + ATM	**0.641** ***	**0.617** ***	**0.589** ***	0.467 ***	**0.397** ***
BoP exports + Overnights + Confidence + ATM	0.688 **	0.656 **	0.599 ***	0.503 ***	0.405 ***
AR-M-MIDAS					
Overnights + Confidence	0.634 **	0.619 **	0.582 ***	0.577 ***	0.582 ***
ATM + Confidence	**0.595** ***	**0.577** **	0.544 ***	0.509 ***	0.509 ***
BoP exports + Confidence	0.633 **	0.620 **	0.615 **	0.593 ***	0.355 ***
BoP exports + Overnights	1.088	1.022	0.647 ***	0.492 ***	**0.329** ***
BoP exports + Overnights + Confidence	0.632 ***	0.609 ***	0.550 ***	0.503 ***	0.338 ***
BoP exports + Confidence + ATM	0.632 ***	0.611 ***	0.538 ***	0.484 ***	0.359 ***
BoP exports + Overnights + Confidence + ATM	0.659 ***	0.641 ***	**0.524** ***	**0.465** ***	0.358 ***

TABLE 1. Relative RMSE performance of nowcast pooling against an AR benchmark

Notes: See Section 4 for a detailed description of the variables and the information used for each nowcast. Ratios of the RMSE with respect to an AR model. A ratio lower than 1 denotes a forecasting gain by the bridge and/or MIDAS models. $*$, $**$ and $***$ indicate the forecasts which are statistically superior to the ones from the benchmark at a confidence level of 10, 5 and 1 per cent, respectively, according to the Diebold and Mariano (1995) test modified for short samples by Harvey *et al.* (1997). The numbers in bold denote the minimum relative RMSE within each model and in each period. The numbers dark-shaded and with white font denote the minimum relative RMSE for each evaluation period across all models and combinations of indicators. The light-shaded areas represent the cases where the MIDAS model is statistically superior to the respective bridge equation (at least at a 10 per cent significance level).

three evaluation periods, it is possible to obtain a projection that is statistically superior to the univariate benchmark in 96.4 per cent of the cases.

Focusing on the best combinations within each model in each period (the bold numbers in the table), they all provide results that are statistically superior to the AR benchmark. The good performances are shared by different combinations but a common feature emerges across all models: the consumer confidence indicator is always part of the best performing combination in the first two periods and BoP exports are always included in the best combination in the last three projection moments.

The overall minimum relative RMSE for each evaluation period across all models and combinations of indicators (the numbers dark-shaded with white font in the table) is always produced by a MIDAS model: the AR-M-MIDAS in four cases and the AR-U-MIDAS in one case. Not only the best performing MIDAS specification changes over time, but the best combination of predictors also changes for the different within-quarter information sets of the variables. The AR-M-MIDAS model works particularly well in the three months of the reference quarter: it delivers the best results in first two months of the reference quarter by combining ATM/POS transactions and consumer confidence and, in the last month of the quarter, by combining the four individual indicators. In the last two evaluation periods, when more data is already observed for the reference quarter, the preferred combination is BoP exports and overnight stays, first obtained from the AR-U-MIDAS model and, then, from the AR-M-MIDAS model in the last period.

Another way to compare the alternative mixed-frequency models is to compute the RMSE of the different MIDAS models relative to the RMSE of the bridge equations for each combination of indicators. The light-shaded areas in Table 1 represent the cases where the forecasts from a MIDAS model are statistically superior to the respective forecasts from the bridge equation, at least at a 10 per cent significance level.

The most useful forecast combinations of MIDAS models should outperform both the AR benchmark and the competing bridge equation (Schumacher 2016). In Table 1 , these are cases where the $*$s are light-shaded. The statistically significant improvements of MIDAS models relative to both benchmarks simultaneously occur, in particular, in the first two periods: in around 35 per cent of the cases for AR-U-MIDAS and in more than 70 per cent of the cases for both AR-MIDAS and AR-M-MIDAS. Considering the five evaluation periods and the best seven combinations of indicators, AR-M-MIDAS is the model with best overall performance: it delivers projections that are statistically better that both benchmarks in around 57 per cent of the cases. Overall, and taking into account all evaluation periods, nowcast combinations that comprise 3 and 4 indicators tend to be more reliable, in the sense that they tend to outperform both benchmarks more frequently than combinations with less indicators.

Final remarks

Tourism exports are an extremely important component of Portuguese international trade of goods and services. Short-run forecasts of this variable play a relevant role in the monitoring of Portuguese economic activity and external accounts.

The purpose of this article is to nowcast the real growth of quarterly tourism exports using four different monthly indicators in a recursive pseudo real-time exercise. We resort to two single-equation approaches that deal with mixed-frequency data: bridge equations and MIDAS regressions. Bridge equations are one of the most used techniques to link monthly and quarterly variables. In these models, the variables on both sides of the equation are on the same (low) frequency: in our case, monthly indicators are aggregated to their corresponding quarterly values. In contrast, in MIDAS regressions, the observations of the low-frequency dependent variable are linked directly to high-frequency observations of the predictors without any previous temporal aggregation. Different specifications of bridge and MIDAS models with single indicators and combination of nowcasts are evaluated in this article.

The results obtained suggest that, as expected, using mixed-frequency models with short-term indicators contributes to increase nowcast accuracy in comparison to a univariate benchmark. In general, MIDAS models tend to fare better than traditional bridge models for the majority of the predictors and evaluation periods, but the differences are higher when less current-quarter information is available. Nowcast combinations of both bridge and MIDAS regressions always provide gains over single-indicator models. In fact, a general finding common to all mixed-frequency models considered is that the AR benchmark can always be outperformed by the best performing combination of nowcasts in every evaluation period and that the differences in terms of relative RMSE are statistically significant. Overall, the best performing nowcast is always obtained from a combination of projections of a MIDAS variant with AR dynamics, which suggests the use of this class of mixed-frequency models for short-term forecasting tourism exports.

References

Andraz, Jorge L.M., Pedro M.D.C.B. Gouveia, and Paulo M.M. Rodrigues (2009). "Modelling and forecasting the UK tourism growth cycle in Algarve." *Tourism Economics*, 15(2), 323–338.

Andreou, Elena, Eric Ghysels, and Andros Kourtellos (2011). "Forecasting with Mixed-Frequency Data." In *The Oxford Handbook of Economic Forecasting*, edited by Michael P. Clements and David F. Hendry, chap. 8, pp. 225–267. Oxford University Press.

Armesto, Michelle T., Kristie M. Engemann, and Michael T. Owyang (2010). "Forecasting with mixed frequencies." *Federal Reserve Bank of St. Louis Review*, (Nov/Dec), 521–536.

Athanasopoulos, George, Rob J. Hyndman, Haiyan Song, and Doris C. Wu (2011). "The tourism forecasting competition." *International Journal of Forecasting*, 27(3), 822–844.

Baffigi, Alberto, Roberto Golinelli, and Giuseppe Parigi (2004). "Bridge models to forecast the euro area GDP." *International Journal of Forecasting*, 20(3), 447–460.

Banbura, Marta, Domenico Giannone, Michele Modugno, and Lucrezia Reichlin (2013). "Now-Casting and the Real-Time Data Flow." In *Handbook of Economic Forecasting*, vol. 2, Part A, edited by Graham Elliott and Allan Timmermann, chap. 4, pp. 195–237. Elsevier.

Banbura, Marta, Domenico Giannone, and Lucrezia Reichlin (2011). "Nowcasting." In *The Oxford Handbook of Economic Forecasting*, edited by Michael P. Clements and David F. Hendry, chap. 7, pp. 193–224. Oxford University Press.

Barhoumi, Karim, Olivier Darné, Laurent Ferrara, and Bertrand Pluyaud (2012). "Monthly GDP forecasting using bridge models: Application for the French economy." *Bulletin of Economic Research*, 64, s53–s70.

Bulligan, Guido, Massimiliano Marcellino, and Fabrizio Venditti (2015). "Forecasting economic activity with targeted predictors." *International Journal of Forecasting*, 31(1), 188–206.

Cardoso, Fátima and Cláudia Duarte (2006). "The use of qualitative information for forecasting exports." *Banco de Portugal Economic Bulletin*, Winter, 67–74.

Chen, Xilong and Eric Ghysels (2011). "News – Good or Bad – and Its Impact on Volatility Predictions over Multiple Horizons." *Review of Financial Studies*, 24(1), 46–81.

Clements, M. P. and A. B. Galvão (2008). "Macroeconomic Forecasting With Mixed-Frequency Data: Forecasting Output Growth in the United States." *Journal of Business and Economic Statistics*, 26(4), 546–554.

Daniel, Ana Cristina M. and Francisco F. R. Ramos (2002). "Modelling inbound international tourism demand to Portugal." *International Journal of Tourism Research*, 4(3), 193–209.

Diebold, Francis X. and Roberto S. Mariano (1995). "Comparing Predictive Accuracy." *Journal of Business and Economic Statistics*, 13(3), 253–63.

Diron, Marie (2008). "Short-term forecasts of euro area real GDP growth: an assessment of real-time performance based on vintage data." *Journal of Forecasting*, 27(5), 371–390.

Duarte, Cláudia (2014). "Autoregressive augmentation of MIDAS regressions." Working Paper 01/2014, Banco de Portugal.

Duarte, Cláudia, Paulo M. M. Rodrigues, and António Rua (2016). "A Mixed Frequency Approach to Forecast Private Consumption with ATM/POS Data." *International Journal of Forecasting*, Forthcoming.

Esteves, Paulo Soares (2009). "Are ATM/POS Data Relevant When Nowcasting Private Consumption?" Working Paper 25/2009, Banco de Portugal.

Esteves, Paulo Soares and António Rua (2012). "Short-term forecasting for the Portuguese economy: A methodological overview." *Banco de Portugal Economic Bulletin*, Autumn, 137–149.

Foroni, C. and M. Marcellino (2014). "A comparison of mixed frequency approaches for nowcasting Euro area macroeconomic aggregates." *International Journal of Forecasting*, 30, 554–568.

Foroni, Claudia, Massimiliano Marcellino, and Christian Schumacher (2015). "Unrestricted mixed data sampling (MIDAS): MIDAS regressions with unrestricted lag polynomials." *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 178(1), 57–82.

Ghysels, E., P. Santa-Clara, and R. Valkanov (2004). "The MIDAS touch: Mixed data sampling regressions." Discussion paper, UNC and UCLA.

Ghysels, Eric and Massimiliano Marcellino (2016). "The econometric analysis of mixed frequency data sampling." *Journal of Econometrics*, 193(2), 291 – 293.

Ghysels, Eric and Nazire Ozkan (2015). "Real-time forecasting of the US federal government budget: A simple mixed frequency data regression approach." *International Journal of Forecasting*, 31(4), 1009 – 1020.

Ghysels, Eric, Arthur Sinko, and Rossen Valkanov (2007). "MIDAS Regressions: Further Results and New Directions." *Econometric Reviews*, 26(1), 53–90.

Giannone, Domenico, Lucrezia Reichlin, and David Small (2008). "Nowcasting: The real-time informational content of macroeconomic data." *Journal of Monetary Economics*, 55(4), 665–676.

Gouveia, Pedro M.D.C.B. and Paulo M.M. Rodrigues (2005). "Dating and synchronizing tourism growth cycles." *Tourism Economics*, 11(4), 501–515.

Harvey, David, Stephen Leybourne, and Paul Newbold (1997). "Testing the equality of prediction mean squared errors." *International Journal of Forecasting*, 13(2), 281 – 291.

Kuzin, Vladimir, Massimiliano Marcellino, and Christian Schumacher (2011). "MIDAS vs. mixed-frequency VAR: Nowcasting GDP in the euro area." *International Journal of Forecasting*, 27(2), 529 – 542.

Kuzin, Vladimir, Massimiliano Marcellino, and Christian Schumacher (2013). "Pooling versus model selection for nowcasting GDP with many predictors: Empirical evidence for six industrialized countries." *Journal of Applied Econometrics*, 28(3), 392–411.

Maria, José Ramos and Sara Serra (2008). "Forecasting investment: A fishing contest using survey data." Working Paper 18/2008, Banco de Portugal.

Rodrigues, Paulo M.M. and Pedro M.D.C.B. Gouveia (2004). "An application of PAR models for tourism forecasting." *Tourism Economics*, 10(3), 281–303.

Schumacher, Christian (2016). "A comparison of MIDAS and bridge equations." *International Journal of Forecasting*, 32(2), 257–270.

Serra, Jaime, Antónia Correia, and Paulo M.M. Rodrigues (2014). "A comparative analysis of tourism destination demand in Portugal." *Journal of Destination Marketing and Management*, 2(4), 221–227.

Song, Haiyan and Rob J. Hyndman (2011). "Tourism forecasting: An introduction." *International Journal of Forecasting*, 27(3), 817–821.

Song, Haiyan and Gang Li (2008). "Tourism demand modelling and forecasting - A review of recent research." *Tourism Management*, 29(2), 203–220.

Song, Haiyan, Gang Li, Stephen F. Witt, and George Athanasopoulos (2011). "Forecasting tourist arrivals using time-varying parameter structural time series models." *International Journal of Forecasting*, 27(3), 855–869.

Stock, James H. and Mark W. Watson (2004). "Combination forecasts of output growth in a seven-country data set." *Journal of Forecasting*, 23(6), 405–430.

Teixeira, João Paulo and Paula Odete Fernandes (2012). "Tourism Time Series Forecast - Different ANN Architectures with Time Index Input." *Procedia Technology*, 5, 445–454.

Teixeira, João Paulo and Paula Odete Fernandes (2014). "Tourism time series forecast with artificial neural networks." *Tékhne*, 12(1-2), 26–36.

Appendix: Detailed results of nowcast pooling for all models considered

	1st m Qt	2nd m Qt	3rd m Qt	1st m Qt+1	2nd m Qt+1
Bridge models					
Overnights + ATM	1.144	1.056	0.731***	0.679***	0.679***
Overnights + Confidence	0.861	0.845	0.755*	0.769*	0.786
ATM + Confidence	0.754*	0.727	0.623***	0.599***	0.599***
BoP exports + Confidence	0.745*	0.691**	0.694**	0.493***	0.450***
BoP exports + Overnights	1.046	0.981	0.672***	0.464***	0.376***
BoP exports + ATM	1.103	0.984	0.619***	0.481***	0.421***
Overnights + Confidence + ATM	0.822	0.801*	0.652***	0.624***	0.627***
BoP exports + Overnights + Confidence	0.805*	0.762**	0.633**	0.473***	0.406***
BoP exports + Confidence + ATM	0.794*	0.747**	0.592***	0.483***	0.435***
BoP exports + Overnights + ATM	1.094	0.995	0.638***	0.492***	0.417***
BoP exports + Overnights + Confidence + ATM	0.848*	0.802**	0.604***	0.488***	0.427***
AR-MIDAS					
Overnights + ATM	1.196	1.186	1.127	0.813***	0.758**
Overnights + Confidence	0.611**	0.599**	0.656**	0.655**	0.738*
ATM + Confidence	0.609**	0.604**	0.657**	0.614**	0.614**
BoP exports + Confidence	0.683**	0.631**	0.662**	0.662**	0.433***
BoP exports + Overnights	0.995	0.822**	0.770**	0.574***	0.388***
BoP exports + ATM	1.022	0.857**	0.807**	0.561***	0.421***
Overnights + Confidence + ATM	0.624***	0.623***	0.689***	0.642***	0.649**
BoP exports + Overnights + Confidence	0.668**	0.613***	0.624***	0.569***	0.422***
BoP exports + Confidence + ATM	0.657**	0.624***	0.637***	0.561***	0.434***
BoP exports + Overnights + ATM	1.033	0.881**	0.830**	0.583***	0.426***
BoP exports + Overnights + Confidence + ATM	0.679***	0.637***	0.653***	0.557***	0.439***
AR-U-MIDAS					
Overnights + ATM	1.313	1.195	0.910	0.800**	0.728***
Overnights + Confidence	0.823	0.814	0.922	0.754*	0.810
ATM + Confidence	0.715*	0.666*	0.662***	0.588***	0.588***
BoP exports + Confidence	0.655**	0.668**	0.631***	0.495***	0.407***
BoP exports + Overnights	0.880	0.849	0.658***	0.425***	0.407***
BoP exports + ATM	0.846*	0.815**	0.673***	0.511***	0.415***
Overnights + Confidence + ATM	0.794	0.751*	0.712**	0.628***	0.613***
BoP exports + Overnights + Confidence	0.679**	0.676**	0.629***	0.518***	0.407***
BoP exports + Confidence + ATM	0.641***	0.617***	0.589***	0.467***	0.397***
BoP exports + Overnights + ATM	0.899	0.835*	0.658***	0.553***	0.417***
BoP exports + Overnights + Confidence + ATM	0.688**	0.656**	0.599***	0.503***	0.405***
AR-M-MIDAS					
Overnights + ATM	1.154	1.226	0.945	0.662***	0.668***
Overnights + Confidence	0.634**	0.619**	0.582***	0.577***	0.582***
ATM + Confidence	0.595***	0.577**	0.544***	0.509***	0.509***
BoP exports + Confidence	0.633**	0.620**	0.615**	0.593***	0.355***
BoP exports + Overnights	1.088	1.022	0.647***	0.492***	0.329***
BoP exports + ATM	0.978	0.981	0.649***	0.489***	0.359***
Overnights + Confidence + ATM	0.620***	0.613***	0.545***	0.516***	0.520***
BoP exports + Overnights + Confidence	0.632***	0.609***	0.550***	0.503***	0.338***
BoP exports + Confidence + ATM	0.632***	0.611***	0.538***	0.484***	0.359***
BoP exports + Overnights + ATM	1.044	1.035	0.647***	0.483***	0.363***
BoP exports + Overnights + Confidence + ATM	0.659***	0.641***	0.524***	0.465***	0.358***

TABLE A.1. Relative RMSE of nowcast pooling against the AR benchmark

Notes: See Section 4 for a detailed description of the variables and information used for each nowcast. Ratios of the RMSE with respect to an AR model. A ratio lower than 1 denotes a forecasting gain by the bridge and/or MIDAS models. $*$, $**$ and $***$ indicate the forecasts which are significantly more accurate than the benchmark at a confidence level of 10, 5 and 1 per cent, respectively, according to the Diebold and Mariano (1995) test modified for short samples by Harvey *et al.* (1997).

	1st m Qt	2nd m Qt	3rd m Qt	1st m Qt+1	2nd m Qt+1
AR-MIDAS					
Overnights + ATM	1.045	1.123	1.542	1.197	1.116
Overnights + Confidence	0.709***	0.708***	0.869	0.852*	0.938
ATM + Confidence	0.808***	0.830**	1.056	1.026	1.026
BoP exports + Confidence	0.917	0.912	0.954	1.343	0.964
BoP exports + Overnights	0.951	0.837	1.146	1.237	1.032
BoP exports + ATM	0.927	0.871	1.304	1.165	0.999
Overnights + Confidence + ATM	0.760***	0.778***	1.057	1.029	1.035
BoP exports + Overnights + Confidence	0.829**	0.804**	0.986	1.203	1.041
BoP exports + Confidence + ATM	0.828**	0.835*	1.076	1.162	0.999
BoP exports + Overnights + ATM	0.945	0.886	1.301	1.185	1.020
BoP exports + Overnights + Confidence + ATM	0.801**	0.794**	1.081	1.141	1.029
AR-U-MIDAS					
Overnights + ATM	1.147	1.132	1.245	1.179	1.072
Overnights + Confidence	0.956	0.963	1.222	0.980	1.030
ATM + Confidence	0.948	0.916	1.064	0.982	0.982
BoP exports + Confidence	0.879	0.967	0.909	1.004	0.905*
BoP exports + Overnights	0.841	0.865	0.979	0.916	1.084
BoP exports + ATM	0.767*	0.829	1.088	1.061	0.987
Overnights + Confidence + ATM	0.966	0.938	1.091	1.006	0.978
BoP exports + Overnights + Confidence	0.843*	0.887	0.992	1.095	1.002
BoP exports + Confidence + ATM	0.807**	0.827*	0.995	0.966	0.912
BoP exports + Overnights + ATM	0.822*	0.839	1.031	1.125	1.000
BoP exports + Overnights + Confidence + ATM	0.812**	0.819*	0.992	1.031	0.948
AR-M-MIDAS					
Overnights + ATM	1.008	1.161	1.294	0.974	0.984
Overnights + Confidence	0.736***	0.733***	0.771***	0.750**	0.740
ATM + Confidence	0.789***	0.793***	0.874	0.850**	0.850**
BoP exports + Confidence	0.850**	0.897	0.887	1.202	0.788***
BoP exports + Overnights	1.040	1.042	0.963	1.060	0.876
BoP exports + ATM	0.887	0.998	1.049	1.016	0.853**
Overnights + Confidence + ATM	0.754***	0.765***	0.836*	0.827**	0.829**
BoP exports + Overnights + Confidence	0.785***	0.800***	0.868*	1.062	0.833**
BoP exports + Confidence + ATM	0.797***	0.819***	0.908	1.003	0.827***
BoP exports + Overnights + ATM	0.955	1.040	1.014	0.982	0.870
BoP exports + Overnights + Confidence + ATM	0.777***	0.800***	0.867	0.953	0.838**

TABLE A.2. Relative RMSE of nowcast pooling against bridge models

Notes: See Section 4 for a detailed description of the variables and the information used for each nowcast. Ratios of the RMSE with respect to bridge models. A ratio lower than 1 denotes a forecasting gain by the MIDAS models. $*$, $**$ and $***$ indicate the forecasts which are significantly more accurate than the benchmark at a confidence level of 10, 5 and 1 per cent, respectively, according to the Diebold and Mariano (1995) test modified for short samples by Harvey *et al.* (1997).