

Universidade de Lisboa
ISEG - Instituto Superior de Economia e Gestão



**Essays on Mixed-Frequency Data:
Forecasting and Unit Root Testing**

CLÁUDIA FILIPA PIRES DUARTE

Orientador(es): Prof. Doutor Paulo Manuel Marques Rodrigues

Prof. Doutor João Carlos Henriques da Costa Nicolau

Tese especialmente elaborada para obtenção do grau de Doutor em Economia

2016

Universidade de Lisboa
ISEG - Instituto Superior de Economia e Gestão



Essays on Mixed-Frequency Data: Forecasting and Unit Root Testing

CLÁUDIA FILIPA PIRES DUARTE

Orientador(es): Prof. Doutor Paulo Manuel Marques Rodrigues

Prof. Doutor João Carlos Henriques da Costa Nicolau

Tese especialmente elaborada para obtenção do grau de Doutor em Economia

Júri:

Presidente: Prof. Doutor Manuel Fernando Cília de Mira Godinho, Professor Catedrático e Presidente do Conselho Científico do Instituto Superior de Economia e Gestão da Universidade de Lisboa

Vogais: Doutor Antonio Manuel Pedro Afonso, Professor Catedrático do Instituto Superior de Economia e Gestão da Universidade de Lisboa

Doutor Luís Miguel Rainho Catela Nunes, Professor Associado da Faculdade de Economia da Universidade Nova de Lisboa

Doutor Paulo Manuel Marques Rodrigues (orientador) Professor Associado Convidado da Faculdade de Economia da Universidade Nova de Lisboa

Doutor Luís Filipe Farias de Sousa Martins (relator) Professor Auxiliar com Agregação do ISCTE - Instituto Universitário de Lisboa

Doutor Nuno Ricardo Martins Sobreira (relator) Professor Auxiliar Convidado do Instituto Superior de Economia e Gestão da Universidade de Lisboa

To my parents

Contents

Resumo	xi
Abstract	xiii
Acknowledgments	xv
Introduction	xvi
1 Macroeconomic Forecasting with MIDAS	1
1.1 Introduction	1
1.2 Overview of the literature	6
1.3 MIDAS modelling	10
1.3.1 Background	11
1.3.2 From the approximate mixed-frequency regression to MIDAS regressions	14
1.3.3 Autoregressive augmentation of MIDAS regressions	20
1.4 Design of the nowcasting and forecasting exercise	28
1.5 Empirical results	32
1.5.1 Single-variable models	32
1.5.2 Pooling forecasts	39
1.5.3 Multi-variable models	42
1.6 Conclusion	44
Acknowledgments	45
Bibliography	46
Appendices	54
2 Covariate-augmented unit root tests with mixed-frequency data	61
2.1 Introduction	61
2.2 Covariate-augmented unit root tests	63
2.2.1 The CADF test	66
2.2.2 Tests with GLS demeaning	70

2.3	Mixed-frequency covariate-augmented unit root tests	74
2.3.1	The MIDAS approach	76
2.3.2	The M-CADF test	78
2.3.3	The M-CADF-GLS test	78
2.4	Monte Carlo simulation	79
2.4.1	Baseline results	84
2.4.2	The case of a larger sample size	91
2.4.3	Different lags	93
2.4.4	Different time frequencies	94
2.5	An application to the US unemployment rate	96
2.6	Conclusion	99
	Acknowledgments	100
	Bibliography	101
	Appendices	106
3	Unit root tests using mixed-frequency VAR models	109
3.1	Introduction	109
3.2	Feasible point optimal unit root tests	111
3.2.1	The ERS test	111
3.2.2	EJ test	114
3.3	Mixed-frequency EJ unit root test	118
3.3.1	VAR-MIDAS approach	118
3.3.2	M-EJ test	121
3.4	Monte Carlo simulation	122
3.4.1	Baseline results	126
3.4.2	The case of a larger sample size	134
3.4.3	Different lags	136
3.5	An application to the US unemployment rate	137
3.6	Conclusion	141
	Acknowledgments	141
	Bibliography	142
	Appendices	147
	Conclusion	150

List of Tables

A.1	Single-variable models: relative performance, in terms of RMSFE, against an AR benchmark, for forecast horizon h	56
A.2	Single-variable models: relative performance, in terms of RMSFE, against a traditional low-frequency quarterly benchmark, for forecast horizon h	57
A.3	Single-variable model for different time frequencies, against an autoregressive benchmark, for forecast horizon h	58
A.4	Pooled forecasts: relative performance, in terms of RMSFE, against an AR benchmark, for forecast horizon h	59
A.5	Multi-variable models, against an AR benchmark, for forecast horizon h	60
2.4.1	Simulation design and median of the lag order selected by MAIC (Perron and Qu, 2007)	83
2.4.2	Finite sample size for unit root tests considering nominal size of 5 per cent, $T = 100$	85
2.4.3	Size-adjusted power of unit root tests, Constant only, $T = 100$. . .	88
2.4.4	Size-adjusted power of unit root tests, Time trend included, $T = 100$	89
2.4.5	Estimated R^2 and δ parameters and critical values, Constant only .	90
2.5.1	Unit root tests for US monthly unemployment rate, using jobless claims as covariate	98
2.5.2	Unit root tests for US monthly unemployment rate, using jobless claims as covariate	99
B.1	Asymptotic critical values for CADF t -statistics	108
B.2	Asymptotic critical values for the CADF-GLS test	108
3.4.1	Simulation design and median of the lag order selected by AIC . . .	125
3.4.2	Finite sample size for unit root tests considering nominal size of 5 per cent, $T = 100$	127
3.4.3	Size-adjusted power of unit root tests, Constant only, $T = 100$. . .	130
3.4.4	Size-adjusted power of unit root tests, Time trend included, $T = 100$	131

3.4.5	Estimated R^2 and critical values of EJ-family unit root tests, Constant only	133
3.5.1	Unit root tests for US monthly unemployment rate, using jobless claims as covariate	139
C.1	Asymptotic critical values for the EJ test	149

List of Figures

1.1	Simulated sequences of the coefficients associated with $x_t^{(3)}$ and its lags	22
1.2	Forecast performance of MIDAS single-variable regressions compared to an AR benchmark, relative RMSFE	33
1.3	Forecast performance of MIDAS single-variable regressions compared to a quarterly benchmark, relative RMSFE	33
1.4	Best performing single-variable regression, for each forecast horizon, relative RMSFE compared to AR benchmark	34
1.5	Single-variable model - Dow Jones Euro Stoxx index - for each forecast horizon, relative RMSFE compared to AR benchmark	36
1.6	Relative RMSFE of MIDAS regressions compared to an AR benchmark	37
1.7	Relative RMSFE of MIDAS regressions compared to a quarterly benchmark	37
1.8	Pooled forecasts - best performing MIDAS models against an AR benchmark for each forecast horizon, relative RMSFE	41
1.9	MIDAS regressions - single-variable models vs. pooled forecasts against an AR benchmark, for each forecast horizon, relative RMSFE	42
1.10	MIDAS models: pooled forecasts vs. multi-variable models - for each forecast horizon, relative RMSFE against an AR benchmark	43
2.1	Finite sample size distortions vis-à-vis a nominal size of 5 per cent, $T = 100$, Constant only	86
2.2	Differences in size-adjusted power between mixed- and low-frequency tests with GLS demeaning for $c = -5, -10$ and -15 , Constant only, $T = 100$	87
2.3	Finite sample size distortions vis-à-vis a nominal size of 5 per cent, $T = 500$	92
2.4	Differences in size-adjusted power between mixed- and low-frequency tests with GLS demeaning for $c = -5, -10$ and -15 , Constant only, $T = 500$	93

2.5	Differences in size-adjusted power between mixed- and low-frequency tests with GLS demeaning for $c = -5$ by number of lags included in the test regression, Constant only	94
2.6	Size-adjusted power of mixed- and low-frequency tests for different m , with GLS demeaning, aggregate $\alpha = 0.99$, Constant only, $T = 500$	95
2.7	Data on US unemployment	97
3.1	Finite sample size, $T = 100$, Constant only	128
3.2	Differences in size-adjusted power between mixed- and low-frequency EJ tests for $\alpha = 0.95, 0.90$ and 0.85 , Constant only, $T = 100$	132
3.3	Finite sample size, $T = 500$, Constant only	135
3.4	Differences in size-adjusted power between mixed- and low-frequency EJ tests for $\alpha = 0.99, 0.98$ and 0.97 , Constant only, $T = 500$	135
3.5	Differences in size-adjusted power between mixed- and low-frequency tests with GLS demeaning by number of lags included in the test regression, Constant only	137
3.6	Data on US unemployment	138

Resumo

Nas últimas décadas, os investigadores têm tido acesso a bases de dados cada vez mais abrangentes, que incluem séries com frequências temporais mais elevadas e que são divulgadas mais atempadamente. Em contraste, algumas variáveis, nomeadamente alguns dos principais indicadores macroeconómicos, são divulgados com um desfasamento temporal significativo e com baixa frequência.

Esta situação levanta questões sobre como lidar com séries com frequências temporais diferentes, mistas. Ao longo do tempo, várias técnicas têm sido propostas. Esta tese debruça-se sobre uma técnica em particular — a abordagem *MI(xed) DA(ta) S(ampling)*, proposta por Ghysels et al. (2004).

No Capítulo 1 eu utilizo a técnica *MIDAS* para prever o crescimento do PIB na área do euro com base num pequeno conjunto de indicadores, cobrindo séries com diferentes frequências temporais e divulgadas com diferentes desfasamentos. Eu comparo o desempenho de um conjunto alargado de regressões *MIDAS*, utilizando a raiz quadrada do erro quadrático médio de previsão e tomando como ponto de referência quer regressões autoregressivas, quer multivariadas (*bridge models*). A questão sobre a forma de introduzir termos autoregressivos nas equações *MIDAS* é dirimida. São consideradas diferentes combinações de variáveis, obtidas através da agregação de previsões ou de regressões multivariadas, assim como diferentes frequências temporais. Os resultados sugerem que, em geral, a utilização de regressões *MIDAS* contribui para o aumento da precisão das previsões.

Adicionalmente, nesta tese são propostos novos testes de raízes unitárias que exploram informação com frequências mistas. Tipicamente, os testes de raízes unitárias têm baixa potência, especialmente em amostras pequenas. Uma forma de combater esta dificuldade consiste em recorrer a testes que exploram informação

adicional de um regressor estacionário incluído na regressão de teste. Eu avalio se é possível melhorar o desempenho de alguns testes deste tipo ao explorar dados com frequências temporais mistas, através de regressões *MIDAS*. No Capítulo 2 eu proponho uma nova classe de testes da família Dickey-Fuller (DF) com regressores adicionais de frequência temporal mista, tomando por base os testes DF com regressores adicionais (CADF) propostos por Hansen (1995) e uma versão modificada proposta por Pesavento (2006), semelhante ao filtro GLS aplicado ao teste ADF univariado em Elliott et al. (1996).

Em alternativa aos testes da família DF, Elliott and Jansson (2003) propõem um teste de raízes unitárias viável que retém propriedades ótimas mesmo na presença de variáveis determinísticas (EJ), tomando por base a versão univariada proposta por Elliott et al. (1996). No Capítulo 3 eu alargo o âmbito de aplicação destes testes de forma a incluir dados com frequência temporal mista. Dado que para implementar o teste EJ é necessário estimar modelos VAR, eu proponho um modelo VAR-*MIDAS* não restrito, parcimonioso, que inclui séries de frequência temporal mista e é estimado com técnicas econométricas tradicionais.

Os resultados de um exercício de Monte Carlo indicam que os testes com dados de frequência temporal mista têm um desempenho em termos de potência melhor do que os testes que agregam todas as variáveis para a mesma frequência temporal (necessariamente a frequência mais baixa). Os ganhos são robusto à dimensão da amostra, à escolha do número de defasamentos a incluir nas regressões de teste e às frequências temporais concretas. Adicionalmente, os testes da família EJ tendem a ter um melhor desempenho do que os testes da família CADF, independentemente das frequências temporais consideradas. Para ilustrar empiricamente a utilização destes testes, analisa-se a série da taxa de desemprego nos EUA.

Palavras-Chave: Regressões MIDAS, Dados de frequência alta, Previsão, Precisão das previsões, PIB da área do euro, Raízes unitárias, Testes de hipóteses, Dados de frequência mista, VAR-MIDAS não restrito, Taxa de desemprego dos EUA.

Abstract

Over the last decades, researchers have had access to more comprehensive datasets, which are released on a more frequent and timely basis. Nevertheless, some variables, namely some key macroeconomic indicators, are released with a significant time delay and at low frequencies.

This situation raises the question on how to deal with series released at different, mixed time frequencies. Over the years and for different purposes, several techniques have been put forward. This essay focuses on a particular technique — the MI(xed) DA(ta) S(ampling) framework, proposed by Ghysels et al. (2004).

In Chapter 1 I use MIDAS for forecasting euro area GDP growth using a small set of selected indicators in an environment with different sampling frequencies and asynchronous releases of information. I run a horse race between a wide set of MIDAS regressions and evaluate their performance, in terms of root mean squared forecast error, against AR and quarterly bridge models. The issue on how to include autoregressive terms in MIDAS regressions is disentangled. Different combinations of variables, through forecast pooling and multi-variable regressions, and different time frequencies are also considered. The results obtained suggest that, in general, using MIDAS regressions contributes to increase forecast accuracy.

In addition, I propose new unit root tests that exploit mixed-frequency information. Unit root tests typically suffer from low power in small samples. To overcome this shortcoming, tests exploiting information from stationary covariates have been proposed. I assess whether it is possible to improve the power performance of some of these tests by exploiting mixed-frequency data, through the MIDAS approach. In Chapter 2 I put forward a new class of mixed-frequency covariate-augmented Dickey-Fuller (DF) tests, extending the covariate-augmented

DF test (CADF test) proposed by Hansen (1995) and its modified version, similar to the GLS generalisation of the univariate ADF test in Elliott et al. (1996), proposed by Pesavento (2006).

Alternatively to the CADF tests, Elliott and Jansson (2003) proposed a feasible point optimal unit root test in the presence of deterministic components (EJ test, hereafter), which extended the univariate results in Elliott et al. (1996). In Chapter 3 I go one step further and include mixed-frequency data in the EJ testing framework. Given that implementing the EJ test requires estimating VAR models, in order to plug in mixed-frequency data in the test regression I propose an unconstrained, though parsimonious, stacked skip-sampled reduced-form VAR-MIDAS model, which is estimated using standard econometric techniques.

The results from a Monte Carlo exercise indicate that mixed-frequency tests have better power performance than low-frequency tests. The gains are robust to the size of the sample, to the lag specification of the test regressions and to different combinations of time frequencies. Moreover, the EJ-family of tests tends to have a better power performance than the CADF-family of tests, either with low or mixed-frequency data. An empirical illustration using the US unemployment rate is presented.

Keywords: MIDAS regressions, High-frequency data, Forecasting, Forecast accuracy, Euro Area GDP, Unit root, Hypothesis testing, Mixed-frequency data, Unrestricted VAR-MIDAS, US unemployment rate.

Acknowledgments

First of all, I would like to thank my advisers, Professors João Nicolau and Paulo Rodrigues, for guiding me through this unique journey.

I would also like to thank Banco de Portugal for giving me the opportunity of pursuing this project, while at work, and for the financial support. As I learned over time, combining regular work with such a long-haul project is not always an easy task. A kind acknowledgement is also dedicated to Ricardo Félix for his commitment to providing research time amid the regular activities.

This project benefited from thoughtful discussions with Paulo Júlio, José Francisco Maria and Christian Schumacher, and from comments of participants at the 24th (EC)² Conference “The Econometrics Analysis of Mixed Frequency Data”, held at the University of Cyprus, and at “Exchange” (Banco de Portugal seminars). Detailed acknowledgements follow at the end of each chapter. A special thanks is owed to Fátima Teodoro for software assistance.

I would like to express my deepest gratitude to Carlos Robalo Marques, whose comments and suggestions were crucial for the development of this project.

I thank my colleagues and friends, for their encouragement, friendship and for supporting me throughout this process. A special word is dedicated to João Cardoso, for his generosity, encouragement and care.

Finally, I thank my parents for all their love, support and encouragement. For always being there.

Cláudia Duarte

Lisbon, July 2015

Introduction

In the last decades, researchers have had access to more and more comprehensive datasets. Not only the number of variables available has increased significantly, but data are also released on a more frequent and timely basis. Nevertheless, some variables, namely some key macroeconomic indicators, are released with a significant time delay and at low frequencies.

This situation raises the question on how to deal with this diversity of data, in particular on how to deal with series released at different, mixed time frequencies. Over the years and for different purposes, several techniques have been put forward in order to combine data with different frequencies and exploit timely releases of high frequency data. This thesis focuses on one particular technique — the MI(xed) DA(ta) S(ampling) framework, proposed by Ghysels et al. (2004).

In Chapter 1 I use MIDAS for forecasting in an environment with different sampling frequencies and asynchronous releases of information. The goal is forecasting quarterly euro area GDP growth using a small set of selected indicators. I run a horse race between a wide set of MIDAS regressions - original and multiplicative, with and without autoregressive terms, with and without current period information - and evaluate their performance, in terms of root mean squared forecast error, against AR and quarterly bridge models. Different combinations of variables, through forecast pooling and multi-variable regressions, and different time frequencies, namely daily, weekly and monthly frequencies, are also considered.

Adding to the extensive comparison between MIDAS models, I contribute to the existing literature by disentangling the question on how to include autoregressive terms in MIDAS regressions.

The results obtained suggest that, in general, using MIDAS regressions con-

tributes to increase forecast accuracy. Flexible MIDAS weighting schemes are able to exploit the information content of high-frequency series for forecasting purposes, regardless of the exact time frequency of the regressors.

Another issue addressed herein is the power performance when testing for the presence of unit roots. Unit root tests typically suffer from low power in small samples, which results in not rejecting the null hypothesis as often as they should. To overcome this shortcoming, alternative tests exploiting information from covariates have been proposed. I assess whether it is possible to improve the power performance of covariate-augmented unit root tests by exploiting mixed-frequency data through the MIDAS approach.

In Chapter 2 I put forward a new class of covariate-augmented Dickey-Fuller (DF) unit root tests that is able to deal with mixed-frequency data. Hansen (1995) generalised the augmented DF test to include covariates - the CADF test. The intuition is that including a weakly exogenous and stationary variable in the auxiliary test regression may lead to efficiency gains. More recently, Pesavento (2006) proposed a modified version of the CADF test, similar to the GLS generalisation of the ADF test in Elliott et al. (1996). I use MIDAS regressions to extend these two covariate-augmented DF tests to take on board mixed frequency data — the mixed-frequency CADF test (M-CADF) and the mixed-frequency CADF test with GLS demeaning/detrending (M-CADF-GLS).

Alternatively to the CADF tests, Elliott and Jansson (2003) proposed a feasible point optimal unit root test in the presence of deterministic components (EJ test, hereafter), which extended the univariate results in Elliott et al. (1996). In spite of having a slightly worse size performance, the authors concluded that the EJ test outperforms the CADF test in terms of power.

In Chapter 3 I go one step further and include mixed-frequency data in the EJ testing framework. Given that implementing the EJ test requires estimating VAR models, in order to plug in mixed-frequency data in the test regression I use the MIDAS vector autoregressive (VAR-MIDAS) approach. Building on Ghysels (2012), I propose an unconstrained, though parsimonious, stacked skip-sampled reduced-form VAR-MIDAS model, which is estimated using standard econometric techniques.

An extensive comparison of a suite of unit root tests, with low- or mixed-frequency data, is presented. Both in Chapter 2 and Chapter 3 the results from a Monte Carlo exercise indicate that mixed-frequency tests have better power performance than low-frequency tests. Hence, covariate-augmented unit root tests also seem to benefit from exploiting high-frequency data. The gains are robust to the size of the sample, to the lag specification of the test regressions and to different combinations of time frequencies. Moreover, the EJ-family of tests tends to have a better power performance than the CADF-family of tests, either with low or mixed-frequency data. An empirical illustration using the US unemployment rate is also presented, covering all unit root tests mentioned throughout this thesis.

Finally, the main results and lines for future research are summarised in the last chapter.

Chapter 1

Macroeconomic Forecasting with MIDAS

ABSTRACT

Techniques that allow taking full advantage of the timely releases of high-frequency data play a key role in forecasting and policymaking, especially nowadays because researchers have access to larger datasets. In this paper, I focus on the MI(xed) DA(ta) S(ampling) framework for handling different sampling frequencies and asynchronous releases of information. For forecasting quarterly euro area GDP growth using a small set of selected indicators, I run a horse race between a wide set of MIDAS regressions - original and multiplicative, with and without autoregressive terms, with and without current period information - and evaluate their performance, in terms of root mean squared forecast error, against AR and quarterly bridge models. Alternative techniques for the autoregressive augmentation are presented and discussed. Different combinations of variables, through forecast pooling and multi-variable regressions, and different time frequencies are also considered. The results obtained suggest that, in general, using MIDAS regressions contributes to increase forecast accuracy.

JEL Classification: C53, E37.

Keywords: MIDAS regressions, High-frequency data, Forecasting.

1.1 Introduction

Nowadays researchers have access to more ‘comprehensive’ datasets. Not only the number of variables available has increased significantly, but data are also released on a more frequent and timely basis. However, some key macroeconomic

indicators are released with a significant time delay and at low frequencies (e.g., national account aggregates are typically published at quarterly frequency and with a delay of six weeks in the case of euro area GDP flash estimate and four weeks in the case of US GDP advance estimate). Therefore, forecasting models are usually set up with all variables sampled at the same frequency (the lower frequency), not capturing directly the timely releases of higher frequency data (e.g., industrial production, retail turnover and confidence indicators on a monthly basis, and stock prices, oil prices and interest rates on a daily basis).

Taking into account timely releases of high frequency data is quite relevant for forecasting purposes mainly for two reasons. First, forecast accuracy is not invariant to the data frequency used (see Marcellino, 1999, for a discussion on the effects of temporal aggregation on several time series properties).

Second, current period estimates (i.e. nowcasts) can benefit from within-period releases of high-frequency information. For instance, consider the release of the euro area GDP flash estimate for the fourth quarter of 2010. This figure is only available in mid-February of 2011, but until then several daily, weekly and monthly indicators relative to the fourth quarter of 2010 have been released, though not all at the same time. Results from previous studies suggest that estimates that include current period information relative to high frequency series outperform purely low frequency estimates (see, among others, Rünstler et al., 2009, for euro area and individual countries GDP estimates). Nowadays, this is even more relevant because the regular need for estimates on the current state of the economy in order to support policy decisions is more pressing.

In this paper, I focus on the MI(xed) DA(ta) S(ampling) framework, proposed by Ghysels et al. (2004) for handling different sampling frequencies and asynchronous releases of information. Inspired in the distributed lag models, the 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 regressions see Andreou et al., 2011). MIDAS regressions were originally associated with empirical applications to financial series, focusing on volatility predictions; see Ghysels et al. (2004, 2006, 2007). However, MIDAS regressions have also been used in typical macroeconomic

forecasting applications (see, among others, Clements and Galvão, 2008 and Bai et al., 2013 for US GDP growth, Armesto et al., 2009 for the predictive content of the Beige Book, Kuzin et al., 2011 for euro area GDP, Marcellino and Schumacher, 2010 for German GDP, Monteforte and Moretti, 2013 for euro area inflation, and Asimakopoulos et al., 2013 for fiscal variables of a set of European countries).

Over the years other techniques have been used to combine data with different frequencies and exploit timely releases of high frequency data for improving forecast accuracy. Some authors suggested interpolating the lower frequency data in order to specify a model in the higher frequency (see, for example, Liu, 1969, Liu and Hwa, 1974 and Engle et al., 1989). Alternative approaches include pooling forecasts from two models (e.g., a quarterly model and a monthly model) as in Corrado and Greene (1988), Howrey et al. (1991), Donihue and Howrey (1992), Miller and Chin (1996), Stark (2000) and Ruey-Wan and Chung-Hua (1996). In a slightly different vein, Klein and Sojo (1989), Parigi and Schlitzler (1995), Ingenito and Trehan (1996), Pain and Sédillot (2005), Zheng and Rossiter (2006) and Golinelli and Parigi (2007) used bridge models to link quarterly series with monthly indicators aggregated quarterly. Within a single-model framework, alternative contributions were made by Rathjens and Robins (1993), Koenig et al. (2003) and Abeysinghe (1998).

Also among the model-driven (dis)aggregation techniques, and assuming the higher frequency observations as missing, there is a vast literature on estimating such missing observations relying on state-space interpretations and Kalman filtering (see, for example, Howrey, 1991, Casals et al., 2009, Zadrozny, 1990, Mitnik and Zadrozny, 2004 and Hyung and Granger, 2008). For large data sets, factor models can also be used to address the question of forecasting using mixed-frequency data (see Mariano and Murasawa, 2003, Nunes, 2005 and Proietti, 2011, among others). These models can also deal with the existence of asynchronous release schedules of high frequency series, which implies unbalanced panels with different patterns of missing values in the end of the sample - the so-called ragged-edge problem (see Altissimo et al., 2010, Doz et al., 2012 and Pinheiro et al., 2013 for alternative techniques).

In practice, regardless of the modelling approach used, there is not a theoretical

answer to the question on whether using high frequency data improves forecast accuracy over traditional low-frequency models and how that information should be used. Evidence on the improvement of forecast accuracy derived from using disaggregate series is not straightforward when the specification and the parameters of the model are no longer assumed to be known. Results in Lütkepohl (1984) suggest that the forecasts from the aggregated process will be superior to the aggregated forecasts from the disaggregated process for long forecast horizons. Moreover, using data with higher frequency does not necessarily mean using more information (Granger, 1998). For instance, let us consider data observed at daily and quarterly frequencies. On the one hand, daily data may contain useful information for improving the accuracy of quarterly forecasts. Nevertheless, on the other hand, daily data may be too noisy and the signals conveyed by within-quarter releases should be downplayed.

In general, empirical results suggest that taking into account timely releases of information tend to contribute to the increase of the nowcast and forecast accuracy. Moreover, when compared with other techniques, the MIDAS approach is rather appealing because it is a simple, flexible and potentially parsimonious single-regression framework. For example, Bai et al. (2013) and Kuzin et al. (2011) concluded that MIDAS regressions outperform the state space approach, particularly in the short run. Moreover, compared to factor models, small-scale MIDAS regressions are easier to interpret.

Given that empirical exercises play a crucial role in the assessment of the best forecasting practices, the aim of this paper is to provide some useful insights on how to combine different time frequencies for nowcasting and forecasting using the MIDAS framework. The contribution of this paper to the existing literature is twofold. Firstly, for forecasting quarterly euro area GDP growth, I run a horse race between an extended set of MIDAS regressions and evaluate their performance, in terms of root mean squared forecast error (RMSFE), against AR and other traditional quarterly models. Secondly, I address the issue of autoregressive augmentation, discussing alternative perspectives on the implications of introducing autoregressive terms in MIDAS regression. Ghysels et al. (2007) noted that including these terms led to discontinuities in the impulse response function of the

regressors on the variable of interest. To rule out these discontinuities, Clements and Galvão (2008) proposed imposing a common factor restriction on the autoregressive dynamics. In this paper, I question the meaning of these discontinuities for the relevant impulse response function and assess alternatives for taking into account autoregressive dynamics in MIDAS regressions, namely by not imposing a common factor and by using multiplicative and balanced autoregressive MIDAS regressions.

As the evidence in favour of using large information sets is not clear-cut (see, for example, Banerjee and Marcellino, 2006), in this paper I focus on a small set of selected indicators, comparing the forecasting performance of single indicator models with forecast pooling and multiple variable models. In addition to different combinations of variables, different time frequencies are also considered. In particular, the possibility of including in the same model different high-frequency indicators, current period information and different time frequencies enables the assessment of whether there is a marginal gain from adding higher frequency data. This paper also addresses the question on whether daily data is still useful if timely monthly information is included in the model. Moreover, for a given indicator, the best frequency for forecasting purposes - lowest, highest or something in between - is assessed.

The forecasts are obtained through a recursive out-of-sample exercise, which takes into account the ragged-edges of the high-frequency data and the publication delay of GDP. This pseudo real-time exercise mimics the release pattern of the indicators as they become available in real-time situations. To assess whether the gains in terms of forecast accuracy from taking on board high-frequency data are short-lived, past high-frequency data is combined to obtain nowcasts (current period forecasts) and direct forecast for different horizons, up to 4 quarters ahead.

The remainder of this paper is organised as follows. Section 1.2 presents a brief description of the literature on mixed-frequency data. Section 1.3 describes the main topics related with MIDAS modelling and discusses in detail the different techniques for the autoregressive augmentation of MIDAS regressions. The design of the now- and forecasting exercise is presented in Section 1.4, while Section 1.5 focuses on the results. Finally, Section 1.6 concludes.

1.2 Overview of the literature

On the wake of the improvement of data collection and treatment techniques, richer high-frequency datasets represent an opportunity for their users, especially those concerned with improving forecast accuracy, whether they are economic researchers or policymakers. Intuitively, one can put it this way - having larger data sets is an advantage simply because more information is better than less information. The case of datasets comprising a large number of series has been widely discussed in the popular strand of literature on factor models (see, among other, Stock and Watson, 2002, Forni et al., 2000 and Forni et al., 2005). Regarding time frequencies, some theoretical results support the idea that when confronted with the choice between using higher frequency series or using their temporal aggregates one should choose the disaggregate series. For example, Geweke (1978) showed that as the interval between observations goes to zero the discrete time approximation will be closer to the underlying continuous time model, which means that increased temporal disaggregation may help reducing the discrete time aggregation bias.

Moreover, within discrete time models, Lütkepohl (1986) showed that, when the data generation process is a known vector ARMA, using the disaggregate series for forecasting purposes is at least as efficient, in terms of mean squared error, as using the temporally aggregated series, even if one is primarily interested in the less frequent series.

However, using higher frequency data entails important challenges. Evidence on the improvement of forecast accuracy derived from using disaggregate series is not so clear-cut when the specification and the parameters of the model are no longer assumed to be known. In this case, the results in Lütkepohl (1984) suggest that forecasts drawn from the model for the aggregated series will be superior to aggregating forecasts from the model for the disaggregated series for long forecast horizons. In practice, although using data at a higher frequency provides more observations, it is not clear that it also provides more information (Granger, 1998).

Furthermore, as not all variables are available at higher frequencies, handling mixed frequency data is not straightforward. Typically, to deal with this situation one converts the high frequency data, through temporal aggregation or systematic sampling, so that all variables in the database have the same (lower) temporal

frequency. However, this approach may not only lead to information losses and changes in causality relations in the case of temporal aggregation (see Weiss, 1984), but may also defeat the purpose of using readily available variables.

Some authors have also suggested interpolating the lower frequency data in order to specify a model in the higher frequency (see, for example, Liu, 1969, Liu and Hwa, 1974 and Engle et al., 1989). For interpolating the low frequency data different techniques can be used, such as those suggested by Chow and Lin (1971), Fernandez (1981), Litterman (1983), and more recently by Santos Silva and Cardoso (2001). Additionally, Proietti (2006) and Gómez and Aparicio-Pérez (2009) used the state-space framework, while Angelini et al. (2006) considered the factor model technique for interpolation and backdating.

Other alternative approaches have been used over the years. Recognising that monthly information has long been used to try to improve quarterly forecasts based on judgemental adjustments (informal methods), Corrado and Greene (1988) discussed a systematic way of pooling forecasts, based on add-factors, in order to improve the predictive accuracy of quarterly macroeconomic models using monthly information, making this process more systematic and transparent. While Corrado and Greene (1988) aimed at improving the forecasting performance of the quarterly Federal Reserve Board model, Howrey et al. (1991), Howrey (1991) and Donihue and Howrey (1992) applied a similar procedure to the Michigan Quarterly Econometric Model (MQEM) for the US economy. Combining forecasts from two models - a quarterly model and a monthly model - in order to maximise forecast accuracy was also considered by Miller and Chin (1996), Stark (2000) and Ruey-Wan and Chung-Hua (1996). In a slightly different vein, Klein and Sojo (1989), Parigi and Schlitzer (1995), Ingenito and Trehan (1996), Pain and Sédillot (2005), Zheng and Rossiter (2006) and Golinelli and Parigi (2007) used bridge models to link quarterly series with monthly indicators aggregated quarterly. In order to have a full quarter, for the months that have not yet been released the quarterly aggregates of the monthly series are obtained by extrapolation through simple univariate time series models (for example, ARIMA models).

So far, all the techniques mentioned always considered separate models for different frequencies. Now, suppose the aim is the same as before - using monthly

information for improving the forecast accuracy of a quarterly model, but, instead of using different models, one wants to combine both time frequencies into a single model. Rathjens and Robins (1993) proposed an approach, which takes into account within-quarter information. Instead of obtaining quarterly estimates by simply aggregating monthly data (whether observed or forecast), these authors suggested adding a new data series to the low frequency model, associated with the dependent variable, defined as the difference between the third month of the quarter and the simple average of the quarter. This new series does not allow incorporating current quarter information, only information regarding movements within the previous quarters. Their results point to the existence of significant improvements on forecast accuracy from using the within-quarter variable, in the case of one-step ahead forecasts. However, this approach has some caveats: (i) it cannot be implemented if the variable of interest is not observed at both low and high time frequency; (ii) multi-step ahead forecasting is rather complex, as it requires specifying a model for the added variable; and, (iii) as within-quarter developments may be difficult to model, the gains from including this information in forecasting models may be short-lived.

Koenig et al. (2003) proposed a similar approach of exploiting monthly data, also inspired on the fact that lower frequency rates of change are a weighted combination of higher frequency rates of change. The authors regressed quarterly changes in real GDP on a constant and five lagged month-on-month changes of the monthly indicators considered. No autoregressive terms were added and the coefficients were estimated unrestrictedly. If contemporaneous correlation between the variables is relevant, these models can only be estimated when values for all months in the quarter are already known.

Alternatively, Abeysinghe (1998) started with a distributed lag model defined at the higher frequency, which was subsequently multiplied by a specific lag polynomial in order to convert the fractional lags of the dependent variable (not observed in the low frequency variables) into integer lags. Implicitly, the resulting lag structure gives more weight to the more recent high frequency observations. With this technique, adding too many variables has important costs in terms of degrees of freedom. Moreover, if the contemporaneous correlation between current

quarter dependent and independent variables is significant, then the model can be estimated only after the release of data for the third month of the quarter (or monthly forecasts have to be incorporated). Considering data for Singapore's GDP and external trade, the author concluded that the quarterly model with monthly data improved forecast accuracy in terms of root mean squared errors over an univariate benchmark and a quarterly model, but the gains were marginal.

As noted by Abeysinghe and Tay (2000), the above-mentioned technique can be applied to both stock and flow variables. While for stock variables the interpretation is direct, for flow variables an aggregation polynomial has to be added. In the case of flow variables, combining the aggregation of the variables in the model (including the error variable) with the lag structure (lagged terms of the dependent variable) generates an autocorrelation problem, as mentioned by Abeysinghe and Tay (2000) and Abeysinghe (2000). To solve this problem, the use of instrumental variables is suggested but the results are not reported.

Also among the model-driven (dis)aggregation techniques, and assuming the higher frequency observations as missing, there is a vast literature on estimating such missing observations relying on state-space interpretations and Kalman filtering (Howrey, 1991). This technique uses the serial dependence of the data in order to estimate conditional expectations of the missing observations, delivering as a by-product estimates for the unobserved values. Nevertheless, this approach entails some difficulties. In particular, state space models are typically larger and more complex models, which require the explicit specification of the dynamics of the series involved and include a large number of parameters. Casals et al. (2009) used this approach to disaggregate and forecast annual series of value added by industry in Spain, using as indicator the quarterly series of the industrial production index. Zadrozny (1990) also used the Kalman filter to estimate a multivariate ARMA model for obtaining forecasts for quarterly US GNP using monthly observations on total employment. The author concluded that this approach has a better forecasting performance in the short-run than traditional models, such as VAR models, estimated using low frequencies. Moreover, Mittnik and Zadrozny (2004) obtained similar results for German GDP. For forecasting quarterly US GDP using monthly industrial production series, Hyung and Granger (2008) also casted an

ARMA model into state-space representation and estimated it, using the Kalman filter. This model was named by the authors as linked-ARMA model. According to their results, the improvements in terms of forecast accuracy of the linked-ARMA model over a traditional quarterly model and the Rathjens and Robins' (1993) approach are not significant. However, using within-current-quarter information clearly improves forecast accuracy and the linked-ARMA model significantly outperforms the alternatives considered.

Factor models can also be used for forecasting using mixed-frequency data. For example, as in Mariano and Murasawa (2003), Nunes (2005) extended the Stock and Watson (1989) factor model by including indicators available at different frequencies and concluded that this approach delivers better forecasting performance than AR(1) models. For the euro area, Proietti (2011) estimated a large scale factor model combining monthly series with quarterly national accounts series. When combining large amounts of high frequency information and estimating factor models, 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). To tackle this issue, different techniques for estimating the factors have been used, namely realigning the time series in order to obtain a balanced panel (Altissimo et al., 2010), using the Expectation-Maximisation (EM) algorithm within a static principal component analysis Stock and Watson (2002) or in a dynamic factor model framework (see, among others, Doz et al., 2012 and Pinheiro et al., 2013). In general, the results point to an increase in nowcast and forecast accuracy from introducing within-period information (see, for example, Giannone et al., 2008).

1.3 MIDAS modelling

After presenting some theoretical motivation and notation (Section 1.3.1), the MIDAS approach is discussed in Section 1.3.2. A detailed description of the different weighting schemes underlying MIDAS regressions is presented in Section 1.3.2. Finally, in Section 1.3.3 alternative methods for introducing autoregressive terms in MIDAS regressions are reviewed and a new perspective on this issue is proposed.

1.3.1 Background

Consider the traditional low frequency regression,

$$Y_{t+h} = \alpha + \beta Q_t + \varepsilon_{t+h} \quad (1.1)$$

where h denotes the forecast horizon (when $h = 0$ the model delivers nowcasts) and both Y_{t+h} and Q_t are sampled at a low frequency, e.g., quarterly. All the parameters of the model depend on the forecast horizon and forecasts are computed directly, i.e., no additional forecasts of the explanatory variables are needed in order to obtain forecasts for the variable of interest. Equation 1.1 can be extended to include lags of the Y and Q variables, as well as additional regressors and respective lags, as follows

$$G(L)Y_{t+h} = \alpha + \sum_{i=1}^N \beta_i(L)Q_{i,t} + \varepsilon_{t+h} \quad (1.2)$$

where $G(L)$ and $\beta_i(L)$ are finite order lag polynomials, L is the conventional lag operator and N is the number of regressors. This equation is a typical autoregressive distributed lag (ADL) regression.

Now, assume that Y_{t+h} and Q_t are temporal aggregates of higher frequency disaggregated series, e.g., monthly series (y_{t+h} and x_t , respectively). For each low-frequency (quarterly) observation of Y_{t+h} and Q_t there are m (3) observations of the high-frequency (monthly) y_{t+h} and x_t series. The low-frequency variables are characterised by an aggregation scheme $\Gamma(L^{1/m})$, which can be one of the following:

1. stock variable

$$Y_t = \begin{bmatrix} 1 & 0 & 0 \end{bmatrix} \times \begin{bmatrix} y_t^{(3)} \\ y_{t-1/3}^{(3)} \\ y_{t-2/3}^{(3)} \end{bmatrix} \quad (1.3)$$

2. flow variable

(a) sum

$$Y_t = \begin{bmatrix} 1 & 1 & 1 \end{bmatrix} \times \begin{bmatrix} y_t^{(3)} \\ y_{t-1/3}^{(3)} \\ y_{t-2/3}^{(3)} \end{bmatrix} \quad (1.4)$$

(b) equal-weight average

$$Y_t = \begin{bmatrix} 1/3 & 1/3 & 1/3 \end{bmatrix} \times \begin{bmatrix} y_t^{(3)} \\ y_{t-1/3}^{(3)} \\ y_{t-2/3}^{(3)} \end{bmatrix} \quad (1.5)$$

(c) rate of change (assume that $\tilde{y}_t^{(3)}$ are the levels underlying the $y_t^{(3)}$ rates of change)

$$Y_t = \begin{bmatrix} \frac{\tilde{y}_{t-1}^{(3)}}{\sum_{i=0}^2 \tilde{y}_{t-1-(i/3)}^{(3)}} & \frac{\tilde{y}_{t-4/3}^{(3)}}{\sum_{i=0}^2 \tilde{y}_{t-1-(i/3)}^{(3)}} & \frac{\tilde{y}_{t-5/3}^{(3)}}{\sum_{i=0}^2 \tilde{y}_{t-1-(i/3)}^{(3)}} \end{bmatrix} \times \begin{bmatrix} y_t^{(3)} \\ y_{t-1/3}^{(3)} \\ y_{t-2/3}^{(3)} \end{bmatrix} \quad (1.6)$$

(d) unrestricted linear combination

$$Y_t = \begin{bmatrix} \gamma_0 & \gamma_1 & \gamma_2 \end{bmatrix} \times \begin{bmatrix} y_t^{(3)} \\ y_{t-1/3}^{(3)} \\ y_{t-2/3}^{(3)} \end{bmatrix} \quad (1.7)$$

where $y_t^{(3)}$ is a skip-sampled version of the high-frequency y_t series, meaning that $y_t^{(3)}$ and Y_t have the same number of quarterly observations.

In time series analysis, observed time series are often temporal aggregates of unobserved disaggregated series. As before, assume for instance that Q_t is observed monthly (x_t), while Y_{t+h} is only observed quarterly - meaning that y_{t+h} is not observed (flow variable) or has missing observations (stock variable). Although regression analysis would ideally try to approximate the original data generating

process (DGP) by using high-frequency samples, as in the following equation

$$g(L^{1/m})y_{t+h} = a + \sum_{i=1}^N b_i(L^{1/m})x_{i,t} + e_{t+h} \quad (1.8)$$

where $g(L^{1/m})$ and $b_i(L^{1/m})$ are finite order lag polynomials, this is not always possible.

The mixed-frequency approaches suggest mixing a low-frequency dependent variable on the left-hand side with high-frequency regressors on the right-hand side. Assuming that high-frequency y_{t+h} would be well represented by equation 1.8, there is a $\phi(L^{1/m})$ polynomial such that $h(L) = g(L^{1/m})\phi(L^{1/m})$. Multiplying both sides of equation 1.8 by $\phi(L^{1/m})$ and the aggregation scheme $\Gamma(L^{1/m})$ one obtains

$$\begin{aligned} \phi(L^{1/m})g(L^{1/m})\Gamma(L^{1/m})y_{t+h} &= \ddot{a} + \sum_{i=1}^N b_i(L^{1/m})\phi(L^{1/m})\Gamma(L^{1/m})x_{i,t} \\ &\quad + \phi(L^{1/m})\Gamma(L^{1/m})e_{t+h} \\ h(L)Y_{t+h} &= \ddot{a} + \sum_{i=1}^N b_i(L^{1/m})z_{i,t} + \xi_{t+h} \end{aligned} \quad (1.9)$$

where $\ddot{a} = \phi(L^{1/m})\Gamma(L^{1/m})a$, $z_{i,t} = \phi(L^{1/m})\Gamma(L^{1/m})x_{i,t}$ and $\xi_{t+h} = \phi(L^{1/m})\Gamma(L^{1/m})e_{t+h}$ for $t = m, 2m, 3m, \dots$. Equation 1.9 is the exact mixed-frequency model associated with the high-frequency model in equation 1.8, relating the low-frequency Y_{t+h} dependent variable with the high-frequency $x_{i,t}$ regressors.

As discussed in Wei and Stram, 1990, Marcellino, 1998 and Marcellino, 1999, among others, in general it is not possible to uniquely identify $\phi(L^{1/m})$ in this single-equation framework and, so, neither the $g(L^{1/m})$ and $b_i(L^{1/m})$ polynomials. This means that, in general, one can only approximate the mixed-frequency model, as follows

$$\tilde{h}(L)Y_{t+h} = \tilde{a} + \sum_{i=1}^N \tilde{b}_i(L^{1/m})x_{i,t}^{(3)} + u_{t+h} \quad (1.10)$$

where $x_{i,t}^{(3)}$ is the skip-sampled version of the high-frequency $x_{i,t}$ and the orders

of the polynomials $\tilde{h}(L)$ and $\tilde{b}_i(L^{1/m})$ are data-driven (e.g., selected from information criteria). Furthermore, since one cannot recover the high-frequency $g(L^{1/m})$ and $b_i(L^{1/m})$ polynomials, it is also not possible to identify the high-frequency impulse response function of $x_{i,t}$ on y_t . Hence, the approximate mixed-frequency model implies an observable quarterly impulse response function of $x_{i,t}^{(3)}$ on Y_t . Underlying this observable impulse response function are the latent monthly impulse response functions of $x_{i,t}$ on y_t .

1.3.2 From the approximate mixed-frequency regression to MIDAS regressions

Equation 1.10 is a MIDAS regression. More precisely, equation 1.10 as been referred to in the literature as an *unrestricted* MIDAS regression; see e.g. Marcellino and Schumacher (2010), Foroni and Marcellino (2012) and Foroni et al. (2011). This regression is one of the particular cases covered by the general MIDAS framework. Introduced by Ghysels et al. (2004) and initially presented in Ghysels et al. (2006) or Ghysels et al. (2007), the MIDAS approach provides simple, reduced-form models to approximate more elaborate, though unknown, high-frequency models. Original MIDAS regressions assume that the coefficients of $\tilde{b}_i(L^{1/m})$ in equation 1.10 are captured by a known weight function $B(j; \theta)$,

$$Y_{t+h} = \beta_0 + \beta_1 B(L^{1/m}; \theta) x_t^{(3)} + u_{t+h} \quad (1.11)$$

where $B(L^{1/m}; \theta) = \sum_{j=0}^J B(j; \theta) L^{j/m}$ is a polynomial of length J in the $L^{1/m}$ operator, $B(j; \theta)$ represents the weighting scheme used for the aggregation, which is assumed to be normalised to sum to 1, β_1 is the slope coefficient, β_0 is a constant, $L^{j/m} x_t^{(m)} = x_{t-j/m}^{(m)}$ and u_{t+h} is a standard iid error term. Similarly to distributed lag models (for further details on distributed lag models, see Dhrymes, 1972 and Griliches, 1967, among others), Ghysels et al. (2004) show that the aggregation bias also converges to zero in the MIDAS framework when $m \rightarrow 0$. Hence, regardless of the unavailability of Y_{t+h} at a higher frequency, using high-frequency regressors can mitigate the aggregation bias mentioned by Sims (1971) and Geweke (1978).

Although the order of the polynomial $B(L^{1/m}; \theta)$, i.e. J , is potentially infinite,

some restrictions must be imposed for the sake of tractability and, thus, the choice of J is an important step in the model specification (see, for example, Judge et al., 1985). Typically, information criteria, such as AIC and BIC, are used for model specification purposes. The restrictions result from the need of a balance between the gains in terms of additional information (more lags) and the costs of parameter proliferation. Nevertheless, in the original MIDAS regressions, the coefficients of $B(L^{1/m}; \theta)$ are captured by a known weight function $B(j; \theta)$, which depends on a few parameters summarised in vector θ . Thus, in this case, MIDAS regressions are tightly parameterised.

Regardless of the choice of a specific weight function (discussed in more detail in Section 1.3.2 below), MIDAS regressions allow for more flexible, data-driven weighting schemes than commonly used schemes, such as sum or equal-weight average (equations 1.4 and 1.5). For example, as Andreou et al. (2010) analytically demonstrate, traditional equal-weight aggregation is nested in MIDAS regressions, i.e.,

$$Y_{t+h} = \beta_0 + \beta_1 Q_t + \beta_1 \sum_{j=0}^J (B(j; \theta) - \frac{1}{J}) \Delta^{J-(j+1)} x_{t-j/m}^{(m)} + u_{t+h} \quad (1.12)$$

where $\Delta = 1 - L$ is the difference operator. Given that alternative aggregation functions of high-frequency data imply different restrictions on the $B(L^{1/m}; \theta)$ polynomial, the main question is which restrictions fit the data better. Focusing on equation 1.12, in case MIDAS regressions outperform the traditional approach, this implies that averaging may not be the aggregation method that maximises the ability of a high-frequency variable to predict a variable only available at a lower sampling frequency, meaning that $B(j; \theta) \neq (1/J)$ and, so, the coefficients $(B(j; \theta) - \frac{1}{J})$ do not cancel out.

Hence, MIDAS regressions allow for more flexible weighting structures than traditional low-frequency models and also tend to be more parsimonious. Moreover, the MIDAS framework can easily accommodate the timely releases of high-frequency data. In (1.11), where Y_{t+h} and $x_t^{(m)}$ are contemporaneously related, it is assumed that all high-frequency observations over the current low-frequency period are known. Considering quarterly and monthly data, this means that the

three months in the quarter are already available. If instead of a full-quarter, say, only the first month is available, then the MIDAS regression can be rewritten as

$$Y_{t+h} = \beta_0 + \beta_1 B(L^{1/3}; \theta) x_{t-2/3}^{(3)} + u_{t+h}. \quad (1.13)$$

Alternatively, the current-period information can be associated with a different β coefficient. Furthermore, MIDAS regressions can be extended to accommodate different (high) frequencies and additional indicators and, in most cases, without requiring many more parameters to be estimated. Thus, a more general version of MIDAS regressions, which can handle multiple regressors with different time frequencies, is as follows:

$$Y_{t+h} = \beta_0 + \sum_{n=1}^N \beta_n B_n(L^{1/m_n}; \theta) x_{t,n}^{(m_n)} + u_{t+h} \quad (1.14)$$

where N refers to the number of indicators and m_n to the time frequency of each indicator. Moreover, different polynomials $B_n(L^{1/m_n}; \theta)$ for each regressor can also be considered.

Although it is possible to have multivariate MIDAS regressions, the number of regressors included is necessarily limited. However, the MIDAS framework can still be used if one is interested in exploiting large information sets. Combining two important strands of the literature on forecasting using high frequency indicators, Marcellino and Schumacher (2010) merge factor models with MIDAS, introducing the Factor-MIDAS approach, which is applied to nowcast German GDP quarterly developments. Instead of focusing on a single regressor, or a small group of regressors, the authors extracted factors from a large set of variables, and compared the performance of different techniques for extracting factors in the presence of ragged-edge data. Then, the monthly factors were plugged into MIDAS regressions to estimate current quarter GDP. This approach is compared with an integrated state-space approach, as used in Altissimo et al. (2010), and with unrestricted MIDAS regressions. The results presented in Marcellino and Schumacher (2010) suggest that MIDAS regressions perform well when compared with state-space models.

Weighting schemes

Ghysels et al. (2007) considered two alternatives for the weighting function, both assuming that the weights are determined by a few hyperparameters θ : the exponential Almon lag

$$B(j; \theta_1, \theta_2) = \frac{e^{(\theta_1 j + \theta_2 j^2)}}{\sum_{i=1}^J e^{(\theta_1 i + \theta_2 i^2)}} \quad (1.15)$$

and the beta polynomial

$$B(j; \theta_1, \theta_2) = \frac{f(\frac{j}{J}, \theta_1, \theta_2)}{\sum_{i=1}^J f(\frac{i}{J}, \theta_1, \theta_2)} \quad (1.16)$$

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 $\Gamma(\theta) = \int_0^\infty e^{-k} k^{\theta-1} dk$. Given that exponential Almon and beta polynomials have nonlinear functional specifications, in both cases MIDAS regressions have to be estimated using nonlinear methods, namely nonlinear least squares.

Chen and Ghysels (2010) discussed a multiplicative MIDAS framework, which is closer to traditional aggregation. Instead of aggregating all lags in the high frequency variable to a single aggregate, multiplicative MIDAS regressions include m -aggregates of high-frequency data and their lags,

$$Y_{t+h} = \beta_0 + \sum_{i=1}^p \beta_{i+1} x_{t-i}^{mult} + u_{t+h} \quad (1.17)$$

where $x_t^{mult} = \sum_{j=0}^{m-1} B(j; \theta) L^{j/m} x_t^{(m)}$. In case many lags need to be included, a restricted version of the model can be used, where the coefficients associated with the m -aggregates and their lags (β_i) are also determined by a given polynomial, such as the exponential Almon or the beta polynomial. The multiplicative MIDAS model can also include current-period information, but in this case these data are necessarily associated with a different coefficient.

Bai et al. (2013) demonstrated that, in ideal circumstances (namely, in population and assuming correctly specified models), multiplicative MIDAS regressions can match exactly the steady state Kalman filter. In practice, when parameter

estimation and model specification errors also play a part, empirical exercises are needed in order to assess the forecasting performance of MIDAS vs. state space models. To forecast quarterly US GDP growth, Bai et al. (2013) showed that MIDAS regressions outperform the state space approach in short-term forecast horizons. Moreover, within MIDAS regressions, the original and multiplicative regressions have similar performances, neither one clearly dominating the other. Similarly, for euro area GDP, Kuzin et al. (2011) compared the forecasting performance of MIDAS regressions with mixed frequency VAR models, as proposed by Zadrozny (1988) and Mittnik and Zadrozny (2004), using a set of about 20 monthly indicators. These authors also concluded that MIDAS regressions perform better in shorter horizons, while state space models tend to perform better in longer horizons.

Inspired by the Heterogeneous Autoregressive models (HAR) in Corsi (2009), Forsberg and Ghysels (2007) considered an alternative weighting scheme, the so-called step functions. Instead of using a known function determined by a few parameters, step functions correspond to partial sums of the high-frequency variable $x_t^{(m)}$, as follows:

$$Y_{t+h} = \beta_0 + \sum_{s=1}^S \beta_s z_s^{(m)} + u_{t+h}$$

$$z_s^{(m)} = \mathbf{x}^{(m)} I_{j \in [a_{s-1}, a_s]}^s$$

$$I_{j \in [a_{s-1}, a_s]}^s = \begin{cases} 1, & a_{s-1} < j \leq a_s \\ 0, & \text{otherwise} \end{cases} \quad (1.18)$$

where $a_0 = 1 < a_1 < \dots < a_S = J$, S is the number of steps, with $S \leq J$, $I_{j \in [a_{s-1}, a_s]}^s$ is a $(J \times 1)$ vector (there are as many I^s vectors as the number of steps), and $\mathbf{x}^{(m)}$ is a matrix with the J lags of $x_t^{(m)}$. For example, assume that $\mathbf{x}^{(m)}$ represents 30 lags (a full month) of a daily series. Additionally, assume that these data are aggregated through step functions, using 3 steps - $[1 - 10]$, $[11 - 20]$ and $[21 - 30]$. Then, $s = 1, 2$ or 3 . For $s = 1$, I^s is a vector with 30 observations,

which equals 1 for $j = 1$ to 10 and 0 for $j = 11$ to 30. As there is a β_s coefficient for each step function, the more steps the less parsimonious the model will be. Although losing in terms of parsimony for the other two weighting schemes, the step functions are easier to estimate. In this case, ordinary least squares (OLS) can be used for estimating MIDAS regressions because step functions are linear.

There are two additional weighting schemes that also lead to regressions that can be estimated by OLS. The first one is the traditional Almon lag polynomial. This aggregation scheme assumes that J lag weights can be related to d linearly estimable underlying parameters, with $d < J$, as follows:

$$B(j; \theta_i) = \sum_{i=1}^d \theta_i j^i \quad (1.19)$$

where $j = 1, \dots, J$ and $i = 1, \dots, d$.

The second is the aggregation scheme underlying the above-mentioned unrestricted MIDAS regressions. The unrestricted MIDAS regressions are nested in the MIDAS regression with step functions, with a different ‘step’ for each high-frequency lag, such as

$$\begin{aligned} Y_{t+h} &= \beta_0 + B_u(L^{1/m})x_t^{(m)} + u_{t+h} \\ &= \beta_0 + \sum_{j=0}^J \beta_{j+1} L^{j/m} x_t^{(m)} + u_{t+h} \\ &= \beta_0 + \beta_1 x_t^{(m)} + \beta_2 x_{t-1/m}^{(m)} + \dots + \beta_{J+1} x_{t-J/m}^{(m)} + u_{t+h}. \end{aligned} \quad (1.20)$$

When the difference between the low and the high frequency is large, estimating a regression such as (1.20) can involve a large number of parameters. In this case, large differences in sampling frequencies between the variables considered are readily penalised in terms of parsimony. For instance, if Y_{t+h} is sampled quarterly and x_t refers to daily data, then estimating a mixed-frequency regression can easily involve estimating more than 60 parameters. The unrestricted MIDAS regressions are better able to deal with mixing quarterly and monthly data (as in Marcellino and Schumacher, 2010, for example) for which the parameter proliferation shortcoming is less stringent, than mixing quarterly with weekly, daily or intra-daily

data. Regarding its empirical performance, for nowcasting German GDP quarterly developments, Marcellino and Schumacher (2010) found that there are no big differences between alternative - original and unrestricted - MIDAS approaches. Using data for both the euro area and the US, Foroni et al. (2011) concluded that unrestricted MIDAS regression deliver a better in-sample performance than original MIDAS regressions when m is small. The results for the out-of-sample GDP forecasting exercise were not clear-cut and neither model globally dominated the other.

1.3.3 Autoregressive augmentation of MIDAS regressions

Originally, the MIDAS approach was associated with empirical applications to financial series, focusing on volatility predictions; see Ghysels et al. (2004, 2006, 2007). However, MIDAS regressions have also been used in macroeconomic applications. For example, Clements and Galvão (2008) and Bai et al. (2013) assessed whether the MIDAS framework is a good tool for forecasting output growth in the US. Moreover, Armesto et al. (2009) used MIDAS regressions to assess the predictive content of the Beige Book (US Federal Reserve System regional high-frequency information). Monteforte and Moretti (2013) used MIDAS regressions for forecasting monthly euro area inflation (current and one-month ahead) resorting to a monthly index of core inflation and daily data from financial markets.

The inclusion of autoregressive dynamics is an important feature of forecasting models, especially for macroeconomic applications. In this section the initial contributions to the literature regarding the introduction of autoregressive terms in MIDAS regressions are reviewed and, then, an alternative perspective on this subject is presented. For the sake of simplicity and unless otherwise stated, a first order autoregression and $h = 0$ are assumed. Notwithstanding, the notation can be easily extended.

Initial contributions

Ghysels et al. (2007) discussed the implications of introducing autoregressive terms in MIDAS regressions and suggested two possible ways of doing this, noting that

both solutions had caveats. In the first case, a polynomial in $L^{1/m}$ is considered,

$$Y_t = \beta_0 + \beta_1 B(L^{1/m}; \theta) x_t^{(m)} + \gamma Y_{t-(1/m)} + u_t \quad (1.21)$$

which implicitly assumes that $Y_{t-(1/m)}$ is available. This solution may not be very appealing because if high-frequency lags of Y_t were available then probably it would be possible to set up a high-frequency regression. Moreover, Ghysels et al. (2007) remarked that estimating (1.21) is more challenging because including the term $Y_{t-(1/m)}$ forces one to deal with endogenous regressors and with instrumental variable estimation.

In the second case, a polynomial in L is considered, i.e.

$$\begin{aligned} Y_t &= \beta_0 + \beta_1 B(L^{1/m}; \theta) x_t^{(m)} + \gamma Y_{t-1} + u_t \\ Y_t &= \beta_0^* + \beta_1 \frac{B(L^{1/m}; \theta)}{(1 - \gamma L)} x_t^{(m)} + u_t^* \end{aligned} \quad (1.22)$$

where $B(L^{1/m}; \theta) = \sum_{j=0}^J B(j; \theta) L^{j/m}$, $\beta_0^* = \beta_0 / (1 - \gamma)$ and $u_t^* = u_t / (1 - \gamma L)$. Considering the distributed lag representation, instead of a polynomial in $L^{1/m}$ one obtains a mixture $\frac{B(L^{1/m}; \theta)}{(1 - \gamma L)}$. Equation 1.23 shows more clearly the shape of the polynomial, for example assuming $J = m - 1 = 2$ for the sake of simplicity, we can see that

$$\begin{aligned} Y_t &= \beta_0^* + \beta_1 (B_0 x_t^{(3)} + B_1 x_{t-1/3}^{(3)} + B_2 x_{t-2/3}^{(3)} + \gamma B_0 x_{t-1}^{(3)} + \gamma B_1 x_{t-4/3}^{(3)} + \gamma B_2 x_{t-5/3}^{(3)} \\ &\quad + \gamma^2 B_0 x_{t-2}^{(3)} + \gamma^2 B_1 x_{t-7/3}^{(3)} + \gamma^2 B_2 x_{t-8/3}^{(3)} + \dots) + u_t^* \end{aligned} \quad (1.23)$$

Ghysels et al. (2007) and Andreou et al. (2011) pointed out that the autoregressive distributed lag MIDAS regression in equation 1.22 entails an undesirable property - the $\frac{B(L^{1/m}; \theta)}{(1 - \gamma L)}$ polynomial displays geometrically declining spikes at distance m , mimicking a seasonal pattern. This autoregressive augmentation of MIDAS regressions should be used if a seasonal pattern in $x_t^{(m)}$ is detected (Ghysels et al., 2007).

To illustrate these so-called spikes, Figure 1.1 plots simulated sequences of the coefficients associated with $x_t^{(3)}$ and its lags (up to 26 lags). The solid black line refers to equation 1.23, with $J = m - 1 = 2$, meaning that the high-frequency




Figure 1.1: Simulated sequences of the coefficients associated with $x_t^{(3)}$ and its lags

Note: In panels 1.1a and 1.1b the assumptions for the B_i coefficients are the following: if $J = 2$, $B_0 = 0.60$, $B_1 = 0.25$ and $B_2 = 0.15$; if $J = 3$, $B_0 = 0.45$, $B_1 = 0.30$, $B_2 = 0.15$ and $B_3 = 0.10$; finally, if $J = 5$, $B_0 = 0.25$, $B_1 = 0.40$, $B_2 = 0.20$, $B_3 = 0.075$, $B_4 = 0.05$ and $B_5 = 0.025$. In panels 1.1a and 1.1c $\gamma = 0.7$, while in panel 1.1b $\gamma = 0.2$. In all cases, $\beta_1 = 1$ and $\sum_{i=0}^J B_i = 1$.

terms cover a full low-frequency period ($m = 3$). The dashed black line represents the case where $J = 3$ and the solid grey line is for $J = 5$ (two low-frequency periods). Two alternatives for the autoregressive γ coefficient are considered: in panels 1.1a and 1.1c $\gamma = 0.7$, while in panel 1.1b $\gamma = 0.2$. Moreover, different weighting schemes are also assessed, as in Ghysels et al. (2007). Panels 1.1a and 1.1b display rapidly declining weights for $J = 2$, slower declining weights for $J = 3$ and hump-shaped weights for $J = 5$. Panel 1.1c plots the equal-weight case.

By looking at Figure 1.1 one can see that all three simple sets of weights that try to mimic the polynomial weighting schemes display a *spiky* pattern and, as expected, this pattern is softened with smaller γ . Note that with unrestricted MIDAS the pattern could be even more irregular, as the weights/coefficients are estimated unrestrictedly, not obeying to a known polynomial function. In the traditional case of equal-weight schemes, the sequence of coefficients have a step-wise pattern, except when the number of high-frequency terms do not fully cover low-frequency periods (e.g. when $J = 3$), also exhibiting spikes.

Clements and Galvão (2008) suggested an alternative way of introducing autoregressive dynamics in MIDAS regressions. The authors proposed interpreting the dynamics in Y_t as a common factor (Hendry and Mizon, 1978). This as-

assumption rests on the hypothesis that Y_t and $x_t^{(m)}$ share the same autoregressive dynamics, though, as Hendry and Mizon (1978) pointed out, a common factor may not always be found. Hence, considering

$$\begin{aligned} Y_t &= \beta_0 + \beta_1 B(L^{1/m}; \theta) x_t^{(m)} + u_t \\ u_t &= \gamma u_{t-1} + \varepsilon_t \end{aligned} \quad (1.24)$$

and replacing u_t in (1.24) it follows that

$$(1 - \gamma L)Y_t = \beta_0(1 - \gamma) + \beta_1(1 - \gamma L)B(L^{1/m}; \theta)x_t^{(m)} + \varepsilon_t \quad (1.25)$$

An analogous equation for h higher than 0 can be written as

$$(1 - \gamma L^h)Y_{t+h} = \beta_0(1 - \gamma) + \beta_1(1 - \gamma L^h)B(L^{1/m}; \theta)x_t^{(m)} + \varepsilon_{t+h} \quad (1.26)$$

or, in case the forecaster decides to take full advantage from the fact that high-frequency data are released sooner than the low frequency series, then equation 1.26 can be written as

$$(1 - \gamma L^{h+1})Y_{t+h} = \beta_0(1 - \gamma) + \beta_1(1 - \gamma L^{h+1})B(L^{1/m}; \theta)x_t^{(m)} + \varepsilon_{t+h} \quad (1.27)$$

and can be estimated before Y_t is released, but full information of x_t for the current quarter is already available.

When writing (1.25), (1.26) or (1.27) in the distributed lag representation, the polynomial in L cancels out, leaving the polynomial in $L^{1/m}$. The coefficients β_0 , β_1 , θ and γ are estimated together, using nonlinear least squares in a two-step approach. First, the initial value for γ can be estimated as

$$\hat{u}_t = \gamma \hat{u}_{t-1} + \varepsilon_t \quad (1.28)$$

where \hat{u}_t are the residuals from a static MIDAS regression (without autoregressive terms) and ε_t is an error term iid $(0, \sigma^2)$. The estimates for γ are used to filter the original Y_t and $x_t^{(m)}$ variables. Second, the filtered variables are used to estimate β_0 , β_1 and θ in a static MIDAS regression. Although the initial work

by Clements and Galvão (2008) and subsequent empirical applications (see, for example, Marcellino and Schumacher, 2010, Kuzin et al., 2011, Foroni and Marcellino, 2012, Jansen et al., 2012, Monteforte and Moretti, 2013 and Kuzin et al., 2013) only consider a single autoregressive term, in this paper this technique was extended to allow for more autoregressive terms. This can be done by including the additional lags of the residuals \hat{u}_t in (1.28), as follows

$$\hat{u}_t = \sum_{i=1}^p \gamma_i \hat{u}_{t-i} + \varepsilon_t. \quad (1.29)$$

Alternative perspective

The literature presented so far suggests that adding autoregressive terms to MIDAS regressions in such a way that generates a $\frac{B(L^{1/m}; \theta)}{(1-\gamma L)}$ polynomial implies declining spikes at distance m in the infinite sequence of coefficients associated with $x_t^{(m)}$ and its lags. However, should this sequence be perceived as the relevant impulse response function from the high-frequency variable on the low-frequency variable?

To answer this question let us start by assuming that x_t and y_t are both observed at the high frequency, say monthly. Furthermore, assume that the DGP is known and resumes to an autoregressive term of order one, the contemporaneous variable x_t and two lags ($x_{t-1/3}$ and $x_{t-2/3}$). To ease the exposition, the time index remains unaltered, so that the monthly time index is expressed by $t = 0, 1/3, 2/3, 1, 4/3, \dots, T$.

$$y_t = \alpha + b_0 x_t + b_1 x_{t-1/3} + b_2 x_{t-2/3} + \lambda y_{t-1/3} + u_t \quad (1.30)$$

The distributed lag version of this equation can be written as

$$y_t = \alpha^* + b_0 x_t + (b_1 + \lambda b_0) x_{t-1/3} + \sum_{i=0}^{\infty} \lambda^i (b_2 + \lambda b_1 + \lambda^2 b_0) x_{t-(2+i)/3} + u_t^* \quad (1.31)$$

where $\alpha^* = \alpha / (1 - \lambda L^{1/3})$ and $u_t^* = u_t / (1 - \lambda L^{1/3})$. Consider a shock equal to 1 in x in January 2012. The response on the y variable in January 2012 is equal to b_0 . The responses in February and March are $b_1 + \lambda b_0$ and $b_2 + \lambda b_1 + \lambda^2 b_0$, respectively. So, considering a generic aggregation scheme $\Gamma(L^{1/m})$ as in (1.7), the

response in the first quarter is equal to $b_0\gamma_0 + (b_1 + \lambda b_0)\gamma_1 + (b_2 + \lambda b_1 + \lambda^2 b_0)\gamma_2$. Similarly, the monthly responses in April, May and June are $\lambda(b_2 + \lambda b_1 + \lambda^2 b_0)$, $\lambda^2(b_2 + \lambda b_1 + \lambda^2 b_0)$ and $\lambda^3(b_2 + \lambda b_1 + \lambda^2 b_0)$, respectively, which results in a response of $(\lambda^3\beta_0 + \lambda^2\beta_1 + \lambda\beta_2)(\gamma_0 + \gamma_1\lambda + \gamma_2\lambda^2)$ in the second quarter.

The exact mixed-frequency regression can be obtained by pre-multiplying (1.30) by the finite order polynomials $\phi(L^{1/3})$ and $\Gamma(L^{1/3})$. Assume that the autoregressive order in the low/quarterly frequency is also one. In this case, $\phi(L^{1/3}) = (1 + \lambda L^{1/3} + \lambda^2 L^{2/3})$, i.e., $\phi(L^{1/3})(1 - \lambda L^{1/3}) = 1 - \lambda^3 L$. So, applying both polynomials one obtains the following mixed-frequency regression

$$(1 - \lambda^3 L)Y_t = \hat{\alpha} + \delta_0 x_t + \delta_1 x_{t-1/3} + \delta_2 x_{t-2/3} + \delta_3 x_{t-1} + \delta_4 x_{t-4/3} + \delta_5 x_{t-5/3} + \delta_6 x_{t-2} + \hat{u}_t \quad (1.32)$$

where $\delta_0 = b_0\gamma_2$, $\delta_1 = b_0(\gamma_1 + \lambda\gamma_2) + b_1\gamma_2$, $\delta_2 = b_0(\gamma_0 + \lambda\gamma_1 + \lambda^2\gamma_2) + b_1(\gamma_1 + \lambda\gamma_2) + b_2\gamma_2$, $\delta_3 = b_0(\lambda\gamma_0 + \lambda^2\gamma_1) + b_1(\gamma_0 + \lambda\gamma_1 + \lambda^2\gamma_2) + b_2(\gamma_1 + \lambda\gamma_2)$, $\delta_4 = b_0\lambda^2\gamma_0 + b_1(\lambda\gamma_0 + \lambda^2\gamma_1) + b_2(\gamma_0 + \lambda\gamma_1 + \lambda^2\gamma_2)$, $\delta_5 = b_1\lambda^2\gamma_0 + b_2(\lambda\gamma_0 + \lambda^2\gamma_1)$ and $\delta_6 = b_2\lambda^2\gamma_0$.

Writing this equation in a distributed lag form one obtains

$$\begin{aligned} Y_t = \tilde{\alpha} + & \delta_0 x_t + \delta_1 x_{t-1/3} + \delta_2 x_{t-2/3} \\ & + (\delta_3 + \lambda^3 \delta_0) x_{t-1} + (\delta_4 + \lambda^3 \delta_1) x_{t-4/3} + (\delta_5 + \lambda^3 \delta_2) x_{t-5/3} \\ & + (\delta_6 + \lambda^3 \delta_3 + \lambda^6 \delta_0) x_{t-2} + \lambda^3 (\delta_4 + \lambda^3 \delta_1) x_{t-7/3} + \lambda^3 (\delta_5 + \lambda^3 \delta_2) x_{t-8/3} \\ & + \lambda^3 (\delta_6 + \lambda^3 \delta_3 + \lambda^6 \delta_0) x_{t-3} + \lambda^6 (\delta_4 + \lambda^3 \delta_1) x_{t-10/3} + \lambda^6 (\delta_5 + \lambda^3 \delta_2) x_{t-11/3} \\ & + \lambda^6 (\delta_6 + \lambda^3 \delta_3 + \lambda^6 \delta_0) x_{t-4} + \lambda^9 (\delta_4 + \lambda^3 \delta_1) x_{t-13/3} + \lambda^9 (\delta_5 + \lambda^3 \delta_2) x_{t-14/3} + \dots + \tilde{u}_t \end{aligned} \quad (1.33)$$

where $\tilde{\alpha} = \hat{\alpha}/(1 - \lambda^3 L)$ and $\tilde{u}_t = \hat{u}_t/(1 - \lambda^3 L)$. Again, consider a shock equal to 1 in x in January 2012. The response on the Y variable in the first quarter of 2012 is equal to δ_2 , which corresponds to $b_0(\gamma_0 + \lambda\gamma_1 + \lambda^2\gamma_2) + b_1(\gamma_1 + \lambda\gamma_2) + b_2\gamma_2$. Rearranging the terms, this response equals the quarterly aggregate response underlying the monthly regression. A similar result is obtained for the following period - the response on the Y variable in the second quarter of 2012 is $\delta_5 + \lambda^3\delta_2$, which also equals the quarterly aggregate of the monthly responses in April, May

and June. The same reasoning is also valid if shocks in other months or combined shocks in more than one month were considered. These results can be extended to different specifications and to other forecast horizons.

As regards impulse response functions from $x_t^{(m)}$ on Y_t , sequentially assessing the coefficients in (1.33) does not seem to be very informative. The sequence $\delta_0, \delta_1, \delta_2, (\delta_3 + \lambda^3\delta_0), \dots$ - with or without spikes - cannot be considered the relevant impulse response function because some of the coefficients (in this case, sets of three non-overlapping parameters) refer to the same time period in the low frequency, i.e., to the same quarter. Furthermore, recall that each coefficient in this sequence already covers the relevant latent monthly impulse responses on y within each quarterly observation of Y , for each monthly shock in x .

Inspired by the periodic model framework (see Hansen and Sargent (2013) and Ghysels (2012), among others), one can say that there are several impulse response functions, one for each m high-frequency period within a low-frequency observation. For example, in (1.33), the observed quarterly impulse response function from a shock in the first month of the quarter on the quarterly Y_t variable is $\delta_2, \delta_5 + \lambda^3\delta_2, \lambda^3(\delta_5 + \lambda^3\delta_2), \lambda^6(\delta_5 + \lambda^3\delta_2)$, and so on and so forth. Similarly, the observed quarterly impulse response function from a shock in the second month of the quarter on the quarterly Y_t variable is $\delta_1, \delta_4 + \lambda^3\delta_1, \lambda^3(\delta_4 + \lambda^3\delta_1), \lambda^6(\delta_4 + \lambda^3\delta_1), \dots$, and so on. Any of the observed quarterly impulse response functions do not exhibit spikes, regardless of the lags of $x_t^{(m)}$ or Y_t included in the regression.

When the number of high-frequency lags in the mixed-frequency regressions is multiple of m minus 1 (or lower than $m - 1$) there is homogeneity in the shape of the impulse response functions on Y_t , regardless of the type of shock in $x_t^{(m)}$. This means that all impulse response functions share the same geometric decay pattern, i.e., the decay starts at the same time. In cases where the number of high-frequency lags is greater than $m - 1$ but not its multiple, such as equation 1.32, the shape of the impulse response functions varies with the high-frequency timing of the shock in $x_t^{(m)}$ - for shocks in the first and second months of the quarter the geometric decay starts after two quarters, while for shocks in the third month of the quarter that decay only starts after three quarters.

Note that, as mentioned before, in a single-equation environment, one cannot

recover the monthly impulse response functions of x_t on y_t departing from the mixed-frequency regression. Moreover, this framework only analyses the transmission of changes in one direction, from the high-frequency variables to the low-frequency variable, not taking into account the possible relation between the high-frequency variables nor the impact of changes in the low-frequency variable into the high-frequency variables.

In light of this discussion, alternatives to the common factor way of introducing autoregressive terms in MIDAS regressions can be considered. In particular, generalizing conventional ADL regressions, autoregressive terms are added to MIDAS regressions, without restrictions - not imposing the common factor, no restrictions on both the lag structure and the order of the autoregressive polynomial - see Andreou et al. (2013) and Guérin and Marcellino (2013). Moreover, no restrictions are imposed on the aggregation scheme - exponential Almon weight function, unrestricted (Forni and Marcellino, 2012) and multiplicative (Francis et al., 2011).

Furthermore, in the same vein of multiplicative MIDAS regressions, which closely map the traditional low frequency model (i.e., reverse engineering the low frequency regressions, replacing the time aggregates by the underlying combinations of high-frequency lags, results in a mixed-frequency regression with a number of high-frequency lags that is always a multiple of m minus 1) the performance of original and unrestricted MIDAS regressions with autoregressive terms and with high-frequency lags multiple of m minus 1 is also analysed.

The latter regressions require full quarter information to be available, including for the current quarter. In order to implement these regressions when the m current-period high-frequency observations have not been released, the series of the regressors with unbalanced m periods were stacked with forecasts obtained from simple autoregressive regressions. Given that MIDAS regressions deliver direct forecasts, the autoregressive extrapolation of regressors is also based on direct forecasting. Note that this *bridge* approach to MIDAS regressions can be easily implemented for, say, monthly regressors. However, this procedure is less feasible when the regressors have time frequencies higher than monthly. This approach somehow mimics the state-space approach, with the simple autoregressive regressions acting as the state equation, and the MIDAS regression as the observation

equation. Nevertheless, as in the traditional bridge model framework, this two-step approach is hindered by orthogonality issues, which can lead to biases in coefficient estimates.

Note that this latter version of MIDAS regressions ensures that the geometric decay pattern starts at the same time in all m impulse response functions. Moreover, when new information becomes available within the current quarter, there is no need to change the forecasting regression in order to update the forecasts. This procedure only involves substituting the stacked regressor forecasts by the newly observed figures.

1.4 Design of the nowcasting and forecasting exercise

The aim of this exercise is to nowcast and forecast quarterly developments in euro area GDP growth, in real terms, using three different indicators: a hard-data series; a soft-data series; and a financial series. Thus, the dataset used contains a quarterly series on the real GDP from 1996Q1 to 2012Q4, the monthly industrial production and the monthly economic sentiment indicator from January 1996 to December 2012, and the Dow Jones Euro Stoxx index on a daily basis from 1 January 1996 to 31 December 2012. All series are seasonally adjusted except the stock market information. Apart from the economic sentiment indicator, the original series were transformed, using the rate of change (based on the first difference), in order to have stationary variables.

The data considered are final data, meaning that they refer to the latest release available when the database was built. While in the case of the economic sentiment indicator and the stock market index final data equal real-time data, as these series are not revised, revisions of GDP and industrial production are not taken into account in this analysis. However, evidence from previous work on data revisions suggests that revisions are typically small for euro area GDP (Marcellino and Musso, 2011).

The database is split in two, for the in-sample estimation and the out-of-sample forecasting exercise. From 1996Q1 to 2006Q4 the sample was used for in-sample

model specification and estimation. Different types of MIDAS regressions were estimated - original, multiplicative, with and without autoregressive terms - based on different information sets - single-variable models, multi-variable models with only monthly data and mixing monthly and daily data.¹ Different lags were considered (up to 4 quarters), also for the autoregressive terms. All regressions were recursively estimated and selected using information criteria, namely the BIC.

In the following analysis the original MIDAS regressions, as in equation 1.11, will be simply denoted as “MIDAS”. Moreover, the multiplicative (equation 1.17) and unrestricted (equation 1.20) regressions are denoted as “M-MIDAS” and “U-MIDAS”, respectively. The MIDAS specification with common factor autoregressive dynamics (equation 1.25) will be labelled “CF-MIDAS”, while without that restriction the prefix “AR-” is added. The case of MIDAS regressions with autoregressive terms and with high-frequency lags multiple of m minus 1 will have the prefix “Balanced”. Apart from U-MIDAS regressions, all other MIDAS regressions were estimated using the exponential Almon polynomial defined as in equation 1.15.² Different initial parameter specifications (including the equal weight hypothesis, i.e. $\theta_1 = \theta_2 = 0$) were tested and the results do not differ significantly (for a discussion on the shapes of different weighting sets, see Ghysels et al., 2007). The hyperparameters θ of the exponential Almon function are restricted to $\theta_1 < 5$ and $\theta_2 < 0$.

The sample from 2007Q1 to 2012Q4 was used for the out-of-sample nowcasting and forecasting exercise. Although an out-of-sample forecast exercise with $P = 24$ quarters has limitations, using euro area data still bears an inevitable trade-off between sample sizes for in-sample and out-of-sample exercises and this forecasting exercise is no exception. For obtaining the forecasts, a recursive exercise was performed, so that throughout the out-of-sample period the estimation sample is recursively expanded by adding one observation at a time. As a new observation is added to the estimation sample, the regressions are re-estimated and, thus, the

¹The codes used to estimate and forecast using MIDAS regressions were written in Matlab. Some functions were taken from the Econometrics Toolbox written by James P. LeSage (<http://www.spatial-econometrics.com>). The MIDAS toolbox used was greatly inspired in a code kindly provided by Arthur Sinko.

²The beta polynomial was also tested and the results were qualitatively similar. All results are available from the author upon request.

coefficients are allowed to change over time. Adding to nowcasts ($h = 0$) direct forecast for up to $h = 4$ quarters ahead are also presented. For each forecast horizon a different model is estimated.

Although the database used is not a real-time database, the different publication lags of the indicators are taken into account when within-quarter information is used. This pseudo real-time exercise relies on mimicking the release pattern of the high-frequency indicators. So, considering single-variable models, it would be possible to have up to 3 different forecasts for a given quarter, for each quarterly forecast horizon, depending on the within-quarter information used - one month (*I*), two months (*II*), or full quarter (*III*).

Furthermore, when combining the different variables, from the end of the first month of quarter t to the end of the first month of quarter $t + 1$ it is possible to have up to 4 different forecasts for quarter t , depending on the information set and the amount of within-quarter information that is taken on board. For example, assume that one is interested in forecasting the euro area GDP growth for the fourth quarter of 2010 (recall that the first release will only be available in mid-February). In the end of October 2010 the economic sentiment indicator and the stock market returns will be available for the month of October (*I*).³ A month later, in the end of November, the economic sentiment indicator and the stock market returns will be available for the month of November and the industrial production will be available for the month of October (*II*). Again, a month later, in the end of December, the economic sentiment indicator and the stock market returns will be available for the month of December and the industrial production will be available for the month of November (*III*). Information on the industrial production for the month of December will be available in the end of January 2011. Hence, by the end of January 2011 full-quarter information is available for all the variables considered (*IV*). If using full-quarter information allows having forecasts for the fourth quarter of 2010 only a couple weeks before the first release of GDP figures, it is possible to have forecasts that exploit partial within-quarter information much earlier than that. However, note that by using MIDAS framework a new regression is needed as new information is released.

³To simplify the description, monthly intervals are also considered for the daily series.

To evaluate the forecasting performance of the different MIDAS regressions is used the root mean squared forecast error (RMSFE). Relative RMSFE are computed to compare the performance of the MIDAS approach with alternative, purely quarterly, benchmark models. Two benchmark models are considered. The first is an autoregressive (AR) model, which is estimated recursively, using a general-to-specific approach, and the lag length (from 0 to 4 lags) is chosen according to information criteria, namely the BIC. The AR benchmark boils down to the sample average when, according to the BIC criterion, including positive lags leads to a worse performance than choosing the lag length equal to 0.

The second is a traditional quarterly single-equation multivariate model, with all the variables in the low frequency. This model includes autoregressive terms (from 0 to 4 lags) and is also estimated recursively, using a general-to-specific approach and the BIC. As different information sets are considered (single-variable or multi-variable) the quarterly multivariate models are adjusted accordingly. Moreover, when full quarter information is not available, forecasts from the quarterly multivariate models are obtained through a bridge model framework, in this case with a direct forecasting approach, similarly to MIDAS approach. So, estimates for the missing monthly observations, obtained from univariate models, are plugged in the monthly data, which are transformed into quarterly series and, then, used for forecasting in the traditional quarterly model, as the one in equation 1.2. To ensure consistency within all forecasts used, the missing monthly observations are also direct forecasts from autoregressive models.

In order to assess the statistical significance of the differences in the forecasting performance between the alternatives considered, the test of equal forecast accuracy on the population-level of direct multi-step forecasts from nested linear models proposed by Clark and McCracken (2005) is used. The null hypothesis is that the benchmark model forecasts (restricted model, denoted as model 1) are as accurate as those of the MIDAS regressions (unrestricted model, denoted as model 2) and the one-sided alternative hypothesis is that the unrestricted model forecasts are more accurate. Following the authors' notation, the test statistic used is

$$MSE - F = P \frac{MSE_1 - MSE_2}{MSE_2} \quad (1.34)$$

where P is the number of forecasts and MSE_i denotes the mean squared forecast error of model i , with $i = 1, 2$. Because this test has a non-standard limiting distribution, a bootstrap procedure was implemented to obtain the critical values. As suggested by Clark and McCracken (2005) and similarly to Kilian (1999), the bootstrap algorithm used starts with the estimation of a large set of simulated samples of the dependent variable (1000 samples), which are computed by drawing with replacement from the sample residuals under the null hypothesis (restricted model). Moreover, a bootstrap-after-bootstrap procedure, as proposed by Kilian (1998), is implemented to obtain small-sample bias-adjusted bootstrapped time series. Based on this simulated data, both restricted and unrestricted direct multi-step forecasts are calculated recursively and the test statistic is computed for each set of forecasts. The critical values are computed as quantiles of the bootstrapped series of test statistics.

1.5 Empirical results

In this section, the empirical results are presented, starting with the results for single-variable models (Section 1.5.1). Then, the role of forecast pooling as a tool for improving forecast accuracy is discussed in Section 1.5.2. Finally, Section 1.5.3 presents the results for multi-variable models.

1.5.1 Single-variable models

Figures 1.2 and 1.3 summarise the results on the forecasting performance of different MIDAS regressions against an AR and traditional low-frequency quarterly benchmark (more detailed results, as well as the significance levels for comparing forecast accuracy, can be found in Tables A.1 and A.2, in the Appendix). The figures reported refer to the relative RMSFE, so figures lesser than one mean that the forecasting performance of the MIDAS model is better, in terms of RMSFE, than the benchmark model - the naive AR or the traditional quarterly model, respectively.

Comparing with an AR benchmark, MIDAS regressions have, in general, a better performance, up to 1 quarter ahead. The relative RMSFE is less than

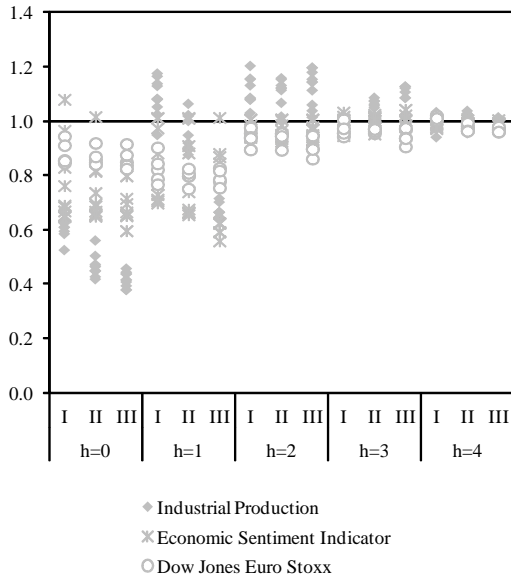


Figure 1.2: Forecast performance of MIDAS single-variable regressions compared to an AR benchmark, relative RMSFE

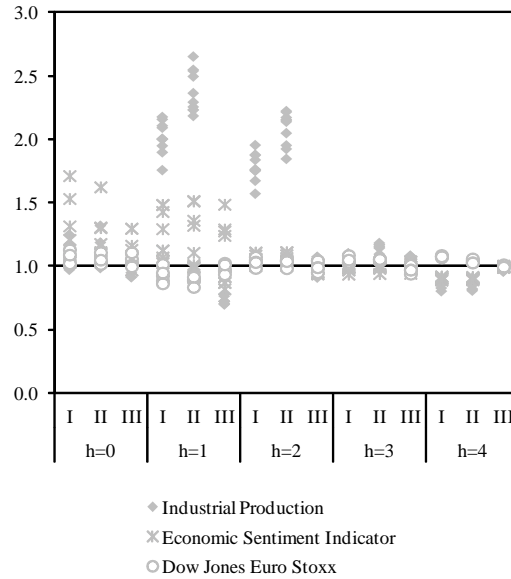


Figure 1.3: Forecast performance of MIDAS single-variable regressions compared to a quarterly benchmark, relative RMSFE

1 in about 78 per cent of the cases for the industrial production, about 88 per cent for the economic sentiment indicator and 100.0 per cent for the Dow Jones Euro Stoxx daily index. On average, the gains from using MIDAS regressions vary from around 16 per cent for the Dow Jones Euro Stoxx index to about 28 per cent for the economic sentiment indicator, reaching 40 per cent in the case of industrial production. From 2 quarters ahead onwards, though there is also evidence in favour of MIDAS, this evidence is less striking. There are less cases where the relative RMSFE is less than 1 and the average gains in those cases are smaller. Nevertheless, the best performing MIDAS regressions clearly outperform the benchmark, as the differences in terms of forecast accuracy are statistically significant.

In some cases, especially for forecasts up to 1 quarter ahead, traditional quarterly models can be seen as a more challenging benchmark than a simple AR model. Considering these quarterly models as benchmark, the results are not clear cut; some MIDAS specifications, for some horizons, have better forecasting performances than the benchmark, while others show a worse performance. Although




Figure 1.4: Best performing single-variable regression, for each forecast horizon, relative RMSFE compared to AR benchmark

Note: For each forecast horizon, the MIDAS figures correspond to the minimum of the respective columns in Table A.1.

MIDAS regressions, as a whole, are not superior to the quarterly benchmark and the average gain (when exists) is smaller than in the case of the AR benchmark, the best performing MIDAS regressions deliver better results than the respective quarterly models and the differences in terms of RMSFE are, in general, statistically significant. So, it can be concluded that exploiting high-frequency data has a significant impact on forecasting performance.

Figure 1.4 summarises the results, showing the relative RMSFE against an AR benchmark for the best performing MIDAS regressions and traditional quarterly models. Note that the figures plotted, for each forecast horizon, correspond to the minimum relative RMSFE, which in the case of MIDAS regressions can be found in the respective columns of Table A.1. Over the forecast horizons, the forecasting performance of MIDAS regressions is slightly better than the one of the traditional quarterly models, especially in shorter forecast horizons - up to 2 quarters ahead, the average gain from MIDAS regressions is 34.6 per cent, which compares with

31.8 per cent for the traditional quarterly models.

So, as in Clements and Galvão (2008), Clements and Galvão (2009) and Marcellino and Schumacher (2010), among others, results presented so far suggest that using MIDAS regressions contributes to increase forecast accuracy in terms of RMSFE. So, let us look into the best performing regressions in more detail. Starting with the choice of the regressors, the results of the single-variable models with the best forecasting performance compared to the autoregressive benchmark suggest that (i) for nowcasting the regressions that include the industrial production index have the best performance, (ii) the economic sentiment indicator delivers the best results in the case of one quarter ahead forecasts, and (iii) for longer forecast horizons the best results are, in general, associated with the Dow Jones Euro Stoxx index (Table A.1 in the Appendix). This pattern is not surprising, given that is typically found a strong contemporaneous correlation between GDP and the industrial production and confidence indicators and stock market indices tend to exhibit forward looking properties. Perhaps, more surprising is the fact that the strictly highest forecast accuracy for four quarters ahead is obtained from a regression that uses the industrial production index.

Now, instead of using a benchmark that does not take into account covariates, let us use traditional quarterly single-variable regressions as benchmark (Table A.2 in the Appendix). Over the forecast horizons considered, no clear pattern can be detected regarding the choice of the regressors. Overall, the best forecasting results are concentrated on the regressions that use the industrial production index and the Dow Jones Euro Stoxx index. On the one hand, this stresses the fact that the best results are not always obtained from using the same regressor. On the other hand, this evidence suggests that the use of MIDAS parsimonious, data-driven weighting schemes to aggregate the high-frequency data is advantageous for forecasting over forecast horizons up to 4 quarters ahead, using either monthly or daily data. Moreover, MIDAS regressions with the highest forecast accuracy also show a good performance when incomplete information for the current quarter is used, beating the results from traditional quarterly model that rely on monthly direct forecasts to construct missing quarterly observations (bridge model framework). Hence, MIDAS seems to be a good and simple tool for using within-quarter




Figure 1.5: Single-variable model - Dow Jones Euro Stoxx index - for each forecast horizon, relative RMSFE compared to AR benchmark

Note: The figures plotted correspond to the minimum of the respective columns in Table A.3.

high-frequency information in order to improve forecast accuracy.

As MIDAS regressions seem to be able to deal with different frequencies within high-frequency data, Figure 1.5 tries to provide further insight on whether the exact frequency of the data used plays a role on improving the forecast accuracy (additional information can be found in Table A.3 in the Appendix). This figure compares the relative RMSFE against an AR benchmark of single-variable models using the Dow Jones Euro Stoxx index in daily, weekly and monthly frequencies. The results are fairly similar over the forecast horizons considered, suggesting that the flexibility of MIDAS lag distribution, rather than the exact time frequency of the series, seems to be responsible for the good forecasting performance of MIDAS regressions.

Regarding the choice within MIDAS regressions, also in this case the best results are not always obtained from the same type of MIDAS weighting scheme. Figures 1.6 and 1.7 plot the relative RMSFE for the several MIDAS regressions, with a different regressor in each panel - panels (a), (b) and (c) display the results

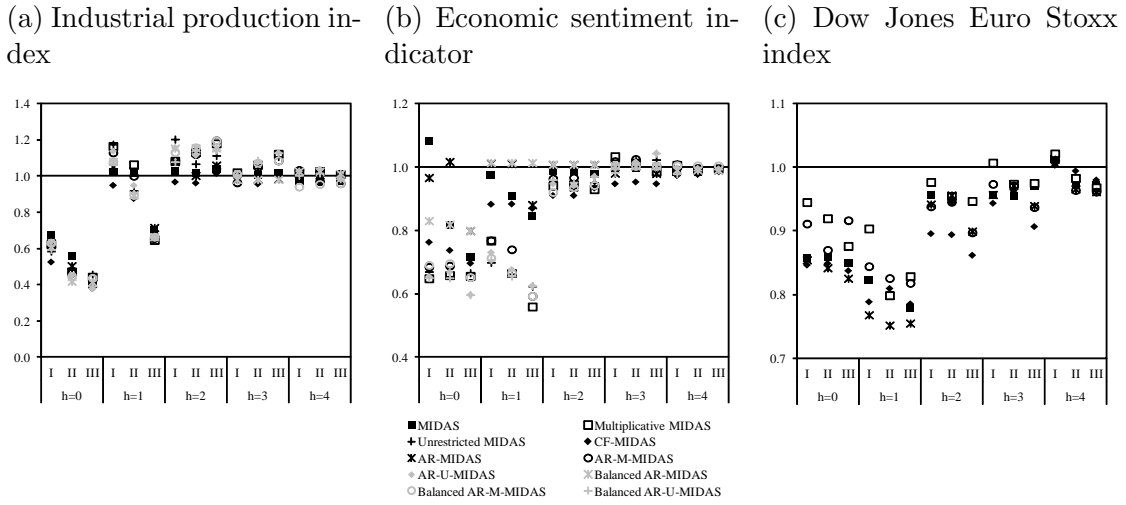


Figure 1.6: Relative RMSFE of MIDAS regressions compared to an AR benchmark

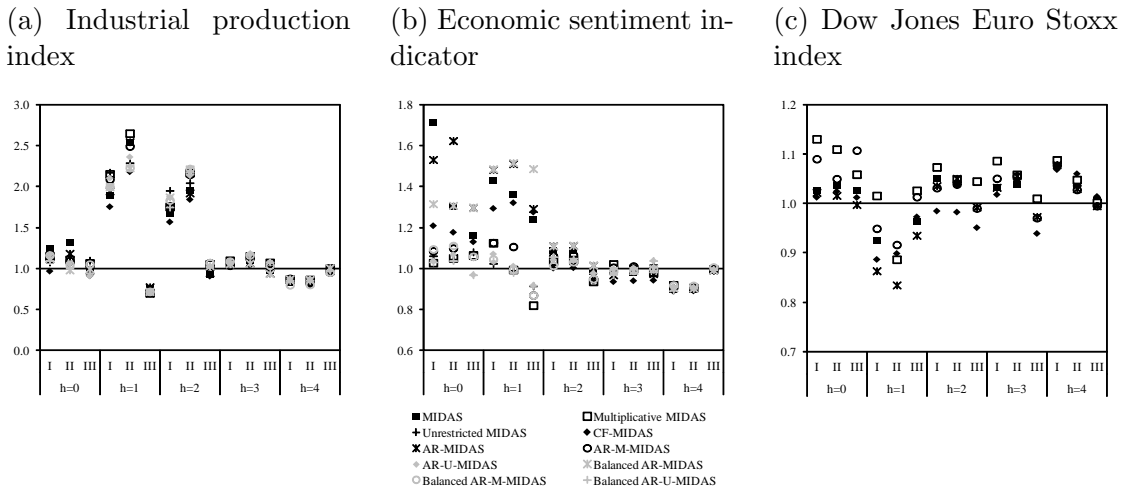


Figure 1.7: Relative RMSFE of MIDAS regressions compared to a quarterly benchmark

of industrial production, economic sentiment indicator and stock market index, respectively. There are five main conclusions that can be drawn from these results. First, as the forecast horizon increases, the differences in the forecasting performance between MIDAS regressions decrease, making the choice among alternative MIDAS weighting schemes less relevant. In contrast, the differences are higher for short-term forecasting, rendering this choice crucial for achieving the best performance.

Second, less parsimonious MIDAS weighting schemes - multiplicative and un-

restricted MIDAS - are often less accurate, in terms of RMSFE, than polynomial schemes. One exception are the short-term forecasts (up to 2 quarters ahead) using the economic sentiment indicator. In this case, multiplicative and unrestricted MIDAS regressions, with or without autoregressive terms display the best forecasting performances, either in comparison with the AR or the quarterly benchmark. This performance may be explained by the fact that the restrictions on the signs of the coefficients are less stringent in multiplicative or unrestricted MIDAS. As noted by Breitung et al. (2012), using exponential (or beta) polynomials to determine the shape of the lag distribution imposes that all $B(j; \theta_1, \theta_2)$ lag coefficients share their sign. Thus, in polynomial MIDAS regressions the sign of the relation between Y_t and $x_t^{(m)}$ is determined by the sign of the β_1 coefficient. On the contrary, the multiplicative scheme allows different signs on the β_i coefficients for each lag of the m -aggregates of the x variables. Similarly, in the unrestricted regressions each high-frequency lag has its own coefficient.⁴ The estimation results from the traditional quarterly models confirm that changing signs in the lag coefficients is an important feature in the regressions using the economic sentiment indicator, for forecast horizons up to 2 quarter ahead.

Third, the best performing MIDAS regressions tend to include autoregressive terms, which is an expected result given that this empirical application uses macroeconomic data (Clements and Galvão, 2008, Marcellino and Schumacher, 2010, Monteforte and Moretti, 2013, among others).

Fourth, focusing on MIDAS regressions with autoregressive terms, the results suggest that it is possible to improve forecasting performance of MIDAS regressions by using weighting schemes alternative to CF-MIDAS. In particular, up to 1 quarter ahead, balanced AR-U-MIDAS consistently outperforms CF-MIDAS in the regressions using the economic sentiment indicator. Note that this performance is observed even when full-quarter information is not available, which suggests that combining balanced MIDAS regressions with autoregressive extrapolation of the regressor can deliver good results in terms of forecast accuracy in the short term. Also for short-term forecasting, AR-MIDAS regressions with the Dow Jones Euro Stoxx index have the best forecasting performance, being a good alternative for

⁴Using the traditional Almon polynomial, instead of exponential or beta polynomials, is another alternative to eliminate the restrictions on the signs of the coefficients.

dealing with autoregressive augmentation. In the case of industrial production, the evidence is more mixed and no clear pattern is detected. Nevertheless, in 7 out of 15 cases the alternative models to CF-MIDAS show the best performance.

Finally, the existence of some degree of variability in the ranking of forecasting performance among alternative MIDAS regressions with autoregressive terms, especially in the short term, suggests that choosing is essentially an empirical question. It may be the case that in some empirical exercises imposing a common factor dynamics between Y_t and $x_t^{(m)}$ can be less benign than using alternative ways of including autoregressive terms in MIDAS regressions. In other cases it may be the opposite.

1.5.2 Pooling forecasts

As discussed in the previous section, the use of the different MIDAS approaches, in practice, requires several decisions regarding the specification and the information set. An adequate selection of the regressors used is crucial for obtaining the best results in terms of forecast accuracy, over the forecast horizons considered. In the case of an information set comprising a small set of selected variables, two different strategies may help overcome this variable selection problem: pooling forecasts and multi-variable models (see section below). Pooling forecasts is typically seen in the literature - also in MIDAS literature, for instance in Kuzin et al. (2011) and Andreou et al. (2013) - as a simple technique for dealing with model uncertainty.

Nowadays, different pooling techniques are available, ranging from simple equal (and constant) weights to performance based weights. Granger and Ramanathan (1984) found evidence in favour of combining forecasts by using simple techniques, as the mean. As simple combination schemes often show good performances, in this paper two different pooling techniques are used, namely the equal-weight mean and, following Stock and Watson (2004), the discounted mean squared forecast error (MSFE) combination, whose weights are as follows:

$$w_{it} = \frac{m_{it}^{-1}}{\sum_{i=1}^n m_{it}^{-1}} \quad m_{it} = \sum_{s=t_0}^T \delta^{T-s} (y_s - \hat{y}_s^i)^2 \quad (1.35)$$

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.⁵

The empirical evidence on forecast pooling is mixed. For example, the results in Hendry and Clements (2004) suggested that forecast pooling performed well for macroeconomic variables. In turn, focusing on the role of leading indicators for forecasting US GDP growth and inflation, Banerjee and Marcellino (2006) concluded that single indicator models showed a good performance when compared with automatically selected models, factor models and pooled forecasts, though the information content of single indicators changes over time. Within MIDAS literature, for example, Kuzin et al. (2013) and Foroni and Marcellino (2012) concluded that pooling tends to outperform single-variable models. Moreover, Clements and Galvão (2008) showed that combinations of MIDAS forecasts are at least as good as combinations of forecasts from bridge models and other mixed frequency models as suggested by Koenig et al. (2003).

Figure 1.8 shows the relative RSMFE against an AR benchmark of the best performing MIDAS regressions for the two different pooling techniques - equal-weight and discounted MSFE combinations - and for two different information sets - only monthly data, and monthly and daily data (additional information can be found in Table A.4 in the Appendix).

The results suggest that the discounted MSFE combination tend to perform better in nowcasting and short-term forecast horizons, being as good as the equal-weight combination for the forecast horizons up to 4 quarters ahead. Furthermore, conditional on the discounted MSFE combination, the results from using the largest information set, with both monthly and daily data, are slightly worse than the results from the monthly data set up to 1 quarter ahead, being in general better for 2 to 4 quarters ahead. So, due to the typically forward looking nature of financial variables, the Dow Jones Euro Stoxx daily index contains useful information on future economic developments, which is particularly useful for longer term forecast horizons, even when timely monthly information is included in the

⁵Regarding the discount parameter, different values were considered and the non-discounting option ($\delta = 1$) showed the best results.

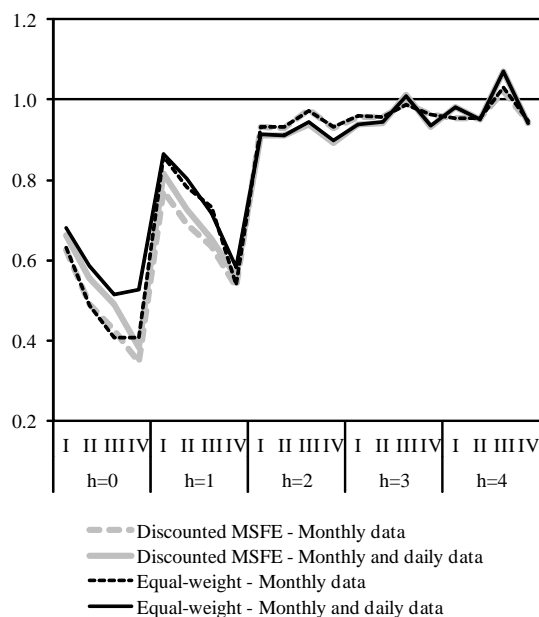


Figure 1.8: Pooled forecasts - best performing MIDAS models against an AR benchmark for each forecast horizon, relative RMSFE

forecasting regressions. Andreou et al. (2013) also concluded for the usefulness of daily financial data. Using a wide set of financial variables available at high frequencies, the authors obtained pooled forecasts from single daily variable models for the US GDP and inflation rate. The authors concluded that adding daily financial data to quarterly macroeconomic information contributes to improve the forecast accuracy.

Using the discounted MSFE combination of forecasts based on the monthly and daily data set, Figure 1.9 compares the three single-variable models above described with the pooled forecasts.

In the case of the single-variable models, the figures plotted correspond to the minimum relative RMSFE of the different MIDAS regressions. For each forecast horizon, in situation *I* there is only available information for the first month of the current quarter for the economic sentiment indicator and the Dow Jones Euro Stoxx index. As there is no information on the industrial production, the figure plotted for the industrial production model corresponds to the relative RMSFE

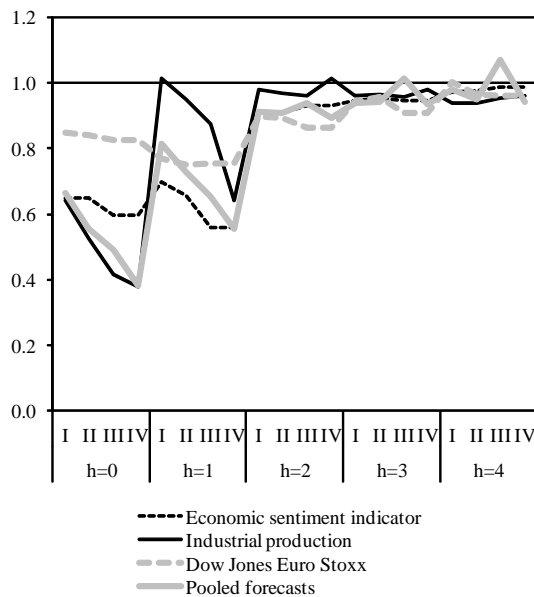


Figure 1.9: MIDAS regressions - single-variable models vs. pooled forecasts against an AR benchmark, for each forecast horizon, relative RMSFE

Note: For each forecast horizon, the relative RMSFE for the pooled forecasts corresponds to the minimum of the different MIDAS models for the discounted MSFE combination with monthly and daily data. In the case of the single-variable models, the figures plotted correspond to the minimum relative RMSFE of the different MIDAS models. For each forecast horizon, in situation "I", there is only available information for the first month of the current quarter for the economic sentiment indicator and the Dow Jones Euro Stoxx index. There is no information on the industrial production. So, the figure plotted for the industrial production model corresponds to the relative RMSFE from forecasts obtained using full-quarter information for the previous quarter.

from forecasts obtained using full-quarter information for the previous quarter. The results suggest that pooling forecasts is a useful tool to circumvent model uncertainty. Although in some circumstances single-variable models can exhibit a better performance (for instance, the economic sentiment indicator model for $h = 1$) pooled forecasts have a good overall performance.

1.5.3 Multi-variable models

Finally, let us address the multi-variable models. Figure 1.10 shows the comparison between the forecasting performance, in terms of RMSFE, of pooled forecasts and multi-variable models, within the MIDAS framework (additional information can

be found in Table A.5 in the Appendix).

The results are borderline, with small differences between combining information from different variables through pooling forecasts and putting several variables into a single model, in terms of relative RMSFE. For nowcasting multi-variable models are slightly better, for 1 and 2 step ahead forecasts the contrary is true and both are fairly similar for 3 and 4 step ahead forecasts. Clements and Galvão (2009) also concluded that using multi-variable MIDAS regressions performed well when compared to forecast combination, in particular for predicting the direction of change.




Figure 1.10: MIDAS models: pooled forecasts vs. multi-variable models - for each forecast horizon, relative RMSFE against an AR benchmark

Note: For each forecast horizon, the relative RMSFE for the pooled forecasts from MIDAS models corresponds to the minimum of the different MIDAS models for the discounted MSFE combination with monthly and daily data. Regarding the multi-variable model, the relative RMSFE corresponds to the minimum of the different models for the monthly and daily datasets.

1.6 Conclusion

Nowadays, techniques that allow taking full advantage of the timely releases of high-frequency data play a key role in forecasting and policymaking. Having started on the financial field, MIDAS regressions have been gaining an increasing attention in macroeconomic forecasting. This technique is a simple, flexible and potentially parsimonious way of taking into account timely releases of high-frequency data, in particular for forecasting a low-frequency series. In this paper, the performance in terms of RMSFE of a wide set of MIDAS regressions is assessed through a recursive forecasting exercise. Alternative ways of dealing with autoregressive augmentation of MIDAS regression are also suggested. The benchmarks used are a simple autoregressive model and traditional quarterly models. In the latter case, a bridge model framework was put in place whenever full-quarter information was not available.

The results obtained suggest that using MIDAS regressions contributes to increase forecast accuracy. The benefits from this data-driven, and potentially more parsimonious, weighting scheme to aggregate the high-frequency data are higher for forecast horizons up to 1 quarter ahead. For longer horizons, the best performing MIDAS regressions also deliver good results, which are significantly better than the benchmarks considered. Flexible MIDAS weighting scheme are able to exploit the information content of high-frequency series for forecasting purposes, regardless of the exact time frequency of the regressors. Furthermore, MIDAS regressions also have a good performance when incomplete current quarter information is used.

This evidence stresses, however, the importance of choosing the best MIDAS model for each specific situation. Hence, the auxiliary choices of the forecaster when applying this technique are crucial for the success in nowcasting and forecasting exercises. Although there is no one-fits-all recipe, the results suggests that the multiplicative and unrestricted MIDAS seem to be a good alternative to the original (polynomial) MIDAS regressions when restrictions on signs of the coefficients play an important role. Moreover, focusing on MIDAS regressions with autoregressive terms, imposing a common factor dynamics between the dependent variable and the regressors (CF-MIDAS) can be, in some cases, too strict. The other ways of introducing autoregressive terms in MIDAS regressions proposed in

this paper - AR-MIDAS, AR-M-MIDAS, AR-U-MIDAS and the respective balanced versions - proved to be good alternatives and in some cases are the best performing MIDAS regression.

Finally, as regards exploiting the information in a small set of selected variables, pooling forecasts (in particular, using the discounted MSFE combination) delivers good results. Although in some cases single-variable models can exhibit a better performance, pooled forecasts have a good overall performance. Multi-variable models also show a good forecasting performance, especially for nowcasting. Moreover, even when timely monthly information is included in the forecasting models, the daily financial variable is still useful, especially for longer term forecast horizons.

Acknowledgments

I am grateful to Carlos Robalo Marques, João Nicolau, Paulo Rodrigues and Christian Schumacher for thoughtful discussions. Comments and suggestions by José Ramos Maria are also gratefully acknowledged. This paper also benefits from comments of participants at the 24th (EC)² Conference “The Econometrics Analysis of Mixed Frequency Data” held at the University of Cyprus. A special thanks to Fátima Teodoro for software assistance. The usual disclaimers apply.

Bibliography

- ABEYSINGHE, T., “Forecasting Singapore’s quarterly GDP with monthly external trade,” *International Journal of Forecasting* 14 (1998), 505–513.
- , “Modeling variables of different frequencies,” *International Journal of Forecasting* 16 (2000), 117–119.
- ABEYSINGHE, T. AND A. S. TAY, “Dynamic Regressions with Variables Observed at Different Frequencies,” Econometric Society World Congress 2000 Contributed Papers 0752, Econometric Society, 2000.
- ALTISSIMO, F., R. CRISTADORO, M. FORNI, M. LIPPI AND G. VERONESE, “New Eurocoin: Tracking Economic Growth in Real Time,” *The Review of Economics and Statistics* 92 (October 2010), 1024–1034.
- ANDREOU, E., E. GHYSELS AND A. KOURTELLOS, “Regression models with mixed sampling frequencies,” *Journal of Econometrics* 158 (2010), 246 – 261.
- , *The Oxford Handbook of Economic Forecasting*, chapter 8, Forecasting with mixed-frequency data (Oxford University Press, 2011), 225–267.
- , “Should macroeconomic forecasters use daily financial data and how?,” *Journal of Business and Economic Statistics* 2 (2013), 240–251.
- ANGELINI, E., J. HENRI AND M. MARCELLINO, “Interpolation and backdating with a large information set,” *Journal of Economic Dynamics and Control* 30 (2006), 2693–2724.
- ARMESTO, M. T., R. HERNÁNDEZ-MURILLO, M. T. OWYANG AND J. PIGER, “Measuring the Information Content of the Beige Book: A Mixed Data Sampling Approach,” *Journal of Money, Credit and Banking* 41 (2009), 35–55.
- ASIMAKOPOULOS, S., J. PAREDES AND T. WARMEDINGER, “Forecasting fiscal time series using mixed frequency data,” Working Paper Series 1550, European Central Bank, May 2013.

BIBLIOGRAPHY

- BAI, J., E. GHYSELS AND J. H. WRIGHT, “State Space Models and MIDAS Regressions,” *Econometric Reviews* 32 (2013), 779–813.
- BANERJEE, A. AND M. MARCELLINO, “Are there any reliable leading indicators for US inflation and GDP growth?,” *International Journal of Forecasting* 22 (2006), 137–151.
- BREITUNG, J., S. ELENGIKAL AND C. ROLING, “Forecasting inflation rates using daily data: A nonparametric MIDAS approach,” mimeo, Institute for Macroeconomics and Econometrics, Bonn University, March 2012.
- CASALS, J., M. JEREZ AND S. SOTOCA, “Modelling and forecasting time series sampled at different frequencies,” *Journal of Forecasting* 28 (2009), 316–342.
- CHEN, X. AND E. GHYSELS, “News — Good or Bad — and Its Impact on Volatility Predictions over Multiple Horizons,” *The Review of Financial Studies* 23 (2010), 46–81.
- CHOW, G. C. AND A.-L. LIN, “Best Linear Unbiased Interpolation, Distribution, and Extrapolation of Time Series by Related Series,” *The Review of Economics and Statistics* 53 (1971), 372–375.
- CLARK, T. AND M. MCCracken, “Evaluating Direct Multistep Forecasts,” *Econometric Reviews* 24 (2005), 369–404.
- CLEMENTS, M. AND A. GALVÃO, “Macroeconomic forecasting with mixed-frequency data: Forecasting output growth in the United States,” *Journal of Business and Economic Statistics* 26 (2008), 546–554.
- , “Forecasting US output growth using leading indicators: An appraisal using MIDAS models,” *Journal of Applied Econometrics* 24 (2009), 1187–1206.
- CORRADO, C. AND M. GREENE, “Reducing uncertainty in short-term projections: Linkage of monthly and quarterly models,” *Journal of Forecasting* 7 (1988), 77–102.
- CORSI, F., “A Simple Approximate Long-Memory Model of Realized Volatility,” *Journal of Financial Econometrics* 7 (Spring 2009), 174–196.
- DHRYMES, P. J., “Distributed Lags: A Survey,” UCLA Economics Working Papers 024, UCLA Department of Economics, October 1972.
- DONIHUE, M. AND E. HOWREY, “Using mixed frequency data to improve macroeconomic forecasts of inventory investment,” *International Journal of Production Economics* 26 (1992), 33–41.

- DOZ, C., D. GIANNONE AND L. REICHLIN, “A Quasi–Maximum Likelihood Approach for Large, Approximate Dynamic Factor Models,” *The Review of Economics and Statistics* 94 (November 2012), 1014–1024.
- ENGLE, R., C. GRANGER AND J. HALLMAN, “Merging short- and long-run forecasts: An application of seasonal cointegration to monthly electricity sales forecasting,” *Journal of Econometrics* 40 (1989), 45–62.
- FERNANDEZ, R. B., “A Methodological Note on the Estimation of Time Series,” *The Review of Economics and Statistics* 63 (1981), 471–476.
- FORNI, M., M. HALLIN, M. LIPPI AND L. REICHLIN, “The Generalized Dynamic-Factor Model: Identification And Estimation,” *The Review of Economics and Statistics* 82 (November 2000), 540–554.
- , “The Generalized Dynamic Factor Model: One-Sided Estimation and Forecasting,” *Journal of the American Statistical Association* 100 (September 2005), 830–840.
- FORONI, C. AND M. MARCELLINO, “A Comparison of Mixed Frequency Approaches for Modelling Euro Area Macroeconomic Variables,” Economics Working Papers ECO2012/07, European University Institute, 2012.
- FORONI, C., M. MARCELLINO AND C. SCHUMACHER, “U-MIDAS: MIDAS regressions with unrestricted lag polynomials,” Discussion Paper Series 1: Economic Studies 2011,35, Deutsche Bundesbank, Research Centre, 2011.
- FORSBERG, L. AND E. GHYSELS, “Why Do Absolute Returns Predict Volatility So Well?,” *Journal of Financial Econometrics* 5 (2007), 31–67.
- FRANCIS, N., E. GHYSELS AND M. T. OWYANG, “The low-frequency impact of daily monetary policy shocks,” Working Papers 2011-009, Federal Reserve Bank of St. Louis, 2011.
- GEWEKE, J., “Temporal Aggregation in the Multiple Regression Model,” *Econometrica* 46 (1978), 643–661.
- GHYSELS, E., “Macroeconomics and the Reality of Mixed Frequency Data,” mimeo, University of North Carolina (UNC) at Chapel Hill - Department of Economics, May 2012.
- GHYSELS, E., P. SANTA-CLARA AND R. VALKANOV, “The MIDAS touch: Mixed data sampling regression models,” CIRANO Working Papers 2004s-20, CIRANO, 2004.

- , “Predicting volatility: getting the most out of return data sampled at different frequencies,” *Journal of Econometrics* 131 (2006), 59–95.
- GHYSELS, E., A. SINKO AND R. VALKANOV, “MIDAS regressions: Further results and new directions,” *Econometric Reviews* 26 (2007), 53–90.
- GIANNONE, D., L. REICHLIN AND D. SMALL, “Nowcasting: The real-time informational content of macroeconomic data,” *Journal of Monetary Economics* 55 (May 2008), 665–676.
- GÓMEZ, V. AND F. APARICIO-PÉREZ, “A new state-space methodology to disaggregate multivariate time series,” *Journal of Time Series Analysis* 30 (2009), 97–124.
- GOLINELLI, R. AND G. PARIGI, “The use of monthly indicators to forecast quarterly GDP in the short run: an application to the G7 countries,” *Journal of Forecasting* 26 (2007), 77–94.
- GRANGER, C., “Extracting information from mega-panels and high-frequency data,” *Statistica Neerlandica* 52 (1998), 258–272.
- GRANGER, C. AND R. RAMANATHAN, “Improved methods of combining forecasts,” *Journal of Forecasting* 3 (1984), 197–204.
- GRILICHES, Z., “Distributed Lags: A Survey,” *Econometrica* 35 (1967), 16–49.
- GUÉRIN, P. AND M. MARCELLINO, “Markov-Switching MIDAS Models,” *Journal of Business & Economic Statistics* 31 (January 2013), 45–56.
- HANSEN, L. AND T. J. SARGENT, *Recursive Models of Dynamic Linear Economies* (Princeton University Press, 2013).
- HENDRY, D. F. AND M. P. CLEMENTS, “Pooling of forecasts,” *Econometrics Journal* 7 (06 2004), 1–31.
- HENDRY, D. F. AND G. E. MIZON, “Serial Correlation as a Convenient Simplification, Not a Nuisance: A Comment on a Study of the Demand for Money by the Bank of England,” *The Economic Journal* 88 (1978), 549–563.
- HOWREY, E., “New methods for using monthly data to improve forecast accuracy,” in L. Klein, ed., *Comparative performance of US economic models* chapter 8 (Oxford University Press, New York, 1991), 227–247.
- HOWREY, E., S. HYMANS AND M. DONIHUE, “Merging monthly and quarterly forecasts: Experience with MQEM,” *Journal of Forecasting* 10 (1991), 255–268.

- HYUNG, N. AND C. GRANGER, “Linking series generated at different frequencies,” *Journal of Forecasting* 27 (2008), 95–108.
- INGENITO, R. AND B. TREHAN, “Using monthly data to predict quarterly output,” *Economic Review* 3 (1996), 3–11.
- JANSEN, J., X. JIN AND J. DE WINTER, “Forecasting and nowcasting real GDP: comparing statistical models and subjective forecasts,” DNB Working Paper 365, De Nederlandsche Bank, December 2012.
- JUDGE, G., W. GRIFFITHS, R. HILL, H. LÜTKEPOHL AND T. LEE, *The theory and practice of econometrics*, Second edition (John Wiley and Sons, 1985).
- KILIAN, L., “Small-Sample Confidence Intervals For Impulse Response Functions,” *The Review of Economics and Statistics* 80 (May 1998), 218–230.
- , “Exchange Rates and Monetary Fundamentals: What do we Learn from Long-Horizon Regressions?,” *Journal of Applied Econometrics* 14 (1999), pp. 491–510.
- KLEIN, L. AND E. SOJO, “Economics in Theory and Practice: An Eclectic Approach,” *Combinations of High and Low Frequency Data in Macroeconometric Models*, chapter 1 (Kluwer Academic Publishers, 1989), 3–16.
- KOENIG, E., S. DOLMAS AND J. PIGER, “The use and abuse of real-time data in economic forecasting,” *The Review of Economics and Statistics* 85 (2003), 618–628.
- KUZIN, V., M. MARCELLINO AND C. SCHUMACHER, “MIDAS versus mixed-frequency VAR: Nowcasting GDP in the euro area,” *International Journal of Forecasting* 27 (2011), 529–542.
- , “Pooling versus model selection for nowcasting GDP with many predictors: empirical evidence for six industrialized countries,” *Journal of Applied Econometrics* 28 (2013), 392–411.
- LITTERMAN, R. B., “A random walk, Markov model for the Distribution of Time Series,” *Journal of Business and Economic Statistics* 1 (1983), 169–173.
- LIU, T.-C., “A Monthly Recursive Econometric Model of United States: A Test of Feasibility,” *The Review of Economics and Statistics* 51 (1969), 1–13.
- LIU, T.-C. AND E.-C. HWA, “A Monthly Econometric Model of the U. S. Economy,” *International Economic Review* 15 (1974), 328–365.

BIBLIOGRAPHY

- LÜTKEPOHL, H., “Linear aggregation of vector autoregressive moving average processes,” *Economics Letters* 14 (1984), 345–350.
- , “Forecasting temporally aggregated vector ARMA processes,” *Journal of Forecasting* 5 (1986), 85–95.
- MARCELLINO, M., “Temporal disaggregation, missing observations, outliers and forecasting: a unifying non-model based procedure,” *Advances in Econometrics* 13 (1998), 181–202.
- , “Consequences of temporal aggregation in empirical analysis,” *Journal of Business and Economic Statistics* 17 (January 1999), 129–136.
- MARCELLINO, M. AND A. MUSSO, “The reliability of real-time estimates of the euro area output gap,” *Economic Modelling* 28 (July 2011), 1842–1856.
- MARCELLINO, M. AND C. SCHUMACHER, “Factor MIDAS for Nowcasting and Forecasting with Ragged-Edge Data: A Model Comparison for German GDP,” *Oxford Bulletin of Economics and Statistics* 72 (08 2010), 518–550.
- MARIANO, R. S. AND Y. MURASAWA, “A new coincident index of business cycles based on monthly and quarterly series,” *Journal of Applied Econometrics* 18 (2003), 427–443.
- MILLER, P. J. AND D. M. CHIN, “Using monthly data to improve quarterly model forecasts,” *Quarterly Review* 20 (1996), 16–33.
- MITTNIK, S. AND P. A. ZADROZNY, “Forecasting Quarterly German GDP at Monthly Intervals Using Monthly IFO Business Conditions Data,” CESifo Working Paper Series 1203, CESifo Group Munich, May 2004.
- MONTEFORTE, L. AND G. MORETTI, “Real-Time Forecasts of Inflation: The Role of Financial Variables,” *Journal of Forecasting* 32 (2013), 51–61.
- NUNES, L. C., “Nowcasting quarterly GDP growth in a monthly coincident indicator model,” *Journal of Forecasting* 24 (2005), 575–592.
- PAIN, N. AND F. SÉDILLOT, “Indicator models of real GDP growth in the major OECD economies,” *OECD Economic Studies* 2005 (2005), 167–217.
- PARIGI, G. AND G. SCHLITZER, “Quarterly forecasts of the Italian business cycle by means of monthly economic indicators,” *Journal of Forecasting* 14 (1995), 117–141.

- PINHEIRO, M., A. RUA AND F. DIAS, “Dynamic Factor Models with Jagged Edge Panel Data: Taking on Board the Dynamics of the Idiosyncratic Components,” *Oxford Bulletin of Economics and Statistics* 75 (2013), 80–102.
- PROIETTI, T., “Temporal disaggregation by state space methods: Dynamic regression methods revisited,” *Econometrics Journal* 9 (2006), 357–372.
- , “Estimation of Common Factors under Cross-Sectional and Temporal Aggregation Constraints,” *International Statistical Review* 79 (December 2011), 455–476.
- RATHJENS, P. AND R. ROBINS, “Forecasting quarterly data using monthly information,” *Journal of Forecasting* 12 (1993), 321–330.
- RÜNSTLER, G., K. BARHOUMI, S. BENK, R. CRISTADORO, A. DEN REIJER, A. JAKAITIENE, P. JELONEK, A. RUA, K. RUTH AND C. VAN NIEUWENHUYZE, “Short-term forecasting of GDP using large datasets: a pseudo real-time forecast evaluation exercise,” *Journal of Forecasting* 28 (2009), 595–611.
- RUEY-WAN, L. AND S. CHUNG-HUA, “The use of high frequency data to improve macroeconomic forecast,” *International Economic Journal* 10 (1996), 65–83.
- SANTOS SILVA, J. M. C. AND F. N. CARDOSO, “The Chow-Lin method using dynamic models,” *Economic Modelling* 18 (2001), 269 – 280.
- SIMS, C. A., “Discrete Approximations to Continuous Time Distributed Lags in Econometrics,” *Econometrica* 39 (1971), 545–563.
- STARK, T., “Does current-quarter information improve quarterly forecasts for the U.S. economy?,” Working Papers 00-2, Federal Reserve Bank of Philadelphia, 2000.
- STOCK, J. H. AND M. W. WATSON, “New Indexes of Coincident and Leading Economic Indicators,” in *NBER Macroeconomics Annual 1989* volume 4 of *NBER Chapters* (National Bureau of Economic Research, Inc, 1989), 351–409.
- , “Macroeconomic Forecasting Using Diffusion Indexes,” *Journal of Business and Economic Statistics* 20 (April 2002), 147–62.
- , “Combination forecasts of output growth in a seven-country data set,” *Journal of Forecasting* 23 (2004), 405–430.
- WEI, W. W. S. AND D. O. STRAM, “Disaggregation of Time Series Models,” *Journal of the Royal Statistical Society. Series B (Methodological)* 52 (1990), pp. 453–467.

BIBLIOGRAPHY

- WEISS, A., “Systematic sampling and temporal aggregation in time series models,” *Journal of Econometrics* 26 (1984), 271–281.
- ZADROZNY, P., “Gaussian likelihood of continuous-time ARMAX models when data are stocks and flows at different frequencies,” *Econometric Theory* 4 (April 1988), 108–124.
- , “Estimating a multivariate ARMA model with mixed-frequency data: An application to forecasting US GNP at monthly intervals,” Working Paper Series 90-6, Federal Reserve Bank of Atlanta, August 1990.
- ZHENG, I. Y. AND J. ROSSITER, “Using Monthly Indicators to Predict Quarterly GDP,” Working Papers 06-26, Bank of Canada, 2006.

Appendices

Table A.1: Single-variable models: relative performance, in terms of RMSFE, against an AR benchmark, for forecast horizon h

	h=0			h=1			h=2			h=3			h=4		
	I	II	III	I	II	III	I	II	III	I	II	III	I	II	III
Industrial Production															
MIDAS	0.676	0.560	0.415	1.024	1.018	0.702	1.030	1.017	1.038	1.017	1.005	1.017	0.970	0.970	0.970
	[0.000]	[0.000]	[0.000]	[0.188]	[0.147]	[0.000]	[0.244]	[0.139]	[0.225]	[0.096]	[0.125]	[0.091]	[0.010]	[0.008]	[0.008]
Multiplicative MIDAS	0.624	0.472	0.443	1.164	1.063	0.646	1.083	1.131	1.179	1.019	1.062	1.121	1.022	1.027	1.010
	[0.000]	[0.000]	[0.000]	[0.705]	[0.277]	[0.000]	[0.471]	[0.510]	[0.327]	[0.159]	[0.265]	[0.306]	[0.176]	[0.129]	[0.057]
Unrestricted MIDAS	0.585	0.464	0.457	1.175	0.919	0.646	1.203	1.067	1.113	1.015	1.056	1.106	1.014	1.020	0.998
	[0.000]	[0.000]	[0.000]	[0.824]	[0.008]	[0.000]	[0.863]	[0.535]	[0.646]	[0.162]	[0.444]	[0.644]	[0.143]	[0.124]	[0.086]
CF-MIDAS	0.525	0.449	0.379	0.948	0.875	0.641	0.967	0.961	1.014	0.987	0.956	1.013	1.026	1.025	1.005
	[0.000]	[0.000]	[0.000]	[0.018]	[0.004]	[0.000]	[0.025]	[0.021]	[0.119]	[0.037]	[0.013]	[0.071]	[0.166]	[0.162]	[0.086]
AR-MIDAS	0.629	0.504	0.407	1.052	1.022	0.715	1.079	1.002	1.059	0.979	0.981	0.984	1.012	1.005	0.979
	[0.000]	[0.000]	[0.000]	[0.307]	[0.182]	[0.000]	[0.477]	[0.050]	[0.293]	[0.024]	[0.026]	[0.022]	[0.103]	[0.057]	[0.026]
AR-M-MIDAS	0.609	0.466	0.437	1.129	1.000	0.664	1.085	1.120	1.196	0.964	1.049	1.085	1.032	0.971	0.960
	[0.000]	[0.000]	[0.000]	[0.243]	[0.215]	[0.000]	[0.232]	[0.252]	[0.338]	[0.006]	[0.065]	[0.296]	[0.168]	[0.011]	[0.018]
AR-U-MIDAS	0.635	0.449	0.378	1.137	0.948	0.664	1.156	1.114	1.143	1.001	1.075	1.127	1.026	1.037	0.983
	[0.000]	[0.000]	[0.000]	[0.583]	[0.024]	[0.000]	[0.635]	[0.651]	[0.668]	[0.439]	[0.471]	[0.644]	[0.174]	[0.167]	[0.012]
Balanced AR-MIDAS	0.609	0.418	0.415	1.080	0.894	0.664	1.154	1.154	1.154	0.980	0.980	0.980	1.026	1.026	1.014
	[0.000]	[0.000]	[0.000]	[0.744]	[0.007]	[0.000]	[0.623]	[0.600]	[0.542]	[0.035]	[0.022]	[0.019]	[0.234]	[0.188]	[0.107]
Balanced AR-M-MIDAS	0.631	0.448	0.437	1.080	0.894	0.664	1.131	1.157	1.196	1.010	1.072	1.085	0.940	0.955	0.960
	[0.000]	[0.000]	[0.000]	[0.744]	[0.007]	[0.000]	[0.267]	[0.304]	[0.338]	[0.125]	[0.313]	[0.296]	[0.010]	[0.017]	[0.018]
Balanced AR-U-MIDAS	0.594	0.426	0.393	1.082	0.904	0.664	1.078	1.122	1.143	1.003	1.086	1.127	1.003	1.015	0.983
	[0.000]	[0.000]	[0.000]	[0.779]	[0.011]	[0.000]	[0.716]	[0.749]	[0.668]	[0.191]	[0.636]	[0.644]	[0.207]	[0.205]	[0.012]
Economic Sentiment Indicator															
MIDAS	1.080	0.815	0.716	0.975	0.910	0.846	0.986	0.986	0.979	0.995	0.995	0.986	0.985	0.984	0.993
	[0.073]	[0.002]	[0.000]	[0.047]	[0.007]	[0.000]	[0.096]	[0.084]	[0.052]	[0.168]	[0.166]	[0.097]	[0.071]	[0.080]	[0.115]
Multiplicative MIDAS	0.648	0.658	0.656	0.767	0.665	0.559	0.940	0.945	0.929	1.033	1.013	0.985	1.006	0.994	0.995
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.014]	[0.021]	[0.008]	[0.220]	[0.217]	[0.091]	[0.108]	[0.113]	[0.141]
Unrestricted MIDAS	0.670	0.655	0.666	0.699	0.662	0.623	0.956	0.926	0.950	1.004	1.019	1.023	0.993	0.994	0.988
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.024]	[0.013]	[0.027]	[0.170]	[0.301]	[0.317]	[0.107]	[0.114]	[0.083]
CF-MIDAS	0.763	0.737	0.696	0.882	0.883	0.869	0.910	0.910	0.934	0.947	0.952	0.947	0.973	0.976	0.987
	[0.001]	[0.000]	[0.000]	[0.003]	[0.004]	[0.000]	[0.010]	[0.011]	[0.010]	[0.035]	[0.036]	[0.036]	[0.057]	[0.073]	[0.095]
AR-MIDAS	0.966	1.017	0.799	1.012	1.010	0.881	1.007	1.007	1.007	0.980	1.003	0.979	0.990	0.990	0.990
	[0.147]	[0.183]	[0.203]	[0.299]	[0.322]	[0.329]	[0.483]	[0.481]	[0.493]	[0.438]	[0.469]	[0.425]	[0.432]	[0.425]	[0.427]
AR-M-MIDAS	0.685	0.688	0.652	0.768	0.740	0.593	0.963	0.966	0.942	1.018	1.025	1.003	1.008	0.996	0.996
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.031]	[0.029]	[0.010]	[0.155]	[0.175]	[0.205]	[0.094]	[0.090]	[0.079]
AR-U-MIDAS	0.654	0.669	0.597	0.730	0.676	0.625	0.947	0.941	0.970	1.008	1.015	1.043	0.997	0.995	0.993
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.022]	[0.021]	[0.044]	[0.169]	[0.212]	[0.371]	[0.097]	[0.101]	[0.083]
Balanced AR-MIDAS	0.830	0.819	0.799	1.013	1.013	1.013	1.007	1.007	1.007	1.003	1.003	1.003	0.990	0.990	0.990
	[0.000]	[0.000]	[0.000]	[0.227]	[0.225]	[0.221]	[0.206]	[0.208]	[0.206]	[0.204]	[0.209]	[0.205]	[0.088]	[0.085]	[0.085]
Balanced AR-M-MIDAS	0.691	0.696	0.652	0.713	0.666	0.593	0.921	0.935	0.942	1.003	1.003	1.003	1.005	1.005	1.005
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.005]	[0.010]	[0.010]	[0.204]	[0.209]	[0.205]	[0.179]	[0.173]	[0.173]
Balanced AR-U-MIDAS	0.650	0.650	0.597	0.707	0.655	0.625	0.960	0.951	0.970	0.983	1.011	1.043	0.978	0.990	0.993
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.039]	[0.023]	[0.044]	[0.112]	[0.221]	[0.371]	[0.063]	[0.092]	[0.083]
Dow Jones Euro Stoxx															
MIDAS	0.857	0.859	0.850	0.822	0.799	0.778	0.956	0.947	0.947	0.957	0.955	0.970	1.013	0.974	0.972
	[0.001]	[0.002]	[0.000]	[0.002]	[0.001]	[0.000]	[0.045]	[0.024]	[0.032]	[0.044]	[0.039]	[0.079]	[0.229]	[0.064]	[0.072]
Multiplicative MIDAS	0.945	0.920	0.876	0.904	0.799	0.829	0.977	0.955	0.947	1.007	0.974	0.975	1.021	0.983	0.968
	[0.009]	[0.001]	[0.000]	[0.008]	[0.000]	[0.000]	[0.081]	[0.036]	[0.035]	[0.246]	[0.070]	[0.091]	[0.317]	[0.074]	[0.056]
CF-MIDAS	0.846	0.849	0.837	0.789	0.809	0.785	0.895	0.894	0.862	0.943	0.963	0.907	1.003	0.994	0.980
	[0.004]	[0.000]	[0.005]	[0.000]	[0.000]	[0.000]	[0.006]	[0.007]	[0.001]	[0.025]	[0.046]	[0.012]	[0.173]	[0.119]	[0.084]
AR-MIDAS	0.854	0.842	0.825	0.768	0.752	0.755	0.942	0.956	0.899	0.957	0.971	0.939	1.012	0.966	0.961
	[0.004]	[0.001]	[0.001]	[0.000]	[0.000]	[0.000]	[0.026]	[0.035]	[0.006]	[0.033]	[0.063]	[0.031]	[0.189]	[0.047]	[0.042]
AR-M-MIDAS	0.911	0.870	0.916	0.844	0.826	0.818	0.938	0.945	0.897	0.974	0.971	0.937	1.011	0.964	0.961
	[0.002]	[0.003]	[0.006]	[0.000]	[0.000]	[0.000]	[0.023]	[0.034]	[0.006]	[0.050]	[0.037]	[0.029]	[0.168]	[0.053]	[0.042]

Note: For each quarterly forecast horizon (h), depending on the within-quarter information used, "I" refers to one month, "II" to two months and "III" to a full quarter. The figures in bold denote the minimum relative RMSFE for each single-variable model. The figures shaded denote the minimum relative RMSFE for each forecast horizon, across the single-variable models. The figures in brackets denote the empirical rejection frequencies for the null hypothesis of equal RMSFE calculated using bootstrapped test statistics.

Table A.2: Single-variable models: relative performance, in terms of RMSFE, against a traditional low-frequency quarterly benchmark, for forecast horizon h

	h=0			h=1			h=2			h=3			h=4		
	I	II	III	I	II	III	I	II	III	I	II	III	I	II	III
Industrial Production															
MIDAS	1.248	1.317	1.001	1.896	2.539	0.762	1.672	1.951	0.930	1.097	1.093	0.975	0.826	0.817	0.966
	[0.784]	[0.839]	[0.197]	[1.000]	[1.000]	[0.002]	[1.000]	[1.000]	[0.037]	[0.774]	[0.892]	[0.086]	[0.095]	[0.093]	[0.079]
Multiplicative MIDAS	1.152	1.110	1.068	2.156	2.651	0.702	1.758	2.170	1.056	1.100	1.154	1.075	0.871	0.866	1.005
	[0.542]	[0.094]	[0.044]	[0.999]	[1.000]	[0.034]	[0.966]	[0.992]	[0.332]	[0.589]	[0.603]	[0.411]	[0.165]	[0.132]	[0.424]
Unrestricted MIDAS	1.080	1.091	1.101	2.175	2.292	0.701	1.953	2.047	0.997	1.095	1.148	1.060	0.863	0.860	0.994
	[0.698]	[0.647]	[0.609]	[1.000]	[1.000]	[0.000]	[1.000]	[1.000]	[0.272]	[0.845]	[0.936]	[0.696]	[0.063]	[0.049]	[0.203]
CF-MIDAS	0.969	1.056	0.913	1.756	2.183	0.697	1.570	1.844	0.908	1.065	1.039	0.971	0.874	0.864	1.001
	[0.025]	[0.265]	[0.021]	[0.998]	[1.000]	[0.001]	[1.000]	[1.000]	[0.030]	[0.655]	[0.583]	[0.087]	[0.046]	[0.058]	[0.307]
AR-MIDAS	1.162	1.184	0.981	1.947	2.547	0.776	1.752	1.922	0.949	1.057	1.067	0.942	0.862	0.847	0.975
	[0.838]	[0.747]	[0.100]	[1.000]	[1.000]	[0.002]	[1.000]	[1.000]	[0.051]	[0.627]	[0.665]	[0.029]	[0.033]	[0.085]	[0.081]
AR-M-MIDAS	1.125	1.096	1.054	2.090	2.493	0.721	1.762	2.149	1.071	1.040	1.140	1.040	0.879	0.818	0.956
	[0.158]	[0.271]	[0.103]	[0.999]	[1.000]	[0.000]	[0.999]	[1.000]	[0.132]	[0.288]	[0.483]	[0.250]	[0.075]	[0.043]	[0.041]
AR-U-MIDAS	1.173	1.056	0.911	2.105	2.363	0.721	1.877	2.137	1.024	1.080	1.169	1.080	0.874	0.874	0.979
	[0.894]	[0.368]	[0.003]	[1.000]	[1.000]	[0.000]	[1.000]	[1.000]	[0.332]	[0.518]	[0.929]	[0.678]	[0.177]	[0.136]	[0.029]
Balanced AR-MIDAS	1.125	0.983	1.000	1.999	2.230	0.721	1.873	2.214	1.034	1.058	1.066	0.939	0.874	0.865	1.009
	[0.737]	[0.140]	[0.147]	[1.000]	[1.000]	[0.000]	[0.999]	[1.000]	[0.221]	[0.683]	[0.674]	[0.028]	[0.129]	[0.099]	[0.352]
Balanced AR-M-MIDAS	1.164	1.053	1.054	1.999	2.230	0.721	1.836	2.221	1.071	1.090	1.166	1.040	0.801	0.805	0.956
	[0.346]	[0.103]	[0.103]	[1.000]	[1.000]	[0.000]	[0.996]	[1.000]	[0.132]	[0.793]	[0.862]	[0.250]	[0.013]	[0.009]	[0.041]
Balanced AR-U-MIDAS	1.098	1.003	0.947	2.003	2.253	0.721	1.751	2.154	1.024	1.083	1.181	1.080	0.854	0.856	0.979
	[0.731]	[0.095]	[0.016]	[1.000]	[1.000]	[0.000]	[1.000]	[1.000]	[0.332]	[0.916]	[0.953]	[0.678]	[0.025]	[0.025]	[0.029]
Economic Sentiment Indicator															
MIDAS	1.713	1.303	1.164	1.430	1.362	1.242	1.087	1.089	0.987	0.984	0.983	0.982	0.900	0.899	0.996
	[0.220]	[0.206]	[0.170]	[0.955]	[0.978]	[0.142]	[0.830]	[0.866]	[0.135]	[0.121]	[0.121]	[0.055]	[0.079]	[0.074]	[0.053]
Multiplicative MIDAS	1.028	1.052	1.066	1.125	0.995	0.821	1.036	1.043	0.937	1.022	1.000	0.981	0.920	0.907	0.997
	[0.462]	[0.717]	[0.824]	[0.992]	[0.981]	[0.002]	[0.710]	[0.660]	[0.040]	[0.245]	[0.275]	[0.076]	[0.092]	[0.077]	[0.212]
Unrestricted MIDAS	1.062	1.046	1.082	1.025	0.991	0.915	1.054	1.022	0.958	0.992	1.007	1.019	0.908	0.907	0.990
	[0.658]	[0.402]	[0.392]	[0.962]	[0.982]	[0.020]	[0.660]	[0.629]	[0.051]	[0.249]	[0.383]	[0.363]	[0.086]	[0.092]	[0.112]
CF-MIDAS	1.209	1.177	1.131	1.294	1.321	1.276	1.003	1.004	0.941	0.936	0.941	0.943	0.889	0.891	0.990
	[0.097]	[0.098]	[0.181]	[0.909]	[0.971]	[0.175]	[0.421]	[0.470]	[0.041]	[0.029]	[0.064]	[0.030]	[0.066]	[0.083]	[0.099]
AR-MIDAS	1.532	1.624	1.298	1.484	1.512	1.293	1.110	1.112	1.016	0.969	0.991	0.975	0.906	0.904	0.993
	[0.453]	[0.398]	[0.367]	[0.983]	[0.996]	[0.299]	[0.914]	[0.927]	[0.455]	[0.070]	[0.072]	[0.070]	[0.095]	[0.093]	[0.187]
AR-M-MIDAS	1.086	1.100	1.060	1.126	1.107	0.870	1.061	1.067	0.950	1.006	1.012	0.999	0.922	0.909	0.997
	[0.653]	[0.733]	[0.528]	[0.985]	[0.991]	[0.006]	[0.734]	[0.713]	[0.044]	[0.202]	[0.212]	[0.217]	[0.078]	[0.067]	[0.133]
AR-U-MIDAS	1.037	1.068	0.970	1.071	1.012	0.918	1.043	1.039	0.978	0.997	1.003	1.039	0.912	0.909	0.995
	[0.193]	[0.276]	[0.049]	[0.969]	[0.980]	[0.018]	[0.609]	[0.641]	[0.085]	[0.225]	[0.307]	[0.445]	[0.080]	[0.080]	[0.125]
Balanced AR-MIDAS	1.316	1.308	1.298	1.486	1.517	1.488	1.110	1.112	1.016	0.992	0.990	0.999	0.906	0.904	0.993
	[0.347]	[0.330]	[0.367]	[0.988]	[0.990]	[0.371]	[0.934]	[0.942]	[0.455]	[0.345]	[0.335]	[0.217]	[0.100]	[0.094]	[0.187]
Balanced AR-M-MIDAS	1.095	1.112	1.060	1.046	0.997	0.870	1.015	1.033	0.950	0.992	0.990	0.999	0.919	0.917	1.007
	[0.831]	[0.818]	[0.528]	[0.958]	[0.980]	[0.006]	[0.505]	[0.612]	[0.044]	[0.345]	[0.335]	[0.217]	[0.113]	[0.107]	[0.353]
Balanced AR-U-MIDAS	1.032	1.039	0.970	1.037	0.981	0.918	1.058	1.051	0.978	0.972	0.999	1.039	0.894	0.903	0.995
	[0.224]	[0.186]	[0.049]	[0.951]	[0.019]	[0.018]	[0.797]	[0.741]	[0.085]	[0.104]	[0.324]	[0.445]	[0.062]	[0.076]	[0.125]
Dow Jones Euro Stoxx															
MIDAS	1.026	1.037	1.027	0.924	0.887	0.964	1.051	1.041	1.045	1.033	1.039	1.005	1.079	1.039	1.006
	[0.118]	[0.145]	[0.146]	[0.070]	[0.068]	[0.073]	[0.925]	[0.931]	[0.336]	[0.739]	[0.784]	[0.243]	[0.905]	[0.858]	[0.336]
Multiplicative MIDAS	1.131	1.110	1.059	1.016	0.887	1.026	1.074	1.049	1.045	1.087	1.059	1.010	1.088	1.048	1.002
	[0.352]	[0.293]	[0.155]	[0.960]	[0.068]	[0.201]	[0.929]	[0.950]	[0.355]	[0.880]	[0.844]	[0.257]	[0.909]	[0.798]	[0.209]
CF-MIDAS	1.013	1.025	1.012	0.887	0.899	0.973	0.985	0.983	0.951	1.018	1.047	0.939	1.069	1.060	1.014
	[0.102]	[0.103]	[0.102]	[0.089]	[0.090]	[0.098]	[0.037]	[0.037]	[0.036]	[0.637]	[0.800]	[0.033]	[0.882]	[0.885]	[0.241]
AR-MIDAS	1.022	1.016	0.998	0.863	0.835	0.935	1.036	1.051	0.993	1.033	1.056	0.973	1.079	1.030	0.995
	[0.535]	[0.450]	[0.355]	[0.026]	[0.025]	[0.028]	[0.855]	[0.932]	[0.143]	[0.672]	[0.822]	[0.056]	[0.875]	[0.789]	[0.155]
AR-M-MIDAS	1.090	1.050	1.108	0.949	0.917	1.013	1.032	1.039	0.990	1.051	1.056	0.971	1.077	1.028	0.995
	[0.840]	[0.633]	[0.867]	[0.117]	[0.113]	[0.125]	[0.796]	[0.908]	[0.092]	[0.719]	[0.765]	[0.044]	[0.859]	[0.742]	[0.155]

Note: For each quarterly forecast horizon (h), depending on the within-quarter information used, "I" refers to one month, "II" to two months and "III" to a full quarter. The figures in bold denote the minimum relative RMSFE for each single-variable model. The figures in brackets denote the empirical rejection frequencies for the null hypothesis of equal RMSFE calculated using bootstrapped test statistics.

Table A.3: Single-variable model for different time frequencies, against an autoregressive benchmark, for forecast horizon h

	h=0			h=1			h=2			h=3			h=4		
	I	II	III	I	II	III	I	II	III	I	II	III	I	II	III
Monthly															
MIDAS	0.909	0.878	0.876	0.873	0.782	0.814	0.959	0.950	0.884	0.957	0.968	0.959	1.002	0.982	0.991
Multiplicative MIDAS	0.943	0.946	0.904	0.802	0.853	0.814	0.965	1.003	0.965	1.028	1.031	1.004	1.031	1.004	0.996
Unrestricted MIDAS	1.102	1.096	1.104	0.837	0.844	0.876	0.957	0.943	0.933	0.983	0.997	0.990	1.011	0.992	0.988
CF-MIDAS	0.913	0.837	0.838	0.779	0.745	0.742	0.892	0.902	0.889	0.944	0.980	0.981	1.017	0.995	0.998
AR-MIDAS	0.869	0.847	0.817	0.793	0.790	0.833	0.953	0.960	0.921	0.958	0.963	0.970	1.001	0.983	0.980
AR-M-MIDAS	0.865	0.852	0.839	0.807	0.822	0.795	1.021	1.033	0.956	1.046	1.054	1.022	1.030	1.010	0.997
AR-U-MIDAS	0.980	0.973	0.979	0.805	0.817	0.790	0.980	0.967	0.953	0.994	1.001	0.998	1.007	0.991	0.992
Balanced AR-MIDAS	0.820	0.819	0.817	0.821	0.852	0.833	0.947	0.949	0.921	0.985	0.984	0.970	1.018	0.983	0.980
Balanced AR-M-MIDAS	0.839	0.841	0.839	0.768	0.810	0.795	0.978	0.973	0.956	1.031	1.027	1.022	1.035	1.002	0.997
Balanced AR-U-MIDAS	0.973	0.981	0.966	0.819	0.825	0.790	1.000	0.995	0.953	1.029	1.011	0.998	1.040	0.989	0.992
Weekly															
MIDAS	0.919	0.904	0.915	0.761	0.774	0.769	0.951	0.946	0.942	0.961	0.956	0.974	1.005	0.978	0.971
Multiplicative MIDAS	0.917	0.870	0.915	1.006	0.777	0.818	0.963	0.961	0.941	1.020	0.992	0.980	1.015	0.982	0.971
CF-MIDAS	0.851	0.840	0.846	0.735	0.747	0.741	0.893	0.893	0.854	0.994	0.993	0.966	1.007	1.005	0.982
AR-MIDAS	0.845	0.833	0.812	0.771	0.758	0.765	0.959	0.953	0.886	0.965	0.966	0.939	1.002	0.974	0.947
AR-M-MIDAS	0.874	0.824	0.889	0.813	0.798	0.836	1.017	1.022	0.888	1.035	1.000	0.986	1.014	0.979	1.045
Daily															
MIDAS	0.857	0.859	0.850	0.822	0.799	0.778	0.956	0.947	0.947	0.957	0.955	0.970	1.013	0.974	0.972
Multiplicative MIDAS	0.945	0.920	0.876	0.904	0.799	0.829	0.977	0.955	0.947	1.007	0.974	0.975	1.021	0.983	0.968
CF-MIDAS	0.846	0.849	0.837	0.789	0.809	0.785	0.895	0.894	0.862	0.943	0.963	0.907	1.003	0.994	0.980
AR-MIDAS	0.854	0.842	0.825	0.768	0.752	0.755	0.942	0.956	0.899	0.957	0.971	0.939	1.012	0.966	0.961
AR-M-MIDAS	0.911	0.870	0.916	0.844	0.826	0.818	0.938	0.945	0.897	0.974	0.971	0.937	1.011	0.964	0.961

For each quarterly forecast horizon (h), depending on the within-quarter information used, "I" refers to one month, "II" to two months and "III" to a full quarter. The figures in bold denote the minimum relative RMSFE for each single-variable model. The figures shaded denote the minimum relative RMSFE for each forecast horizon, across the single-variable models.

Table A.4: Pooled forecasts: relative performance, in terms of RMSFE, against an AR benchmark, for forecast horizon h

	h=0				h=1				h=2				h=3				h=4			
	I	II	III	IV	I	II	III	IV	I	II	III	IV	I	II	III	IV	I	II	III	IV
Equal weight combination (mean)																				
Monthly data																				
MIDAS	0.831	0.581	0.467	0.418	0.995	0.935	0.886	0.692	0.994	0.983	1.001	0.999	0.970	1.003	1.019	0.995	0.988	0.965	1.039	0.968
Multiplicative MIDAS	0.633	0.493	0.406	0.406	0.444	0.950	0.837	0.539	0.539	0.999	0.999	1.038	1.009	1.000	1.042	1.035	0.988	0.988	1.075	0.998
Unrestricted MIDAS	0.631	0.499	0.421	0.432	0.875	0.854	0.750	0.594	1.018	1.022	1.011	1.018	1.038	1.004	1.057	1.057	0.986	0.987	1.061	0.987
CF-MIDAS	0.638	0.493	0.443	0.411	0.927	0.859	0.831	0.665	0.933	0.931	0.973	0.931	0.958	0.958	0.986	0.964	0.988	0.993	1.067	0.994
AR-MIDAS	0.897	0.764	0.473	0.480	1.015	1.022	0.906	0.719	1.007	1.038	0.975	0.998	0.961	0.975	1.001	0.963	0.982	0.983	1.040	0.946
AR-M-MIDAS	0.725	0.557	0.462	0.490	0.933	0.883	0.851	0.564	1.021	1.008	1.028	1.041	1.005	0.968	1.028	1.022	0.988	0.991	1.043	0.960
AR-U-MIDAS	0.690	0.545	0.462	0.406	0.895	0.843	0.861	0.601	1.005	1.022	1.039	1.041	0.996	0.985	1.041	1.078	0.990	0.993	1.083	0.954
Balanced AR-MIDAS	0.787	0.612	0.449	0.478	1.043	1.010	0.733	0.745	1.035	1.035	0.999	1.035	0.968	0.968	1.009	0.968	0.987	0.986	1.052	0.994
Balanced AR-M-MIDAS	0.667	0.545	0.456	0.490	0.860	0.783	0.733	0.564	0.986	1.001	1.034	1.041	1.001	1.002	1.061	1.022	0.955	0.954	1.030	0.961
Balanced AR-U-MIDAS	0.686	0.487	0.427	0.416	0.857	0.781	0.738	0.601	0.997	0.996	1.028	1.041	0.977	0.990	1.074	1.078	0.983	0.988	1.057	0.954
Monthly and daily data																				
MIDAS	0.806	0.616	0.524	0.527	0.903	0.868	0.811	0.686	0.969	0.960	0.961	0.948	0.951	0.967	1.033	0.982	0.988	0.950	1.079	0.953
Multiplicative MIDAS	0.692	0.585	0.514	0.545	0.867	0.804	0.716	0.583	0.969	0.978	0.943	0.933	1.003	0.979	1.040	0.990	0.992	0.974	1.103	0.978
CF-MIDAS	0.681	0.590	0.532	0.538	0.864	0.830	0.785	0.680	0.913	0.912	0.944	0.898	0.940	0.943	1.008	0.937	0.982	0.977	1.095	0.983
AR-MIDAS	0.870	0.774	0.542	0.577	0.920	0.904	0.822	0.709	0.974	0.999	0.944	0.954	0.952	0.959	1.018	0.942	0.985	0.961	1.070	0.940
AR-M-MIDAS	0.778	0.645	0.528	0.570	0.888	0.852	0.790	0.600	0.985	0.979	0.964	0.966	0.988	0.959	1.030	0.986	0.987	0.964	1.079	0.955
Discounted MSFE combination																				
Monthly data																				
MIDAS	0.832	0.582	0.467	0.363	0.991	0.928	0.876	0.680	0.994	0.983	1.001	0.997	0.970	1.003	1.017	0.994	0.988	0.965	1.041	0.968
Multiplicative MIDAS	0.627	0.499	0.426	0.412	0.848	0.719	0.635	0.533	0.994	0.988	1.011	1.009	1.009	1.000	1.042	1.026	0.984	0.988	1.076	0.998
Unrestricted MIDAS	0.626	0.508	0.443	0.403	0.781	0.739	0.693	0.593	1.005	0.998	1.006	1.005	0.998	1.004	1.056	1.054	0.986	0.987	1.063	0.987
CF-MIDAS	0.637	0.490	0.438	0.345	0.922	0.856	0.831	0.636	0.931	0.929	0.976	0.928	0.958	0.957	0.990	0.987	0.992	0.992	1.072	0.993
AR-MIDAS	0.893	0.739	0.456	0.388	1.015	1.022	0.904	0.703	1.006	1.036	0.976	0.996	0.961	0.975	1.000	0.963	0.981	0.983	1.042	0.946
AR-M-MIDAS	0.712	0.569	0.499	0.453	0.853	0.806	0.835	0.560	1.011	0.999	1.022	1.012	1.005	0.968	1.028	1.018	0.988	0.991	1.045	0.959
AR-U-MIDAS	0.674	0.554	0.505	0.367	0.815	0.745	0.850	0.600	0.995	1.007	1.034	1.027	0.968	0.985	1.040	1.075	0.990	0.993	1.085	0.954
Balanced AR-MIDAS	0.786	0.619	0.458	0.391	1.040	1.008	0.724	0.683	1.026	1.026	0.990	1.026	0.968	0.968	1.007	0.968	0.986	0.986	1.057	0.994
Balanced AR-M-MIDAS	0.661	0.551	0.474	0.453	0.780	0.699	0.685	0.560	0.971	0.985	1.022	1.012	1.001	1.002	1.061	1.018	0.953	0.952	1.020	0.960
Balanced AR-U-MIDAS	0.659	0.508	0.478	0.382	0.773	0.691	0.686	0.600	0.992	0.990	1.022	1.027	0.977	0.990	1.073	1.075	0.982	0.988	1.059	0.954
Monthly and daily data																				
MIDAS	0.808	0.612	0.519	0.399	0.886	0.851	0.795	0.674	0.967	0.959	0.960	0.945	0.950	0.966	1.033	0.981	0.988	0.950	1.079	0.953
Multiplicative MIDAS	0.663	0.554	0.491	0.441	0.814	0.728	0.651	0.557	0.960	0.970	0.937	0.913	1.003	0.979	1.040	0.983	0.991	0.975	1.103	0.978
CF-MIDAS	0.669	0.565	0.507	0.382	0.850	0.824	0.778	0.652	0.910	0.909	0.945	0.892	0.939	0.943	1.012	0.933	0.981	0.977	1.100	0.982
AR-MIDAS	0.865	0.759	0.541	0.432	0.884	0.862	0.797	0.695	0.970	0.993	0.944	0.946	0.951	0.958	1.018	0.941	0.984	0.962	1.071	0.940
AR-M-MIDAS	0.755	0.633	0.539	0.474	0.838	0.803	0.783	0.579	0.975	0.971	0.959	0.940	0.987	0.958	1.031	0.979	0.987	0.964	1.081	0.955

Note: For each quarterly forecast horizon (h), depending on the within-quarter information used, "I" refers to one month information for the economic sentiment indicator and the Dow Jones Euro Stoxx index and no current quarter information for the industrial production, "II" to two months of the economic sentiment indicator and the Dow Jones Euro Stoxx index and one month of the industrial production, "III" to three months of the economic sentiment indicator and the Dow Jones Euro Stoxx index and two months of the industrial production, and, finally, "IV" to full-quarter information for all variables. The figures in bold denote the minimum relative RMSFE for each kind of MIDAS model and each pooling technique. The figures shaded denote the minimum relative RMSFE for each forecast horizon and each pooling technique. The figures in a square refer to the overall minima.

Table A.5: Multi-variable models, against an AR benchmark, for forecast horizon h

	h=0				h=1				h=2				h=3				h=4			
	I	II	III	IV	I	II	III	IV	I	II	III	IV	I	II	III	IV	I	II	III	IV
Monthly data																				
MIDAS	0.745	0.525	0.404	0.373	1.000	0.948	0.898	0.612	1.025	1.014	1.011	1.052	1.013	0.982	1.045	0.993	1.004	1.000	1.004	1.002
Multiplicative MIDAS	0.614	0.499	0.451	0.559	0.946	0.930	0.817	0.748	1.068	1.142	1.222	1.310	1.066	1.076	1.124	1.087	1.031	1.052	1.130	1.056
Unrestricted MIDAS	0.620	0.499	0.445	0.533	0.759	0.779	0.608	0.600	1.031	1.053	1.069	1.109	1.021	1.066	1.149	1.140	0.995	1.104	1.092	1.020
CF-MIDAS	0.602	0.434	0.343	0.422	1.159	0.981	1.018	0.611	1.050	0.980	1.031	1.038	1.006	1.040	1.073	1.023	1.033	1.044	1.051	1.066
AR-MIDAS	0.809	0.548	0.392	0.377	1.009	1.017	0.935	0.701	1.032	1.073	1.093	1.191	0.938	0.979	0.959	0.945	0.990	0.980	1.016	1.013
AR-M-MIDAS	0.645	0.512	0.476	0.562	0.868	0.799	0.960	0.730	1.030	1.147	1.138	1.277	1.076	1.082	1.053	1.063	1.020	1.029	1.129	1.045
AR-U-MIDAS	0.821	0.641	0.466	0.501	1.099	1.194	0.948	0.664	1.087	1.214	1.114	1.143	0.999	1.024	1.075	1.127	1.003	1.017	1.041	1.015
Balanced AR-MIDAS	0.721	0.526	0.360	0.375	1.089	1.111	0.940	0.717	1.189	1.198	1.191	1.191	0.960	0.960	0.960	0.960	0.987	0.988	0.988	0.988
Balanced AR-M-MIDAS	0.774	0.588	0.528	0.562	1.030	1.036	0.884	0.730	1.181	1.206	1.246	1.277	0.998	1.002	1.048	1.063	1.021	1.023	1.038	1.045
Balanced AR-U-MIDAS	0.797	0.592	0.506	0.536	0.945	0.882	0.733	0.654	0.972	0.976	1.093	1.151	0.987	1.018	1.182	1.192	0.976	1.003	1.023	1.037
Monthly and daily data																				
MIDAS	0.779	0.728	0.413	0.348	0.851	0.966	0.796	0.676	0.915	1.009	1.083	1.083	0.965	0.982	0.926	0.922	1.009	1.000	0.946	0.957
Multiplicative MIDAS	0.643	0.535	0.444	0.520	1.173	1.037	0.822	0.737	1.150	1.115	1.297	1.004	0.990	1.078	1.105	1.172	1.051	1.111	1.127	1.233
CF-MIDAS	0.702	0.429	0.387	0.391	0.863	0.880	1.008	0.783	0.968	0.971	1.074	1.134	0.971	0.997	0.946	0.943	1.072	1.053	1.058	1.059
AR-MIDAS	0.840	0.772	0.378	0.341	0.960	0.963	0.796	0.735	0.990	1.033	0.958	1.074	0.944	0.973	0.931	0.923	1.049	1.019	1.066	1.066
AR-M-MIDAS	0.670	0.551	0.419	0.446	1.107	0.794	0.831	0.794	1.205	0.984	1.271	1.218	1.169	1.322	1.432	1.252	1.145	1.086	1.314	1.186

Note: For each quarterly forecast horizon (h), depending on the within-quarter information used, "I" refers to one month information for the economic sentiment indicator and the Dow Jones Euro Stoxx index and no current quarter information for the industrial production, "II" to two months of the economic sentiment indicator and the Dow Jones Euro Stoxx index and one month of the industrial production, "III" to three months of the economic sentiment indicator and the Dow Jones Euro Stoxx index and two months of the industrial production, and, finally, "IV" to full-quarter information for all variables. The figures in bold denote the minimum relative RMSFE for each kind of MIDAS model and each information set. The figures shaded denote the minimum relative RMSFE for each forecast horizon, across information sets.

Chapter 2

Covariate-augmented unit root tests with mixed-frequency data

ABSTRACT

Unit root tests typically suffer from low power in small samples, which results in not rejecting the null hypothesis as often as they should. This paper tries to tackle this issue by assessing whether it is possible to improve the power performance of covariate-augmented unit root tests, namely the ADF family of tests, by exploiting mixed-frequency data. We use the mixed data sampling (MIDAS) approach to deal with mixed-frequency data. The results from a Monte Carlo exercise indicate that mixed-frequency tests have better power performance than low-frequency tests. The gains from exploiting mixed-frequency data are greater for near-integrated variables. An empirical illustration using the US unemployment rate is presented.

JEL Classification: C12, C15, C22.

Keywords: Unit root, Hypothesis testing, Mixed-frequency data.

2.1 Introduction

The importance of unit root testing for modelling and forecasting has been well established since the seminal paper by Granger and Newbold (1974), who showed that applying least squares to non-stationary variables can lead to spurious results. In these circumstances, standard errors are biased, the traditional t -ratio significance test does not have the standard limiting distribution and, hence, the analysis of parameter estimates becomes unreliable (Phillips (1986)).

CHAPTER 2. CADF TESTS WITH MIXED-FREQUENCY DATA

Many studies have focused on unit root tests and their properties (see, for example, Schwert, 1989; Stock, 1986; and Haldrup and Jansson, 2006 for reviews on the topic). The augmented Dickey-Fuller (ADF) test for zero frequency unit roots is the most commonly used procedure, though not necessarily the one with best power performance. This test loses in terms of power to other tests, both asymptotically and in finite samples, especially in the context of near integrated processes.¹ Lower power means that the null hypothesis, being false, is not rejected as often as it should be, leading to the wrong conclusion that the variable is non-stationary.

To overcome this shortcoming, alternative tests exploiting information from covariates have been proposed. In particular, Hansen (1995) generalised the ADF test to include covariates - the CADF test. The intuition is that including a weakly exogenous and stationary variable in the auxiliary test regression may lead to efficiency gains. The performance of covariate-augmented unit root tests depends crucially on the relationship between the variable of interest (dependent variable) and the covariates. The higher the correlation between the variables, the greater the potential power gains. In practice, exploiting these correlations may entail some challenges, especially when the variables involved are sampled at different/mixed frequencies. The typical approach of temporally aggregating high-frequency variables to the same (low) frequency as the variable of interest (e.g., by skip-sampling or computing simple averages) can result in information losses (see Silvestrini and Veredas, 2008 for a survey on temporal aggregation and its implications).

This article puts forward a new class of CADF tests that is able to deal with mixed-frequency data.² We assess the impact of this extension on size and power performances. In particular, the MI(xed) DA(ta) S(ampling) framework is used to deal with mixed-frequency data. Inspired in the distributed lag models, MIDAS

¹The ADF test is a tougher competitor in terms of size. Nevertheless, Perron and Ng (1996) showed that the M-tests originally suggested by Stock (1999) have lower size distortions compared to other unit root tests that are available in the literature. However, for the tests to have good size properties it is essential that an autoregressive spectral density estimator is used as to consistently estimate the long run variance.

²This article focuses on unit root tests with non-stationarity under the null hypothesis. See Jansson (2004) for a unit root test with covariates where the null hypothesis is stationarity.

weighting schemes are very flexible, can be quite parsimonious and are able to account for different frequencies (for a brief overview of the main topics related with MIDAS regressions see, for example, Andreou et al., 2011). To the best of our knowledge, this is the first application of MIDAS to covariate-augmented unit root testing.

This new mixed-frequency test framework is applied to the well-known CADF test proposed by Hansen (1995) and also to the more recent test proposed by Pesavento (2006), which is a modified version of the CADF test, similar to the GLS generalisation of the ADF test in Elliott et al. (1996).

Using a Monte Carlo experiment, we show that mixed-frequency covariate-augmented unit root tests have better power performance than traditional low-frequency tests, while the size performance is similar. Moreover, the performance of mixed-frequency tests improves when variables are near-integrated. These results are robust to the size of the sample, to the lag specification of the test regressions and to different combinations of time frequencies.

The remainder of this article is organised as follows. Section 2.2 summarises the covariate-augmented unit root tests — CADF and CADF-GLS — as they were initially presented, while Section 2.3 describes the mixed-frequency approach to unit root testing. Section 2.4 reports a simulation-based study on the power and size implications of this new approach. Section 2.5 compares the performance of the alternative approaches for testing the presence of a unit root in the US unemployment rate. Finally, Section 2.6 concludes.

2.2 Covariate-augmented unit root tests

This section presents two covariate-augmented unit root tests commonly found in the literature: the CADF test proposed by Hansen (1995) and the CADF-GLS test in Pesavento (2006). For the sake of simplicity, the notation was developed for the case of a single covariate but can be readily extended for multiple covariates.

The common analytical framework is as follows. As in Hansen (1995) and Elliott and Jansson (2003), assume that the variable of interest, Y_t , is the sum of

CHAPTER 2. CADF TESTS WITH MIXED-FREQUENCY DATA

a deterministic component, $d_{Y,t}$, and a stochastic component, $u_{Y,t}$, such as

$$Y_t = d_{Y,t} + u_{Y,t}, \quad (2.1)$$

where the deterministic component equals $d_{Y,t} = 0$, $d_{Y,t} = \beta_{Y,0}$, or $d_{Y,t} = \beta_{Y,0} + \beta_{Y,1}t$, with t denoting a linear trend. Similarly, the stationary covariate series, X_t , can be expressed as

$$X_t = d_{X,t} + u_{X,t}. \quad (2.2)$$

Hence, consider the VAR model formulation

$$\begin{bmatrix} Y_t \\ X_t \end{bmatrix} = d_t + u_t, \quad (2.3)$$

where $d_t = z_t'\beta$, $\beta = [\beta_{Y,0} \ \beta_{X,0} \ \beta_{Y,1} \ \beta_{X,1}]'$,

$$z_t' = \begin{bmatrix} 1 & 0 & t & 0 \\ 0 & 1 & 0 & t \end{bmatrix}, \text{ and } u_t = \begin{bmatrix} u_{Y,t} \\ u_{X,t} \end{bmatrix}.$$

Five different combinations of the deterministic variables can be considered:

Case 1: $\beta_{Y,0} = \beta_{Y,1} = \beta_{X,0} = \beta_{X,1} = 0$;

Case 2: $\beta_{Y,1} = \beta_{X,0} = \beta_{X,1} = 0$;

Case 3: $\beta_{Y,1} = \beta_{X,1} = 0$;

Case 4: $\beta_{X,1} = 0$;

Case 5: No restrictions.

In addition, the stochastic component u_t is expressed as

$$A(L) \begin{bmatrix} (1 - \alpha L)u_{Y,t} \\ u_{X,t} \end{bmatrix} = \begin{bmatrix} e_{Y,t} \\ e_{X,t} \end{bmatrix} \quad (2.4)$$

or,

$$\begin{bmatrix} u_{Y,t} \\ u_{X,t} \end{bmatrix} = \begin{bmatrix} (1 - \alpha L) & 0 \\ 0 & 1 \end{bmatrix}^{-1} A^{-1}(L) \begin{bmatrix} e_{Y,t} \\ e_{X,t} \end{bmatrix}$$

where $A(L)$ is a matrix polynomial of order k in the lag operator L . In the following analysis we assume that:

Assumption 1: The roots of $A(L)$ lie outside the unit circle;

Assumption 2: $u_0, u_{-1}, \dots, u_{-k}$ are $O_p(1)$;

Assumption 3: $E_{t-1}(e_t) = 0, E_{t-1}(e_t e_t') = \Sigma$, where Σ is positive definite, and $\sup_t E \|e_t\|^{2+\kappa} < \infty$, for some $\kappa > 0$,

where E_{t-1} denotes the conditional expectation with respect to e_{t-1}, e_{t-2}, \dots , and Σ can be expressed as

$$\Sigma = \begin{bmatrix} \sigma_{YY} & \sigma_{YX} \\ \sigma_{YX} & \sigma_{XX} \end{bmatrix}. \quad (2.5)$$

Assumption 1 is a standard stationarity condition. Assumption 2 implies that the initial values are asymptotically negligible and Assumption 3 implies that e_t satisfies a functional central limit theorem (Phillips, 1987). Additionally, let $v_t = [(1 - \alpha L)u_{y,t} \ u_{x,t}]'$. Note that $v_t = A(L)^{-1}e_t$, with autocovariance function denoted by $\Gamma(k) = E(v_t v_{t+k}')$. Its spectral density at frequency zero (scaled by 2π), denoted by $\Omega = A(1)^{-1}\Sigma A'(1)^{-1}$, is assumed to be bounded away from zero and can be decomposed as

$$\Omega = \begin{bmatrix} \omega_{YY} & \omega_{YX} \\ \omega_{YX} & \omega_{XX} \end{bmatrix}. \quad (2.6)$$

It is further assumed that

Assumption 4: The autocovariance function of v_t , $\Gamma(k)$, is absolutely summable,

$$\sum_{j=-\infty}^{+\infty} \|\Gamma(k)\| < \infty,$$

where $\|\cdot\|$ is the standard Euclidean norm.

2.2.1 The CADF test

The covariate augmented Dickey Fuller (CADF) test aims at combining the generally good size properties of the ADF test with higher power due to the inclusion of covariates. To see how let us start by mixing equations (2.3) and (2.4) to obtain

$$\begin{bmatrix} \Delta Y_t \\ X_t \end{bmatrix} = z_t' \beta^* + \delta \begin{bmatrix} Y_{t-1} \\ 0 \end{bmatrix} + \begin{bmatrix} v_{Y,t} \\ v_{X,t} \end{bmatrix} \quad (2.7)$$

where $\delta = \alpha - 1$ and $\beta^* = [-\delta\beta_{Y,0} + (1 + \delta)\beta_{Y,1} \quad \beta_{X,0} \quad -\delta\beta_{Y,1} \quad \beta_{X,1}]'$.

Under Assumptions 1 to 4 it is possible to write (Saikkonen, 1991 and Brillinger, 2001)

$$v_{Y,t} = \sum_{j=0}^{\infty} b_j^* v_{X,t-j} + \eta_t \quad (2.8)$$

where η_t is a stationary process with zero mean and spectral density at frequency zero (scaled by 2π) equal to $\omega_{\eta\eta} = \omega_{YY} - \omega_{YX}\omega_{XX}^{-1}\omega_{YX}$.³ Furthermore,

$$E(v_{X,t}\eta_{t+k}) = 0 \quad (2.9)$$

for any $|k| = 0, 1, 2, \dots$, meaning that the right-hand side variables in (2.8) are orthogonal to the regression error. Given that b_j^* is absolute summable, then (2.8)

³For simplicity, it is assumed that the polynomial $b^*(L)$ only includes lags. This is not necessary; see Hansen (1995) for more details.

can be approximated by

$$v_{Y,t} = \sum_{j=0}^k b_j^* v_{X,t-j} + \eta_t = b^*(L)v_{X,t} + \eta_t = b^*(L)(X_t - \beta_{X,0}^* - \beta_{X,1}^*t) + \eta_t \quad (2.10)$$

where k is large enough so that $b_j^* \approx 0$ for $j > k$, as in Remark 2.1 of Chang and Park (2002), and $b^*(L)$ is a lag polynomial of order k . Combining (2.7) and (2.10) yields the regression equation

$$\Delta Y_t = \beta_{Y,0}^* + \beta_{Y,1}^*t + \delta Y_{t-1} + b^*(L)(X_t - \beta_{X,0}^* - \beta_{X,1}^*t) + \eta_t. \quad (2.11)$$

This equation resembles the test regression used in Dickey-Fuller (DF) tests, augmented with the covariate X_t . In practice, η_t can be serially correlated. Similarly to ADF tests, (2.11) can be augmented with lags of the dependent variable, so that the error process is approximately white noise. Letting $\psi(L)\eta_t = \xi_t$, where ξ_t is white noise, the augmented version of (2.11) can be expressed as

$$\psi(L)\Delta Y_t = \mu_0 + \mu_1 t + \delta^* Y_{t-1} + b(L)(X_t - \beta_{X,0}^* - \beta_{X,1}^*t) + \xi_t. \quad (2.12)$$

where $\psi(L)$ is a lag polynomial of order p with all roots lying outside the unit circle, $\mu_0 = \psi(1)\beta_{Y,0}^*$, $\mu_1 = \psi(L)\beta_{Y,1}^*$, $\delta^* = \psi(1)\delta$ and $b(L) = \psi(L)b^*(L)$ is a lag polynomial of order q . Accordingly, $v_{Y,t} = b(L)v_{X,t} + \xi_t$. The long-run covariance matrix between $v_{Y,t}$ and ξ_t is

$$\Phi = \begin{bmatrix} \omega_{YY} & \omega_{Y\xi} \\ \omega_{Y\xi} & \omega_{\xi\xi} \end{bmatrix} \quad (2.13)$$

where $\omega_{\xi\xi} = \omega_{YY} - \omega_{YX}\omega_{XX}^{-1}\omega_{YX}$. Define ρ^2 as the long-run zero frequency squared correlation between $v_{Y,t}$ and ξ_t , which can be expressed as

$$\rho^2 = \frac{\omega_{Y\xi}^2}{\omega_{YY}\omega_{\xi\xi}}, \quad (2.14)$$

and let Q^2 be the ratio of variances,

$$Q^2 = \frac{\omega_{\xi\xi}}{\omega_{Y Y}}. \quad (2.15)$$

In a well-specified dynamic regression, ξ_t is uncorrelated with X_{t-k} for all k , so $\omega_{Y\xi} = \omega_{\xi\xi}$ and $\rho^2 = Q^2$.

Equation (2.12) is the covariate-augmented ADF unit root test regression, denoted by CADF(p, q). The null hypothesis for the presence of a zero frequency unit root in Y_t is $\delta^* = 0$, which is tested against the one-sided alternative hypothesis, $\delta^* < 0$. Following Phillips (1987), the asymptotic theory in Hansen (1995) is based on local-to-unity asymptotics, implying that $\delta^* = c/T$, where T represents the sample size and c is a fixed non-centrality parameter. This means that the null hypothesis holds when $c = 0$ and holds locally for $c < 0$ and $T \rightarrow \infty$. However, as noted by Hansen (1995), in a fixed sample this representation is merely a reparameterization.

The test regression is estimated by OLS. The test statistic is the t -statistic associated with the estimated coefficient of interest ($\hat{\delta}^*$) and is distributed as,

$$t(\hat{\delta}^*) \Rightarrow (c/Q) \left(\int_0^1 (W_1^c)^2 \right)^{1/2} + \rho \frac{\int_0^1 W_1^c dW}{\left(\int_0^1 (W_1^c)^2 \right)^{1/2}} + (1 - \rho^2)^{1/2} N(0, 1) \quad (2.16)$$

where W_1^c is an Ornstein-Uhlenbeck process generated by a stochastic differential equation, such as $dW_1^c(r) = cW_1^c + dW(r)$, and W is a standard Brownian motion and the $N(0, 1)$ variable is independent of W .⁴ Under the null hypothesis, the asymptotic distribution of the t -statistic is a convex linear combination of the Dickey-Fuller (DF) distribution of the univariate unit root tests and the standard Normal

$$t(\hat{\delta}^*) \Rightarrow \rho \frac{\int_0^1 W dW}{\left(\int_0^1 W^2 \right)^{1/2}} + (1 - \rho^2)^{1/2} N(0, 1) \quad (2.17)$$

⁴The case above presented does not include deterministic variables. When extending to the cases where these variables are included, the structure of (2.16) remains unchanged except that W_1^c would be appropriately replaced; for more details see Hansen (1995).

where the weights are determined by the nuisance parameter ρ^2 . The parameter ρ^2 (or Q^2) can be interpreted as the relative contribution of ξ_t to explain $v_{Y,t}$ at the zero frequency. On the one hand, if $b(L)$ equals zero, $v_{Y,t} = \xi_t$ and $\rho^2 = 1$. In this case the CADF test is equivalent to the typical ADF test. On the other hand, if the importance of ξ_t to explain the zero-frequency movements in $v_{Y,t}$ decreases, then $\rho^2 \rightarrow 0$ and the relevant distribution becomes closer to the Normal distribution.⁵

Perhaps more intuitively, one can define a third measure, R^2 , such that,

$$R^2 = 1 - Q^2 = \frac{\omega_{YX}\omega_{XX}^{-1}\omega_{YX}}{\omega_{YY}} \quad (2.18)$$

which accounts for the relative contribution of regressor X_t to explain $v_{Y,t}$ at the zero frequency. If X_t does not contribute at all to explain the variation in $v_{Y,t}$, then $R^2 = 0$. Conversely, if X_t has an increasing contribution to explain $v_{Y,t}$, then $R^2 \rightarrow 1$.

Given that the distribution of the test statistic depends on ρ^2 , a consistent estimate of this parameter is needed, in order to select the appropriate critical value. According to Hansen (1995), an estimate ($\hat{\rho}^2$) can be obtained in a non-parametric way from

$$\hat{\rho}^2 = \frac{\hat{\omega}_{Y\xi}^2}{\hat{\omega}_{YY}\hat{\omega}_{\xi\xi}} \quad (2.19)$$

where

$$\hat{\Phi} = \begin{bmatrix} \hat{\omega}_{YY} & \hat{\omega}_{Y\xi} \\ \hat{\omega}_{Y\xi} & \hat{\omega}_{\xi\xi} \end{bmatrix} = \sum_{k=-M}^M w(k/M) \frac{1}{T} \sum_{t=1}^{T-m} \hat{\pi}_t \hat{\pi}'_{t+k} \quad (2.20)$$

and $\hat{\pi}_t = (\hat{v}_{Y,t}, \hat{\xi}_t)'$ are least squares estimates of $\pi_t = (v_{Y,t}, \xi_t)'$ from the appropriate regression model. For example, assuming no intercepts, $\hat{\xi}_t = \hat{\psi}(L)\Delta Y_t - \hat{\delta}^* Y_{t-1} - \hat{b}(L)X_t$ and $\hat{v}_{Y,t} = \hat{b}(L)X_t + \hat{\xi}_t$. The function $w(\cdot)$ is a kernel weight function, such as the Bartlett or Parzen kernels, and M is the bandwidth selected to grow slowly with sample size (Andrews, 1991, Jansson, 2002). Table 1 in Hansen

⁵The case of $\rho^2 = 0$ is excluded, ruling out the situation where the variable of interest is cointegrated with the cumulated stationary covariate (Lupi, 2009).

(1995) presents the relevant asymptotic critical values for a range of ρ^2 values and is made available in Table B.1 of the Appendix.

In theory, the power of unit root tests can be improved by the inclusion of covariates because these contribute to reduce the standard error of the estimate of the autoregressive parameter. Given that δ^* is estimated more precisely, the unit root test for the null hypothesis $H_0 : \delta^* = 0$ will have more power. Greater reductions in the standard error, i.e. higher power, are associated with lower ρ^2 (higher R^2). As analytically shown by Caporale and Pittis (1999), if there is contemporaneous correlation between $v_{Y,t}$ and X_t and Granger causality from $v_{Y,t}$ to X_t , then, in some cases, adding covariates may also lead to an increase in the absolute value of the parameter estimate itself, further enhancing the power of the unit root test.

The discussion above is based on the assumption that the covariates are stationary. The CADF test is no longer valid if X_t is integrated. In case of doubt about the stationarity of the covariates, one should take first differences of these series before proceeding into testing. As discussed in Hansen (1995), this seems to be a sensible approach because over-differencing results in neither significant size distortions nor power loss.

In practice, the presence of correlation between the variable of interest and the covariates, as well as its nature, matter for the performance of the test. In its empirical application, Hansen (1995) concluded that there are important power gains to be obtained from using the CADF test to assess the stationarity of real GNP per capita, industrial production and the unemployment rate for the US. Nevertheless, the CADF test is more prone to size distortions than the ADF test. Caporale and Pittis (1999) performed a similar exercise, analysing a wider set of US macroeconomic series. The authors concluded that the finding of a unit root does not always hold when the more powerful CADF test is used instead of the standard ADF method, although there is evidence of high persistence.

2.2.2 Tests with GLS demeaning

Recognising that the difficulties with the traditional univariate unit root tests (namely, DF and ADF tests) are associated with inefficient estimates of the deter-

ministic component, Elliott et al. (1996) suggested that modifying the estimation of this component could improve their performance. For this purpose, the authors suggested GLS-demeaning/detrending the variable of interest prior to testing for the presence of unit roots (DF-GLS test). Pesavento (2006) proposed a generalisation of the DF-GLS test to include stationary covariates, the so-called CADF-GLS test.

The DF-GLS test

In brief, the DF-GLS test is similar to the ADF-test. The aim of the DF-GLS test is to assess whether $\delta = 0$ (null hypothesis) against the point alternative that $\delta = c/T < 0$. However, in the test regression, instead of using the original Y_t series, the GLS-demeaned/detrended version (Y_t^d) is used. The GLS-demeaned/detrended Y_t^d series is obtained as $Y_t - z_t' \hat{\beta}$ and $\hat{\beta}$ are the coefficient estimates from regressing $Y(\bar{\alpha})$ on $z(\bar{\alpha})$, which are transformed versions of the dependent variable and deterministic variables, respectively. More precisely, $Y(\bar{\alpha}) = (Y_1, Y_2 - \bar{\alpha}Y_1, \dots, Y_T - \bar{\alpha}Y_{T-1})$, $z(\bar{\alpha}) = (z_1, z_2 - \bar{\alpha}z_1, \dots, z_T - \bar{\alpha}z_{T-1})$ and $\bar{\alpha} = 1 + \bar{c}/T$.

The literature shows that values of \bar{c} associated with an asymptotic power of one half yield tests with power functions tangent to the power envelope at that value, and close to the power envelope over a considerable range of alternative values. The appropriate value of \bar{c} depends on the deterministic specification. Simulation results in Elliott et al. (1996) suggest that \bar{c} should equal -13.5 if a trend is included (meaning that when $\bar{c} = -13.5$ the point optimal test is tangent to the power envelope at 0.5 if a constant and trend are estimated), or -7 for constant only (i.e., when \bar{c} equals -7 the point optimal test is tangent to the power envelope at 0.5).

The DF-GLS test statistic is the t -statistic for testing whether $\delta = 0$ in the following regression, without deterministic regressors

$$\Delta Y_t^d = \delta Y_{t-1}^d + a_1 \Delta Y_{t-1}^d + \dots + a_p \Delta Y_{t-p}^d + e_t \quad (2.21)$$

where Δ denotes the first difference, p is the number of lags and e_t is an error term. For constant only, the critical values are those of the conventional DF-tests,

when there is no intercept. In the linear trend case, the critical values can be found in Table 1 in Elliott et al. (1996).

A related issue concerns the choice of the values assigned to the first observation in the GLS demeaning/detrending procedure. Elliott et al. (1996) considered the first observation of the quasi-differenced series as being equal to the first observation in levels (fixed initial observation assumption). Elliott (1999) extended this framework to the case where the initial observation is drawn from its unconditional distribution under the alternative hypothesis. The author concluded that there are differences between the two approaches, but none is the best; the user's choice will depend on the his/her belief as to the correct alternative to be tested. More recently, Westerlund (2015) also assessed the importance of the hypothesis about the first observation in GLS demeaning/detrending. He compared the fixed initial observation assumption (i.e., the first quasi-difference equals the first level) with simply ignoring the first quasi-difference (i.e., equals zero). His results suggest that choosing between these two alternatives matters, and the first observation does not seem to be negligible. Moreover, the first assumption seems to work better.

The CADF-GLS test

Merging Hansen's approach with the GLS demeaning/detrending used in Elliott et al. (1996), the CADF-GLS test is constructed by demeaning/detrending each variable (dependent variable and the covariate) according to the assumptions on the deterministic terms, and then estimating a test regression similar to equation (2.12) but with the demeaned/detrended variables.⁶

Hence, the CADF-GLS test statistic is the t -statistic for testing whether $\delta = 0$ against the point alternative that $\delta = c/T < 0$ in the following regression, without deterministic regressors

$$\Delta Y_t^d = \delta Y_{t-1}^d + \sum_{j=1}^p a_j \Delta Y_{t-j}^d + \sum_{j=0}^q b_j X_{t-j}^{d*} + e_t \quad (2.22)$$

where Y_t^d and X_t^{d*} are the demeaned/detrended versions of the original series,

⁶Adding to lagged terms, leads can also be included. See Pesavento (2006) for more details.

p and q are the respective number of lags (chosen by an information criterion, such as the BIC) and e_t is an error term. The Y_t^d series is obtained as described in the previous section. Following Pesavento (2006) and Christopoulos and León-Ledesma (2008), given that X_t is assumed to be stationary, X_t^{d*} is obtained by OLS demeaning/detrending the original X_t series, depending on the deterministic component included.

The asymptotic test distribution is

$$t(\hat{\delta})^{GLS} \Rightarrow c \left(\int_0^1 J^2 \right)^{1/2} + \frac{\int_0^1 J dW}{\left(\int_0^1 J^2 \right)^{1/2}} \quad (2.23)$$

where J is a Ornstein-Uhlenbeck process, such that

$$J(r) = W(r) + c \int_0^1 e^{(\lambda-s)c} W(s) ds \quad (2.24)$$

where $W(r) = \sqrt{\frac{R^2}{1-R^2}} W_x(r) + W_y(r)$, $\lambda = (1-c)/(1-c+c^2/3)$, and W_x and W_y are independent standard Brownian motions.⁷ In the case of no deterministic variables or only a constant for the dependent variable, (2.23) is equivalent to (2.16) and the critical values are those of the CADF test, when there is no intercept. For the other cases, Pesavento (2006) reported asymptotic critical values for a significance level of 5 per cent and different values of R^2 (from 0 to 0.9) that can be found in Table B.2 of the Appendix. The author also refers that she used \bar{c} equal to -7 for cases 1 to 3 and -13.5 for cases 4 and 5, in order to make a reasonable comparison with previous work. The estimate of R^2 is obtained non-parametrically, as in Hansen (1995).

⁷The case above presented does not include deterministic variables. See Pesavento (2006) for extending to the cases where these variables are included.

2.3 Mixed-frequency covariate-augmented unit root tests

The performance of covariate-augmented unit root tests depends crucially on the relationship between the variable of interest and the covariate used. Economic theory may help in the choice of covariates for unit root testing. However, empirical evidence is not always clear-cut.

In some cases, this may result from ignoring the fact that time series are often analysed at intervals that reflect the timing of data collection. Economic data are sampled at different frequencies. Suppose that the variable of interest and the covariate are described by a high-frequency data generating process (DGP) at a given frequency (e.g., monthly). Moreover, assume that only the covariate is available at that frequency, while the dependent variable is observed at a lower frequency (e.g., annually or quarterly). To deal with this situation the variables are typically temporally aggregated to the same (low) frequency by skip-sampling or averaging.

Regarding the impact of testing for unit roots in the feasible low frequency version of the dependent variable, Granger and Siklos (1995) and Marcellino (1999), among other, showed that zero frequency unit roots are not affected by temporal aggregation. However, temporal aggregation may cause deviations from the test distributions. Assume that a high-frequency variable y_t has the following DGP

$$y_t = \alpha y_{t-1/m} + u_t, \quad (2.25)$$

where $u_t \sim iid(0, \sigma^2)$ and $E(u_t u_{t-i}) = 0$. Under the null hypothesis of a unit root,

$$\Delta^{1/m} y_t = u_t, \quad (2.26)$$

where m is the sampling frequency and $\Delta^{1/m}$ is the high-frequency difference operator. For simplicity, assume that $m = 2$. The first difference of the aggregate

variable y_t^a resumes to

$$\Delta y_t^a = \begin{bmatrix} \phi_1 & \phi_2 \end{bmatrix} \begin{bmatrix} \Delta y_t \\ \Delta y_{t-1/2} \end{bmatrix} = \begin{bmatrix} \phi_1 & \phi_2 \end{bmatrix} \left(\begin{bmatrix} u_t \\ u_{t-1/2} \end{bmatrix} + \begin{bmatrix} u_{t-1/2} \\ u_{t-1} \end{bmatrix} \right), \quad (2.27)$$

where Δ is the low-frequency difference operator and ϕ_i for $i = 1, \dots, m$ represents the aggregation scheme. Now, consider the first-order autocovariance of the aggregated errors under the null hypothesis, denoted as u_t^a . Given that u_t is serially uncorrelated, the autocovariance of the aggregate only contains the product of the terms with the same time subscript. Hence,

$$\text{cov}(u_t^a) = \phi_1 \phi_2 E(u_{t-1}^2) = \phi_1 \phi_2 \sigma^2. \quad (2.28)$$

As noted by Working (1960), if the aggregation scheme is skip-sampling ($\phi_1 = 1$ and $\phi_2 = 0$ for end-of-period sampling or $\phi_1 = 0$ and $\phi_2 = 1$ for beginning-of-period sampling) then u_t^a is not serially correlated.

When the aggregation scheme is some kind of averaging (e.g., flat sampling, with $\phi_1 = \phi_2 = 1/2$), Working (1960) showed that u_t^a is serially correlated. In this case, the serial correlation affects the limiting distribution for testing unit roots in the aggregate variable. This serial correlation cannot be controlled for by adding lagged terms. Following the same reasoning as in Ghysels and Miller (2013), a first-order moving average (MA) polynomial can be written as

$$\begin{aligned} \begin{bmatrix} \Delta y_t \\ \Delta y_{t-1/2} \end{bmatrix} &= \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} u_t \\ u_{t-1/2} \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 1 & -1 \end{bmatrix} \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} u_{t-1} \\ u_{t-3/2} \end{bmatrix} \\ &= u_t^* + \begin{bmatrix} 0 & 0 \\ 1 & -1 \end{bmatrix} u_{t-1}^* \\ &= u_t^* + A u_{t-1}^*. \end{aligned} \quad (2.29)$$

This MA polynomial is not invertible because the matrix A does not fulfil the condition $\det(I + Az) \neq 0$ when $|z| \leq 1$. Therefore, this polynomial cannot be well approximated by a high-order autoregressive polynomial.

Regarding the role of the temporally aggregated covariate for enhancing the

power of the unit root tests, note that the above-mentioned caveats of temporal aggregation also apply. Focusing on the skip-sampling case (for both the variable of interest and the covariate), the assumption in (2.8) ensures that in the aggregate equation the regressors are also orthogonal to the regression error.

However, temporal aggregation entails information losses and may reduce the contribution of the covariate to explain the variability of the variable of interest. Assuming that temporal aggregation of the variable of interest is inevitable, this article contributes to the literature by proposing a unit root test of the CADF family that is able to deal with mixed-frequency data. In particular, we assess whether the mixed-frequency approach contributes to the waning of potential distortions in the correlation between the dependent variable and the covariate generated by temporal aggregation. To deal with mixed-frequency data we use the MIDAS framework, which we briefly describe in the next section.

2.3.1 The MIDAS approach

Introduced by Ghysels et al. (2004) and used in, e.g., Ghysels et al. (2006) and Ghysels et al. (2007), the MIDAS approach provides simple, reduced-form models to approximate more elaborate, though unknown, high-frequency models.

Consider a low-frequency variable Y_t and a high-frequency variable x_t , which has a time frequency m times higher. MIDAS regressions assume that the coefficients associated with the high-frequency variable and its lags are captured by an aggregation lag polynomial $B(L^{1/m})$

$$Y_t = \mu + B(L^{1/m})x_t^{(m)} + u_t \quad (2.30)$$

where μ is a constant, $x_t^{(m)}$ is the skip-sampled version of the high-frequency x_t , $B(L^{1/m}) = \sum_{j=0}^J B(j) L^{j/m}$ is a polynomial of length J in the $L^{1/m}$ lag operator, i.e., $L^{j/m}x_t^{(m)} = x_{t-j/m}^{(m)}$, $B(j)$ is an aggregation weighting scheme, and u_t is a standard *iid* error term. The index j indicates how many high-frequency periods starting from the end of the low-frequency period are taken into account. Note that $B(L^{1/m})x_t^{(m)}$ can be interpreted as a temporally aggregated variable using a more flexible, data-driven weighting scheme compared with commonly used temporal

aggregation schemes, such as skip-sampling or averaging. For more details on MIDAS regressions, see Andreou et al. (2013), among others.

One crucial assumption about the covariate regards its stationarity. In this case, if x_t is a stationary variable, then $x_t^{(m)}$, its skip-sampled version (or, X_t , the general low-frequency aggregate) also is, as discussed above. Notice that we are assuming that x_t does not display a seasonal behaviour, in order to exclude the possibility that the unit root at the zero frequency can arise because of temporal aggregation of a series which has a unit root at some seasonal frequency (Granger and Siklos, 1995).

Regarding the weighting function, there are several possible choices. Ghysels et al. (2007) considered two alternatives, both assuming that the weights are determined by a few hyperparameters: the exponential Almon lag and the beta polynomial. Given that these options have nonlinear functional specifications, in both cases MIDAS regressions are estimated using nonlinear least squares.

Alternatively, there is the aggregation scheme underlying the unrestricted MIDAS regressions (U-MIDAS), used in Marcellino and Schumacher (2010), Foroni and Marcellino (2012) and, Foroni et al. (2011)

$$Y_t = \mu + B_U(L^{1/m})x_t^{(m)} + u_t, \quad (2.31)$$

where $B_U(L^{1/m}) = \sum_{j=0}^J B_j L^{j/m}$.

Equation (2.31) can involve a large number of parameters, namely when the difference between the low and the high frequency is large. Hence, large differences in sampling frequencies between the variables are readily penalised in terms of parsimony in U-MIDAS regressions.

This article focuses on a parameterised weighting scheme that can be estimated by OLS, namely the traditional Almon lag polynomial. This aggregation scheme assumes that J lag weights can be related to d linearly estimable underlying parameters, with $d < J$, as follows:

$$B_A(j) = \sum_{i=0}^d \theta_i j^i, \quad j = 1, \dots, J \quad (2.32)$$

where θ_i , $i = 0, \dots, d$, denotes the hyperparameters. In the following analysis it is assumed that $d = 2$.⁸

Following the notation in Section 2.2, the next two sections describe how MIDAS regressions were used to extend covariate-augmented unit root tests to mixed frequency data: Section 2.3.2 for the mixed-frequency CADF test (M-CADF); and Section 2.3.3 for the mixed-frequency CADF test with GLS detrending (M-CADF-GLS).

2.3.2 The M-CADF test

The CADF test can be extended to account for mixed-frequency data as follows

$$\psi(L)\Delta Y_t = d_{Y,t} + \delta Y_{t-1} + B(L^{1/m})(x_t^{(m)} - \beta_{x,0} - \beta_{x,1}t) + \xi_t \quad (2.33)$$

where $d_{Y,t}$ represents the deterministic component of the dependent variable and ξ_t is white noise. The M-CADF(p, J) test assesses the null hypothesis for the presence of a zero frequency unit root in Y_t ($\delta = 0$), against the alternative hypothesis that Y_t is stationary. Notice that the test regression in (2.33) simply consists in plugging in high-frequency lags of the covariate in the CADF test regression in (2.12), instead of the low-frequency lags already included. Considering the Almon MIDAS regression, (2.33) is estimated by OLS. As in the original CADF test, the test statistic is the t -statistic associated with the estimated $\hat{\delta}$ coefficient and the distribution of the test is as in (2.17).

2.3.3 The M-CADF-GLS test

Similarly to (2.22), the test regression of the M-CADF-GLS(p, J) test is

$$\Delta Y_t^d = \delta Y_{t-1}^d + \sum_{j=1}^p a_j \Delta Y_{t-j}^d + B(L^{1/m})x_t^{(m),d*} + e_t \quad (2.34)$$

where Y_t^d and $x_t^{(m),d*}$ are the demeaned/detrended versions of the original se-

⁸This weighting scheme also works in the cases where m is not fixed (e.g., combining monthly with weekly or daily data). In these cases, instead of having one set of weights, we have a different set of weights for each low-frequency period of the sample.

ries, as in Section 2.2.2, p is the number of autoregressive lags, $B(L^{1/m})$ is a lag polynomial of order J and e_t is white noise. As in the previous section, the GLS version of the M-CADF test also resumes to replacing the low-frequency lags by the high-frequency lags of the covariate. The M-CADF-GLS(p, j) test statistic is the t -statistic for testing whether $\delta = 0$ against the point alternative that $\delta = c/T < 0$ and the distribution of the test is as in (2.23). For improving the comparability with the existing literature, the figures for \bar{c} used in this article are -7 for cases 1 to 3 and -13.5 for cases 4 and 5. In addition, the asymptotical critical values are the same as in Pesavento (2006).

2.4 Monte Carlo simulation

The finite sample size and power performance of the proposed mixed-frequency versions of the covariate-augmented unit root tests is investigated by means of a Monte Carlo simulation exercise. Inspired by Hansen (1995) and Galvao Jr. (2009), this exercise considers the following DGP

$$\begin{aligned} y_t &= d_{y,t} + \alpha y_{t-1} + v_{y,t} \\ x_t &= d_{x,t} + v_{x,t} \end{aligned} \tag{2.35}$$

where y_t and x_t are both in the same (high) time frequency, $\alpha = 1 + c/T$ and d_t represents the deterministic terms. Four alternatives for c were considered, namely 0, -5 , -10 and -15 . Regarding the deterministic terms, the leading cases 3 (constant for both variables) and 5 (constant and time trend for both variables) were considered; see Elliott and Jansson (2003) and Juhl and Xiao (2003). The error process $v_t = [v_{y,t} v_{x,t}]'$ is generated by a VARMA model $A(L)v_t = B(L)\xi_t$, where $A(L) = I_2 - AL$, $B(L) = I_2 + BL$,

$$A = \begin{bmatrix} a_1 & a_2 \\ a_2 & a_1 \end{bmatrix}, \quad B = \begin{bmatrix} b_1 & b_2 \\ b_2 & b_1 \end{bmatrix}, \tag{2.36}$$

and $\xi_t = [\xi_{y,t} \xi_{x,t}]' \sim N(0, \Sigma)$, where I_2 is a (2×2) identity matrix and Σ is

such that the long-run variance matrix of v_t satisfies

$$\Omega = (I_2 - AL)^{-1}(I_2 + BL)\Sigma(I_2 + BL)'(I_2 - AL)^{-1'} = \begin{bmatrix} 1 & R \\ R & 1 \end{bmatrix} \quad (2.37)$$

and, as before, $R^2 \in [0, 1[$. Various values for R^2 are examined, namely $R^2 = \{0.2, 0.5, 0.8\}$. Note that this DGP does not include seasonal features.

Regarding the treatment of the initial condition of the DGP, Hansen (1995) and Galvao Jr. (2009) dropped the first 100 observations to eliminate the start-up effects. However, Müller and Elliott (2003) noted that different initial conditions lead to dramatic changes in the power of unit root tests. In particular, the GLS-family of unit root tests has the best performance when the initial value is near zero. In this article, the initial condition was set at zero.⁹

The low-frequency series Y_t and X_t are obtained by aggregating the generated high-frequency data y_t and x_t . The aggregating scheme considered is skip-sampling (as for stock variables) and $m = 3$, mimicking the combination of quarterly and monthly data. The sample size is set at $T = 100$, and 10,000 replications are used. In the following sections we present a sensitivity analysis to different m and a larger T , namely $T = 500$.

Notice that the aggregation process affects the value of α tested in the alternative hypothesis. As shown in Pierse and Snell (1995), let y_t be a variable generated by the following first-order process

$$y_t = \alpha y_{t-1} + u_t.$$

Then, the m -period aggregated variable denoted as Y_t^a is given by

$$Y_t^a = \alpha^a Y_{t-1}^a + u_t^a,$$

where $\alpha^a = \alpha^m$, whether y_t is a flow or a stock. Hence, the alternative values

⁹A sensitivity analysis to this assumption was performed and the results indicate that the main result — better power performance of mixed-frequency covariate-augmented unit root tests — does not qualitatively change if a different initial value was considered (e.g., by dropping the initial 100 observations), though the relative performance of the GLS-family of tests is significantly affected by this choice.

of α go from 0.95, 0.90 and 0.85, for the high-frequency process, to 0.86, 0.73 and 0.61, respectively, for the low-frequency process. As discussed in Section 2.3, aggregation also affects the correlation between the variable of interest and the covariate and, thus, affects R^2 , though it is much harder to predict the actual impact.

The number of lags is assumed unknown, replicating what happens in practice. The choice of the number of lags is very important for the performance of the test. Choosing a lag order is crucial to find a good enough approximation to the true DGP, which yields unit root tests with size close to the nominal size while retaining acceptable power. The choice of the number of lags is particularly important in the case of (negative) moving average errors.

A common result in the literature is that estimating the number of lags solely by applying the AIC or BIC in a VAR model under the null leads to a very conservative number of lags, which results in noticeable size distortions. An alternative method is the sequential t -test for the significance of the last lag considered, as in Ng and Perron (1995). This procedure has the ability to yield a higher number of lags than the BIC when there are negative moving-average errors and, hence, reduce size distortions. But, the sequential procedure tends to overparameterize in other cases, which also leads to less efficient estimates and subsequently to power losses. Hansen (1995) used ad-hoc rules to choose the number of lags and Pesavento (2006) suggested applying the MAIC approach proposed by Ng and Perron (2001) to an univariate regression (an ADF-type regression) with GLS detrended series for choosing the relevant number of lags for the Y_t variable and, then, using the same number of lags for the covariate.¹⁰

Fossati (2012) analysed the size and power performance of covariate-augmented unit root tests for different selection procedures of truncation lags. The author showed that the approach in Pesavento (2006) could lead to including too many unnecessary lags and, thus, to size distortions. He suggested applying the MAIC to the output of a CADF test regression and dropped the restriction of using the same number of lags for the Y_t variable and the covariate. Moreover, he considered two alternatives within this unrestricted framework: one with the maximum number

¹⁰Ng and Perron (2001) also consider a modified version of the BIC, denoted as MBIC, but discounted this alternative due to the superior properties of the MAIC.

of lags given by the rule $\text{int}(12(T/100)^{0.25})$, where T is the number of observations (Schwert, 1989); and another where the maximum number of lags is selected using the procedure in Pesavento (2006). The first has lower size distortions — leading to tests with an almost exact size — but worse power performance than the second, which is not too far off from the results for the restricted version in Pesavento (2006).

In this article we follow the approach suggested by Fossati (2012). However, instead of using the MAIC proposed by Ng and Perron (2001), we use the multivariate version of MAIC in Perron and Qu (2007). The maximum number of lags allowed for the dependent variable was chosen according to the rule in Schwert (1989) and for the covariate we assumed a maximum number of lags equal to 5. This restriction reduces considerably the computation time while not affecting the results. In order to have a meaningful comparison, the same number of lags is used for all tests.¹¹

Table 2.4.1 shows the range of simulation designs, as well as the number of lags chosen, under the null hypothesis, for each DGP. The values for the initial R^2 refer to the high-frequency processes. Different values were used for the persistence in the high-frequency autoregressive dynamics of the dependent variable, namely $a_1 = \{0.2, 0.5, 0.8\}$, and for the high-frequency moving average dynamics, $b_1 = \{-0.2, -0.5, -0.8\}$.¹² The criterion used to select the number of lags seems to work well, namely in the case of moving average dynamics (DGP 16 to 24), for which the number of lags for the dependent variable is higher than in the case of autoregressive dynamics.

Heteroskedasticity and autocorrelation consistent (HAC) estimates of the elements of the long-run variance-covariance matrix Ω are used. This calculation commonly involves the use of pre-whitening filters based on simple autoregressive models. This procedure may induce bias in the estimation of autoregressive coef-

¹¹The number of lags was chosen with the low frequency dataset. For the mixed-frequency tests the same time span of lagged information for the covariate is covered, which corresponds to a different number of lags in the high time frequency.

¹²Only negative figures were considered for the moving average dynamics because they represent the most difficult case in terms of size distortions, as the reversion to the mean is higher. Regarding autoregressive dynamics, other DGP were tested, namely with a higher persistence in the covariate. The results remain qualitatively unchanged. So, for the sake of brevity, those results will not be reported, but are available from the author upon request.

CHAPTER 2. CADF TESTS WITH MIXED-FREQUENCY DATA

Table 2.4.1: Simulation design and median of the lag order selected by MAIC (Perron and Qu, 2007)

	Simulation design					Median of the lag order							
	Error structure				Initial R^2	$T = 100$				$T = 500$			
	a_1	a_2	b_1	b_2		Constant		Trend		Constant		Trend	
					Y_t	X_t	Y_t	X_t	Y_t	X_t	Y_t	X_t	
1	0	0	0	0	0.20	0	0	1	0	0	0	1	0
2	0	0	0	0	0.50	0	0	1	0	0	0	1	0
3	0	0	0	0	0.80	0	0	1	0	0	0	1	0
4	0.2	0	0	0	0.20	0	0	0	0	1	0	0	0
5	0.2	0	0	0	0.50	0	0	0	0	1	0	0	0
6	0.2	0	0	0	0.80	0	0	0	0	1	1	0	1
7	0.5	0	0	0	0.20	1	0	0	0	2	0	2	0
8	0.5	0	0	0	0.50	1	0	0	0	2	1	2	1
9	0.5	0	0	0	0.80	1	1	0	1	2	1	2	1
10	0.8	0	0	0	0.20	2	0	2	0	2	2	2	2
11	0.8	0	0	0	0.50	2	1	2	1	2	2	2	2
12	0.8	0	0	0	0.80	2	1	2	1	2	2	2	2
13	0.2	0.2	0	0	0.20	0	0	0	0	1	1	0	1
14	0.2	0.2	0	0	0.50	0	1	0	0	1	1	1	1
15	0.2	0.2	0	0	0.80	0	1	0	1	1	1	1	1
16	0	0	-0.2	0	0.20	1	0	1	0	1	0	2	0
17	0	0	-0.2	0	0.50	1	0	1	0	1	1	2	1
18	0	0	-0.2	0	0.80	1	0	1	0	1	1	2	1
19	0	0	-0.5	0	0.20	2	0	3	0	3	1	4	1
20	0	0	-0.5	0	0.50	2	1	3	1	3	1	4	1
21	0	0	-0.5	0	0.80	2	1	3	1	3	1	4	1
22	0	0	-0.8	0	0.20	6	1	6	1	8	1	10	1
23	0	0	-0.8	0	0.50	6	1	6	1	8	2	10	2
24	0	0	-0.8	0	0.80	6	1	6	2	8	2	10	2

ficients, which is transmitted to the recolouring filter. To mitigate the potential bias associated with these filters, recursive demeaning/detrending procedures were assessed, as in Taylor (2002), Sul et al. (2005) and Rodrigues (2006).

In order to implement the unit root tests we use finite-sample critical values. For $\alpha = 1$, the observed rejection rates of each test were based on critical values from the limiting distribution obtained for the DGP with the simplest dynamics, i.e., DGP 1, 2 and 3 in Table 2.4.1. For $\alpha < 1$ the size-adjusted power of the tests was based on critical values estimated from the simulated data generated under the null hypothesis ($\alpha = 1$).¹³ Following the suggestion in Elliott and Jansson (2003), the critical values were interpolated for the estimated figures of R^2 .

2.4.1 Baseline results

Table 2.4.2 reports the probability of rejecting the null hypothesis under the unit root case, i.e., the finite sample size for unit root tests considering nominal size 5 per cent.¹⁴ Overall, size distortions are larger when there is a linear trend in the regression. The same occurs when b_1 is nonzero, i.e., in the presence of a negative moving average root, which is a result that is commonly found in the unit root literature; see, among others, Schwert (1989). Downward distortions are mainly for stronger and more complex autoregressive dynamics, as also reported in Hansen (1995) and Galvão (2013).

In most cases the size differences between the two sets of tests — mixed- and low-frequency — are not substantial. Figure 2.1 shows the difference between the finite sample size of the unit root tests and the nominal size of 5 per cent for each DGP.¹⁵ When downward distortions exist, they tend to be less marked for mixed-frequency tests. In the case of strong negative moving average dynamics the upward size distortions are also smaller for mixed-frequency tests than for the low-frequency ones. However, when the moving average parameter is smaller (in absolute terms) the opposite happens.

¹³As mentioned by Haug (2002), size-unadjusted power is rather misleading, so those results are not reported.

¹⁴The codes were written in Matlab. Some functions were taken from the Econometrics Toolbox by James P. LeSage (<http://www.spatial-econometrics.com>). The procedure to perform the CADF unit root tests was greatly inspired in the code made available by Bruce E. Hansen (http://www.ssc.wisc.edu/~bhansen/progs/et_95.html). The MIDAS toolbox was inspired in a code kindly provided by Arthur Sinko.

¹⁵The figure shows the figures for the case of only a constant term; results are similar if also a time trend is included.

CHAPTER 2. CADF TESTS WITH MIXED-FREQUENCY DATA

Table 2.4.2: Finite sample size for unit root tests considering nominal size of 5 per cent, $T = 100$

	Error structure					Constant				Trend			
	a_1	a_2	b_1	b_2	Initial R^2	CADF	CADF GLS	M CADF	M CADF GLS	CADF	CADF GLS	M CADF	M CADF GLS
1	0	0	0	0	0.20	5.2	5.1	5.1	5.2	5.4	5.3	5.3	5.3
2	0	0	0	0	0.50	5.1	5.0	5.2	5.2	4.9	5.1	5.1	5.0
3	0	0	0	0	0.80	4.7	4.9	4.8	4.8	4.6	4.7	4.4	4.9
4	0.2	0	0	0	0.20	4.7	4.5	4.8	4.6	4.2	4.1	4.5	4.3
5	0.2	0	0	0	0.50	4.2	4.0	4.7	4.6	3.5	3.7	4.3	4.0
6	0.2	0	0	0	0.80	3.3	3.4	3.9	4.2	2.7	2.6	3.0	3.6
7	0.5	0	0	0	0.20	4.0	2.5	4.8	3.7	2.4	1.6	3.4	2.7
8	0.5	0	0	0	0.50	2.8	2.4	5.6	4.7	1.5	1.0	4.5	4.2
9	0.5	0	0	0	0.80	1.5	1.5	4.9	4.6	0.5	0.5	4.3	4.3
10	0.8	0	0	0	0.20	4.2	3.0	4.7	3.4	2.6	1.3	3.6	2.1
11	0.8	0	0	0	0.50	4.1	3.5	6.1	5.1	2.5	2.1	6.1	4.8
12	0.8	0	0	0	0.80	3.0	3.2	7.8	6.7	1.4	2.0	7.9	7.6
13	0.2	0.2	0	0	0.20	4.0	3.8	4.3	4.1	3.6	3.3	4.2	4.0
14	0.2	0.2	0	0	0.50	3.2	3.0	3.6	4.0	2.6	2.1	3.8	3.6
15	0.2	0.2	0	0	0.80	2.0	2.1	2.8	3.7	1.1	1.0	2.7	2.7
16	0	0	-0.2	0	0.20	5.6	5.4	5.6	5.9	6.2	6.0	6.3	6.5
17	0	0	-0.2	0	0.50	5.6	5.2	6.5	6.1	6.0	5.9	6.7	7.0
18	0	0	-0.2	0	0.80	5.4	5.2	7.6	7.0	5.6	5.8	8.0	7.9
19	0	0	-0.5	0	0.20	5.5	5.6	6.2	6.2	6.9	6.8	7.2	7.2
20	0	0	-0.5	0	0.50	5.7	5.7	7.4	6.9	7.0	6.8	8.5	8.4
21	0	0	-0.5	0	0.80	5.9	5.6	8.4	7.9	6.9	6.7	9.9	10.1
22	0	0	-0.8	0	0.20	8.0	8.4	7.8	8.7	12.7	13.6	12.5	13.0
23	0	0	-0.8	0	0.50	9.1	9.3	8.6	9.3	14.5	15.2	13.5	14.1
24	0	0	-0.8	0	0.80	9.9	9.7	8.2	9.5	16.2	16.3	13.5	14.6

Tables 2.4.3 and 2.4.4 report the empirical rejection frequency of the null hypothesis under the alternative, i.e., the power of the unit root tests. Recall that




Figure 2.1: Finite sample size distortions vis-à-vis a nominal size of 5 per cent, $T = 100$, Constant only

the power is size-adjusted. As expected, when a time trend is included in the test regression power is in general lower than when only a constant is present. Moreover, the tests with GLS detrending tend to have a better power performance than the tests with OLS detrending. This is true whether only a constant is considered or a time trend is also included. As reported in Hansen (1995), the power of the tests increases as more correlated covariates are included in the unit root test regressions.

In the vast majority of cases, the differences in size-adjusted power between the mixed- and the low-frequency unit root tests are positive, meaning that the power of mixed-frequency unit root tests is higher than that of low-frequency tests. This is true for both tests with OLS or GLS demeaning. The power differences between the low- and mixed-frequency with GLS demeaning are presented in Figure 2.2.

While mixed-frequency regressions tend to capture the original correlation between the dependent variable and the covariate, this correlation is hampered, in most cases, by time aggregation. Hence, the actual (and estimated) R^2 for the low-frequency tests may differ significantly from the initial R^2 of the time dis-




Figure 2.2: Differences in size-adjusted power between mixed- and low-frequency tests with GLS demeaning for $c = -5, -10$ and -15 , Constant only, $T = 100$

aggregated DGP, as shown in the top panel of Table 2.4.5.¹⁶ This information loss penalises the power performance of the low-frequency tests. Time aggregation has a milder impact on the correlation between the aggregated data when strong autoregressive dynamics is in place. In this scenario, the performance of the mixed-frequency tests is similar or, in some cases, slightly worse than the performance of the low-frequency tests.

Adding to greater efficiency, the power performance of tests that exploit mixed-frequency data also benefits from less parameter bias (middle panel in Table 2.4.5). As expected, the average estimate of the parameter of interest, δ , is closer to zero for tests with GLS demeaning and with higher R^2 . In addition, the mixed-frequency estimates are also closer to zero than their low-frequency counterparts.

The bottom panel of Table 2.4.5 shows the finite-sample and asymptotic critical values for each test. The asymptotic values were obtained from a simulation exercise with 1,500 observations and 60,000 replications. Separate exercises were

¹⁶Table 2.4.5 shows the results for the DGP with the simplest dynamics, i.e., DGP 1, 2 and 3 in Table 2.4.1. Moreover, it covers the case of including a constant only. Results are qualitative similar if a time trend is also included and are available from the author upon request.

CHAPTER 2. CADF TESTS WITH MIXED-FREQUENCY DATA

Table 2.4.3: Size-adjusted power of unit root tests, Constant only, $T = 100$

	Error structure					Estimated R^2						$c = -5$			$c = -10$			$c = -15$						
	a_1	a_2	b_1	b_2	Initial R^2	CADF	CADF	M	M	M	CADF	CADF	M	M	M	CADF	CADF	M	M	M	CADF	CADF	M	M
1	0	0	0	0	0.20	0.10	0.10	0.23	0.23	35.6	64.4	44.9	70.3	60.2	75.0	67.2	79.9	67.2	77.0	73.4	81.1			
2	0	0	0	0	0.50	0.20	0.20	0.50	0.49	42.7	69.7	71.5	85.0	66.3	79.7	85.8	90.7	73.5	81.9	90.2	91.4			
3	0	0	0	0	0.80	0.30	0.31	0.78	0.76	50.5	75.4	95.3	96.7	73.2	84.4	98.8	98.2	80.7	86.4	99.4	97.8			
4	0.2	0	0	0	0.20	0.13	0.13	0.24	0.23	39.6	70.4	47.5	73.9	65.0	79.9	70.3	82.7	70.8	81.6	75.2	83.8			
5	0.2	0	0	0	0.50	0.27	0.27	0.51	0.50	51.1	78.3	74.7	88.0	73.3	86.3	87.6	92.7	79.3	87.6	91.1	93.2			
6	0.2	0	0	0	0.80	0.41	0.42	0.79	0.78	64.6	84.9	96.1	97.5	82.3	91.7	99.1	98.6	87.8	92.7	99.5	98.5			
7	0.5	0	0	0	0.20	0.18	0.19	0.25	0.25	38.9	81.2	41.6	78.6	74.5	89.9	74.0	88.8	79.4	90.8	79.7	89.8			
8	0.5	0	0	0	0.50	0.41	0.42	0.52	0.52	65.2	90.0	76.6	91.0	86.7	95.5	89.6	95.6	89.8	95.9	92.0	95.9			
9	0.5	0	0	0	0.80	0.64	0.64	0.80	0.79	86.0	95.7	97.0	98.6	95.7	98.4	99.1	99.4	97.5	98.6	99.5	99.3			
10	0.8	0	0	0	0.20	0.22	0.26	0.28	0.28	17.3	58.8	19.8	57.4	40.6	86.1	37.9	82.8	63.9	93.0	55.7	91.1			
11	0.8	0	0	0	0.50	0.45	0.54	0.54	0.54	56.6	83.6	60.0	82.4	80.9	96.1	72.6	92.9	88.2	97.5	76.1	95.4			
12	0.8	0	0	0	0.80	0.73	0.80	0.80	0.80	91.4	97.4	91.3	96.4	98.3	99.7	96.3	98.8	99.0	99.8	96.7	98.6			
13	0.2	0.2	0	0	0.20	0.12	0.13	0.25	0.24	40.0	73.7	45.4	75.6	66.3	81.6	68.8	83.0	71.4	82.8	73.6	84.4			
14	0.2	0.2	0	0	0.50	0.31	0.31	0.52	0.52	56.7	82.9	75.2	88.7	77.6	89.4	86.7	93.4	81.6	90.5	90.2	94.0			
15	0.2	0.2	0	0	0.80	0.53	0.53	0.81	0.80	76.7	91.7	96.4	97.9	90.1	96.0	98.9	99.2	93.7	96.5	99.5	99.1			
16	0	0	-0.2	0	0.20	0.08	0.08	0.23	0.23	32.5	59.7	42.0	66.3	56.2	71.0	64.2	76.5	64.6	73.7	71.4	77.9			
17	0	0	-0.2	0	0.50	0.16	0.15	0.49	0.48	36.3	63.7	66.8	81.6	60.9	75.0	83.1	87.7	69.5	77.5	88.1	88.1			
18	0	0	-0.2	0	0.80	0.23	0.22	0.77	0.75	41.0	67.8	93.6	95.2	65.7	78.6	98.3	97.0	75.0	81.1	99.2	96.7			
19	0	0	-0.5	0	0.20	0.07	0.07	0.22	0.22	30.8	50.7	38.3	57.1	54.2	62.7	61.3	66.9	63.4	65.9	69.5	69.4			
20	0	0	-0.5	0	0.50	0.12	0.11	0.49	0.47	32.3	52.9	61.0	72.3	57.3	65.1	79.9	79.6	67.7	68.6	85.5	80.1			
21	0	0	-0.5	0	0.80	0.17	0.15	0.78	0.75	34.3	54.7	88.9	89.4	60.8	67.3	97.0	92.5	72.3	71.0	98.5	91.8			
22	0	0	-0.8	0	0.20	0.10	0.09	0.21	0.19	32.2	32.2	37.1	34.1	57.2	49.7	61.7	51.4	67.2	58.3	71.5	59.6			
23	0	0	-0.8	0	0.50	0.19	0.17	0.47	0.44	35.2	33.5	51.5	43.0	61.6	51.8	75.5	59.0	72.8	60.8	82.5	65.9			
24	0	0	-0.8	0	0.80	0.28	0.24	0.77	0.74	38.0	34.6	78.3	60.4	67.5	53.7	92.2	71.1	80.2	63.9	95.8	75.0			

Table 2.4.4: Size-adjusted power of unit root tests, Time trend included, $T = 100$

	Error structure			Estimated R^2																	
				$c = -5$				$c = -10$				$c = -15$									
				a_1	a_2	b_1	b_2	Initial R^2	CADF	M	CADF	GLS	M	CADF	GLS	M	CADF	GLS	M	CADF	GLS
1	0	0	0	0	0.20	0.09	0.10	0.20	0.23	29.2	45.4	36.9	52.7	59.1	69.4	64.9	74.3	68.1	74.9	72.9	79.0
2	0	0	0	0	0.50	0.19	0.20	0.42	0.49	35.2	51.4	63.3	74.2	65.2	74.3	82.9	88.3	73.4	79.7	87.6	90.9
3	0	0	0	0	0.80	0.28	0.31	0.62	0.76	42.2	58.7	92.0	94.1	71.2	79.6	97.7	98.1	79.3	84.9	98.9	98.8
4	0.2	0	0	0	0.20	0.12	0.13	0.20	0.23	32.6	51.6	39.8	57.3	65.6	75.3	69.6	78.7	72.8	79.0	76.1	82.1
5	0.2	0	0	0	0.50	0.25	0.27	0.43	0.50	43.9	61.7	67.9	78.6	73.5	81.8	85.4	90.8	79.6	85.7	89.2	92.8
6	0.2	0	0	0	0.80	0.38	0.42	0.63	0.78	56.8	72.6	93.5	95.6	81.7	88.4	98.1	98.8	86.6	91.8	99.0	99.3
7	0.5	0	0	0	0.20	0.16	0.19	0.21	0.25	34.6	62.4	34.6	57.3	75.0	87.5	73.1	85.7	83.7	89.4	82.9	88.5
8	0.5	0	0	0	0.50	0.37	0.42	0.42	0.52	59.9	78.5	71.4	81.9	87.3	94.0	89.5	93.9	90.7	95.4	92.4	95.2
9	0.5	0	0	0	0.80	0.59	0.64	0.62	0.79	82.2	91.0	96.0	97.1	95.1	98.2	98.7	99.2	97.1	99.0	99.3	99.5
10	0.8	0	0	0	0.20	0.19	0.27	0.23	0.28	15.8	34.2	16.3	30.7	35.6	70.2	30.5	57.1	57.8	88.3	43.9	77.4
11	0.8	0	0	0	0.50	0.44	0.55	0.39	0.53	54.9	68.4	57.3	67.9	78.6	92.0	67.6	81.2	85.9	96.1	64.6	84.4
12	0.8	0	0	0	0.80	0.72	0.81	0.50	0.79	89.8	93.3	92.2	93.0	98.1	99.3	96.3	97.2	98.8	99.6	94.9	96.1
13	0.2	0.2	0	0	0.20	0.11	0.13	0.21	0.24	32.2	54.7	36.5	56.2	67.5	77.4	69.7	78.2	73.9	80.6	75.3	81.3
14	0.2	0.2	0	0	0.50	0.28	0.31	0.45	0.51	48.3	68.5	67.6	79.2	77.6	85.8	85.2	90.7	81.8	88.3	88.2	92.7
15	0.2	0.2	0	0	0.80	0.49	0.53	0.68	0.79	70.2	83.7	93.9	96.0	89.3	94.9	97.8	98.8	92.2	96.7	98.8	99.3
16	0	0	-0.2	0	0.20	0.08	0.08	0.20	0.22	26.0	40.9	33.6	46.2	53.7	64.5	60.2	69.0	65.0	71.5	69.9	75.3
17	0	0	-0.2	0	0.50	0.15	0.15	0.41	0.48	29.5	44.1	57.2	67.1	58.2	68.2	79.6	84.2	69.2	75.3	85.8	87.8
18	0	0	-0.2	0	0.80	0.21	0.22	0.61	0.75	32.7	47.9	88.7	90.9	62.8	72.2	97.0	97.2	73.7	79.3	98.5	98.0
19	0	0	-0.5	0	0.20	0.07	0.07	0.19	0.22	23.1	34.0	29.8	39.4	49.7	57.6	55.9	62.3	62.3	66.1	66.6	69.5
20	0	0	-0.5	0	0.50	0.12	0.11	0.40	0.47	24.4	35.8	50.7	56.8	52.6	59.9	75.4	77.0	65.6	69.3	82.8	81.5
21	0	0	-0.5	0	0.80	0.16	0.16	0.63	0.75	25.1	37.1	82.4	83.9	55.0	62.6	94.5	93.7	68.8	72.1	97.3	95.2
22	0	0	-0.8	0	0.20	0.09	0.09	0.18	0.19	24.1	24.2	26.9	26.2	51.0	48.5	54.5	50.7	63.7	60.4	66.1	62.0
23	0	0	-0.8	0	0.50	0.18	0.17	0.39	0.44	26.1	25.9	35.6	34.1	54.8	51.2	63.9	58.7	68.2	64.0	74.4	68.5
24	0	0	-0.8	0	0.80	0.27	0.25	0.63	0.74	27.8	27.2	62.1	55.6	59.2	54.3	84.0	75.2	74.2	68.4	89.2	81.0

CHAPTER 2. CADF TESTS WITH MIXED-FREQUENCY DATA

Table 2.4.5: Estimated R^2 and δ parameters and critical values, Constant only

	Initial R^2					
	OLS detrending			GLS detrending		
	0.2	0.5	0.8	0.2	0.5	0.8
Estimated R^2						
Low-frequency tests						
$T = 100$	0.11	0.23	0.34	0.12	0.25	0.38
$T = 500$	0.09	0.20	0.32	0.10	0.23	0.37
Mixed-frequency tests						
$T = 100$	0.23	0.50	0.78	0.23	0.49	0.77
$T = 500$	0.20	0.48	0.77	0.20	0.48	0.77
Estimated δ						
Low-frequency tests						
$T = 100$	-0.046	-0.041	-0.036	-0.019	-0.016	-0.014
$T = 500$	-0.010	-0.009	-0.008	-0.003	-0.003	-0.002
Mixed-frequency tests						
$T = 100$	-0.041	-0.027	-0.013	-0.016	-0.008	0.000
$T = 500$	-0.009	-0.006	-0.003	-0.003	-0.001	0.000
Finite-sample critical values						
Low-frequency tests						
$T = 100$	-2.61	-2.57	-2.52	-1.81	-1.80	-1.77
$T = 500$	-2.67	-2.64	-2.60	-1.80	-1.77	-1.73
<i>memo</i> : Asymptotical critical values						
	-2.81	-2.75	-2.70	-1.91	-1.85	-1.78
Mixed-frequency tests						
$T = 100$	-2.62	-2.58	-2.48	-1.84	-1.75	-1.55
$T = 500$	-2.69	-2.61	-2.49	-1.80	-1.74	-1.70
<i>memo</i> : Asymptotical critical values						
	-2.75	-2.60	-2.32	-1.87	-1.72	-1.52

Notes: All estimates were obtained from the DGP with the simplest dynamics, i.e., DGP 1, 2 and 3 in Table 2.4.1. The estimated R^2 and δ correspond to the average values over the replications. The asymptotical critical values are interpolated for the estimated R^2 .

performed for the low- and mixed-frequency approaches, delivering similar results. The results are also similar to the values collected from the original papers. The asymptotic critical values presented in the table are interpolated for the values of R^2 . Finite-sample critical values converge to their asymptotic values, though at a lower pace in the case of GLS demeaning. For comparable values of R^2 the mixed-frequency critical values are closer to their asymptotic value than the low frequency ones. This result is underpinned by the greater efficiency and less biased estimates of mixed-frequency tests.

On average, the power gains from taking on board mixed-frequency data are quite substantial, reaching 9.6, 6.4 and 5.2 per cent for $c = -5, -10$ and -15 , respectively. Recall that the results above mentioned correspond to $\alpha = 0.95, 0.90$ and 0.85 on the disaggregated process, meaning that time aggregation with $m = 3$ leads to $\alpha = 0.86, 0.73$ and 0.61 , respectively. The power gains increase as the alternative hypothesis are more demanding, i.e., are closer to the unit root. This results from losses in power performance, which are smaller for the mixed-frequency tests than for the low-frequency ones.

2.4.2 The case of a larger sample size

Now consider a sample with 500 observations. Recall that the values of c , corresponding to $\alpha = 0.99, 0.98$ and 0.97 , are for the disaggregated processes. Time aggregation with $m = 3$ leads to $\alpha = 0.97, 0.94$ and 0.91 , respectively. Not only do we consider a larger sample but the alternative hypothesis is more demanding, being closer to the unit root case.

In general, there are less size distortions using this larger sample.¹⁷ Again, as in the case of the sample with 100 observations, when downward distortions exist, they tend to be less marked for mixed-frequency tests (Figure 2.3). Though upward biased, the performance of the mixed-frequency tests is more favourable with strong moving average dynamics. This is not the case when the moving average coefficient is smaller (in absolute terms).

¹⁷The results presented in this section refer to the case of only a constant included in the test regressions. The results are qualitatively similar when also a time trend is considered and are available from the author upon request.

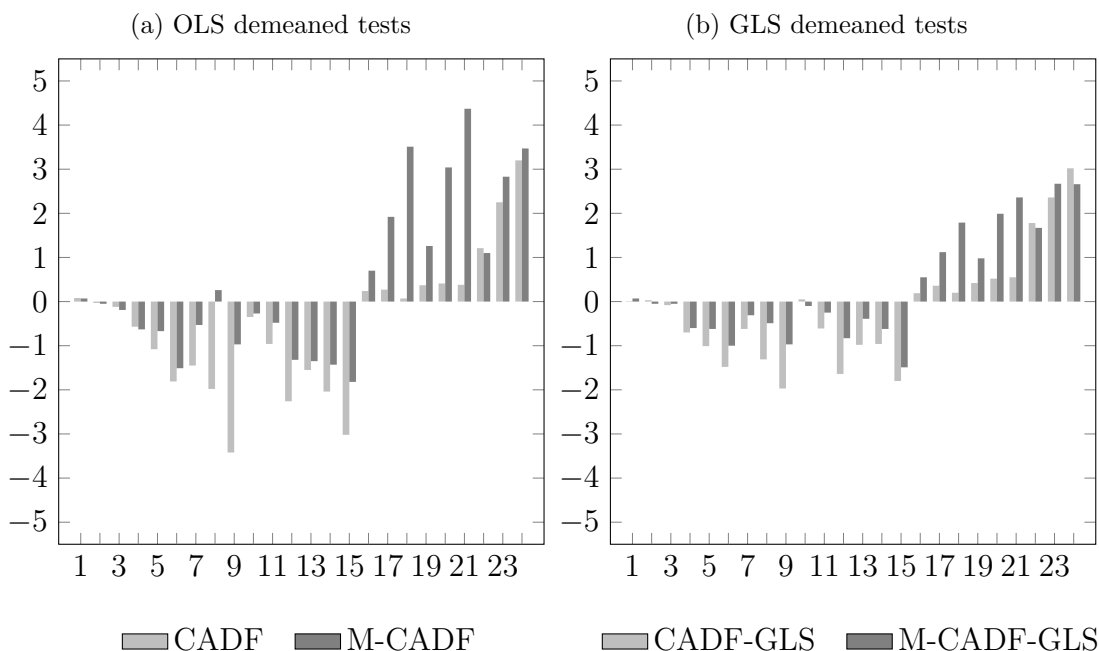


Figure 2.3: Finite sample size distortions vis-à-vis a nominal size of 5 per cent, $T = 500$

Regardless of the DGP, power is always higher in the sample with 500 observations. Moreover, size-adjusted power of mixed-frequency tests is at least as high as the one of low-frequency tests. This is true for both tests with OLS or GLS demeaning. The power differences between the low- and mixed-frequency tests, for the GLS case, are presented in Figure 2.4. In particular, size-adjusted power of mixed-frequency tests is substantially higher for processes with moving average dynamics.

Mixed-frequency tests also deal better with near-integration. On average, the power gains from exploiting mixed-frequency data increase as c increases, from 1.7 to 2.1 and 5.4 per cent for $c = -15, -10$ and -5 , respectively. Hence, as the series becomes near-integrated, the advantage of exploiting mixed-frequency information increases progressively.




Figure 2.4: Differences in size-adjusted power between mixed- and low-frequency tests with GLS demeaning for $c = -5, -10$ and -15 , Constant only, $T = 500$

2.4.3 Different lags

In this section we present a sensitivity analysis to the choice of the truncation lag. Figure 2.5 shows the differences between the size-adjusted power of mixed- and low-frequency tests with an ad-hoc number of lags included in the test regression (3, 4, 6 and 8 lags, respectively).¹⁸ The same number of lags is used in each replication and for each variable (the dependent variable and the covariate). These results illustrate the case of $c = -5$, only a constant added to the test regression and of GLS demeaning.¹⁹ A positive bar means that the size-adjusted power of mixed-frequency tests is higher than the one of low-frequency tests.

In spite of high costs in terms of size (as expected), this exercise shows that regardless of the particular choice of lags, the mixed-frequency tests tend to outperform the low-frequency ones. This is due to the fact that exploiting mixed-

¹⁸The number of lags refers to the low frequency. For the mixed-frequency tests the same time span of lagged information of the covariate is covered, which corresponds to a different number of lags in the high time frequency.

¹⁹Results for other c , including a time trend and OLS demeaning are not qualitatively different and are available from the author upon request.




Figure 2.5: Differences in size-adjusted power between mixed- and low-frequency tests with GLS demeaning for $c = -5$ by number of lags included in the test regression, Constant only

frequency data enables us to capture the (stronger) underlying correlation between the dependent variable and the covariate.

In addition, note that the gains tend to increase with the number of lags included. The performance of the low-frequency tests is more severely affected by the inclusion of unnecessary lags. In contrast, the mixed-frequency tests are better able to deal with this kind of misspecification issues, because the weights of the high-frequency lags are data-driven. Hence, in case of great uncertainty about the choice of the truncation lag (as is typically the case in empirical applications), using the mixed-frequency framework may contribute to reduce the impact of potential misspecification in the power of the unit root tests.

2.4.4 Different time frequencies

To assess the impact of different combinations of time frequencies, alternative figures for m are considered. In addition to $m = 3$, now I will also consider $m = 12, 24, 36, 60$ and 120 .

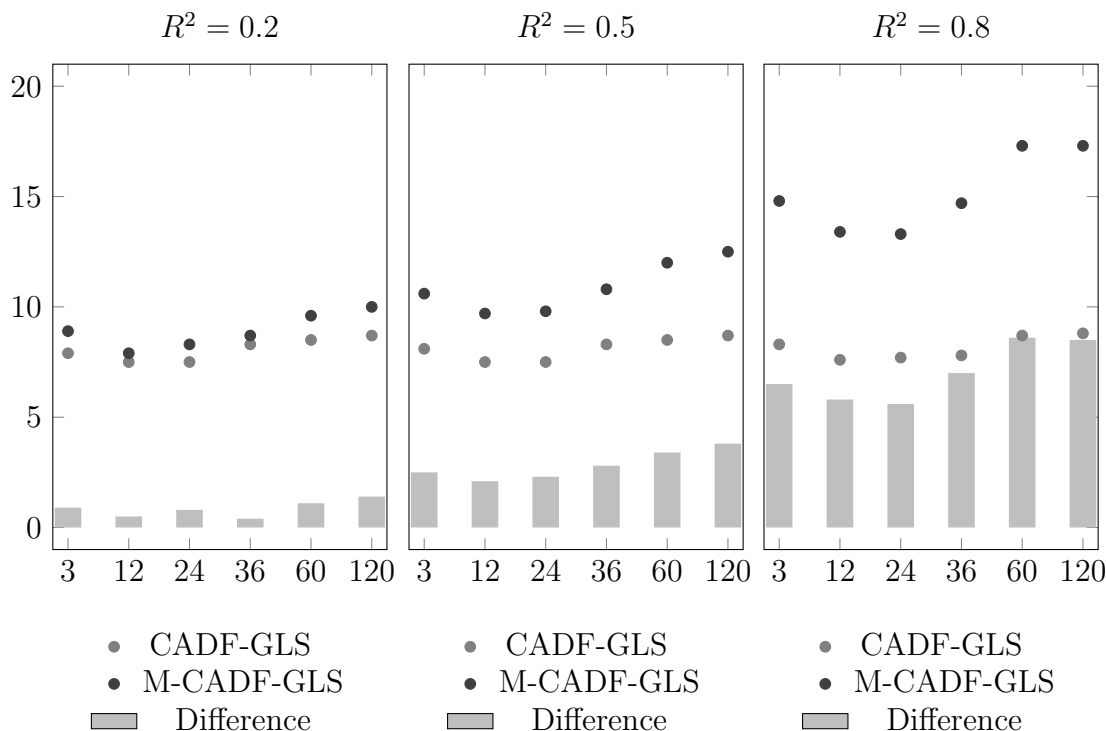


Figure 2.6: Size-adjusted power of mixed- and low-frequency tests for different m , with GLS demeaning, aggregate $\alpha = 0.99$, Constant only, $T = 500$

Hence, to have a meaningful comparison across different m , we simulate samples with $T = 100$ and 500 observations such that the aggregate figure for the first-order autoregressive parameter equals 0.95 and 0.99 , respectively, regardless of m . Recall that for doing this we need to adjust the values of c accordingly.

Figure 2.6 shows the results in terms of power performance of low- and mixed-frequency unit root tests, with GLS demeaning, using a sample of 500 observations.²⁰

Once again, the mixed-frequency test shows better power performance for all m considered in a situation of near-integration. The mixed-frequency test is at least as good as the low-frequency test and, in many cases, substantially better, especially when R^2 is higher. The gains from using the mixed-frequency test are

²⁰Results for OLS demeaning or with the smaller sample ($T = 100$) are qualitatively similar and were omitted for the sake of brevity. All results are available from the author upon request.

fairly stable across different figures for m , for each R^2 . However, as noted before, the gains tend to increase significantly with R^2 .

2.5 An application to the US unemployment rate

There is a rich discussion in the literature about the order of integration of the unemployment rate. Initial contributions by Phelps (1967) and Friedman (1968) described movements in the unemployment rate as fluctuations around a natural rate, which would be generally defined as the equilibrium rate. Given that temporary shocks would have only temporary effects, these traditional theories imply that the unemployment rate is level or, perhaps, trend stationary, evolving around the natural level.²¹

In contrast, Blanchard and Summers (1986, 1988) resort to the concept of hysteresis, meaning that unemployment rates depend sensitively on the shocks an economy experienced in the past, and eventually the unemployment rate should exhibit a unit root. Unifying both strands, the structuralist theories of unemployment assumes that most shocks cause temporary movements of the unemployment rate around the natural rate, but some shocks can cause permanent changes in the natural rate. Hence, the unemployment rate would be stationary around a natural rate, which itself could be subject to structural breaks (for a brief discussion, see Phelps, 1995).

The purpose of this exercise is to illustrate the used of mixed-frequency unit root tests, providing additional evidence on the persistence of US unemployment rate. We apply the above described covariate-augmented unit root tests — both low- and mixed-frequency tests — to assess whether US unemployment rate has a unit root.²² We exploit insights provided by the correlation between the variable

²¹There is no consensus in the literature about including or not a trend when modelling unemployment, existing an ongoing discussion based on sample-driven and theoretical arguments.

²²The difference between the unemployment rate and the natural rate of unemployment is also often analysed. In this exercise we focused on the level of the unemployment rate. Estimating a natural rate of unemployment is beyond the scope of this article. Simply using an estimate of the natural rate of unemployment collected elsewhere (e.g., the Congressional Budget Office estimates a quarterly natural rate of unemployment, which is made available by the St. Louis Federal Reserve Bank) would lead to bias on hypothesis testing, as shown by Murphy and Topel (1985).




Figure 2.7: Data on US unemployment

of interest and the continued jobless claims.

The overall unemployment rate (in logs) is a monthly series taken from the US Bureau of Labor Statistics and is seasonally adjusted. The covariate is the continued jobless claims, also expressed in logs, which is a weekly series. The continued jobless claims are released by the US Department of Labor and are also seasonally adjusted. The stationarity of the covariate was confirmed by univariate unit root tests, namely ADF and ADF-GLS tests.²³ The data cover the period from January 1980 to June 2014. For the low-frequency version of the covariate we use the beginning of the period value. Figure 2.7 shows the series used, in quarterly frequency (the common frequency) and in logs.

In order to choose the truncation lag of the variables, we used the multivariate version of MAIC in Perron and Qu (2007), applied to the output of an unrestricted version of the CADF test regression (for more details, see section 2.4). Tables 2.5.1 and 2.5.2 present the results for the unit root tests to the US unemployment rate. These tables provide the estimates of δ , the t -statistics and the estimated R^2 for

²³The tests were performed for weekly, monthly and quarterly frequencies and, as expected, the result was always the same.

the covariate-augmented regressions.

We consider two different samples, one from 1980 Q1 to 2014 Q2 in Table 2.5.1 and another excluding the great recession period, from 1980 Q1 to 2006 Q4, in Table 2.5.2. By looking at Figure 2.7 one can intuitively see that it seems to be relevant to include a time trend in the test regression for the level of the unemployment rate, namely in the shorter sample. Therefore, for each sample, the top panel in Tables 2.5.1 and 2.5.2 shows results for only including a constant and the bottom panel shows the results for also including a time trend.

In the shorter sample, all tests for the level of the unemployment rate with only a constant included in the test regression agree in not rejecting the null hypothesis, suggesting that the series is $I(1)$. The conclusions from the univariate tests remain unchanged when a time trend is included. However, all covariate-augmented tests agree in rejecting the null hypothesis, suggesting that the unemployment rate is trend stationary in that sample.

Table 2.5.1: Unit root tests for US monthly unemployment rate, using jobless claims as covariate

Sample: 1980Q1 to 2014Q2						
Constant only						
	ADF	ADF-GLS	CADF	CADF-GLS	M-CADF	M-CADF-GLS
δ	-0.009	-0.009	-0.024	-0.024	-0.022	-0.022
t -statistic	-1.882	-1.879	-3.439 **	-3.428*	-3.309 **	-3.293*
R^2			0.42	0.43	0.96	0.95
Time trend included						
	ADF	ADF-GLS	CADF	CADF-GLS	M-CADF	M-CADF-GLS
δ	-0.009	-0.009	-0.041	-0.027	-0.041	-0.025
t -statistic	-1.876	-1.896	-4.626 **	-3.657*	-4.861 **	-3.574*
R^2			0.39	0.43	0.96	0.96

Note: For the low-frequency covariate-augmented unit root tests the frequency of the covariate equals the frequency of the dependent variable. For the mixed-frequency tests, the covariate has a weekly frequency. * significant at a 5 per cent asymptotic level. ** significant at a 1 per cent asymptotic level. For the covariate-augmented GLS family of tests, Pesavento (2006) only presents 5 per cent asymptotic significance levels.

CHAPTER 2. CADF TESTS WITH MIXED-FREQUENCY DATA

Table 2.5.2: Unit root tests for US monthly unemployment rate, using jobless claims as covariate

Sample: 1980Q1 to 2006Q4						
Constant only						
	ADF	ADF-GLS	CADF	CADF-GLS	M-CADF	M-CADF-GLS
$\hat{\delta}$	-0.007	-0.007	-0.015	-0.012	-0.014	-0.010
t -statistic	-1.217	-1.089	-1.989	-1.628	-1.888	-1.406
R^2			0.44	0.44	0.95	0.92
Time trend included						
	ADF	ADF-GLS	CADF	CADF-GLS	M-CADF	M-CADF-GLS
$\hat{\delta}$	-0.018	-0.009	-0.088	-0.035	-0.078	-0.032
t -statistic	-1.913	-1.209	-5.147 **	-3.148*	-4.787 **	-3.043*
R^2			0.29	0.36	0.87	0.88

Note: For the low-frequency covariate-augmented unit root tests the frequency of the covariate equals the frequency of the dependent variable. For the mixed-frequency tests, the covariate has a weekly frequency. * significant at a 5 per cent asymptotic level. ** significant at a 1 per cent asymptotic level. For the covariate-augmented GLS family of tests, Pesavento (2006) only presents 5 per cent asymptotic significance levels.

When using the longer sample, univariate unit root tests continue to suggest that the level of US unemployment is not stationary. In contrast, all covariate-augmented tests reject the null hypothesis, whether or not a time trend is included. Hence, the level of unemployment rate seems to be stationary, though highly persistent, as shown by the small values of $\hat{\delta}$. Notice that in all cases, the estimates of R^2 are higher when the mixed-frequency approach is used, suggesting that in this case combining information with different time frequencies allows us to take better advantage of the covariate-augmented framework of unit root tests.

2.6 Conclusion

Unit root tests typically have low power, especially in near-integrated cases, which results in the over-acceptance of the unit root null. This paper tries to tackle this issue by merging two strands of the literature. In particular, we try to improve the

power performance of CADF unit root tests by exploiting mixed-frequency data. The results of a simulation exercise show that there is room for improvement.

Since Hansen (1995), covariate-augmented unit root tests have been proposed has a more powerful version of traditional univariate unit root tests, such as the ADF tests. The main idea is that using a stationary covariate, which is well correlated with the variable of interest, in the test regression contributes to increase the precision of the estimates of the test statistic and, hence, to increase the power of the test.

In this article we assume as main premise that temporal aggregation of the variable of interest is unavoidable. It is well known that time aggregation and the sampling frequency do not affect the long-run properties of time series, namely the presence of unit roots, but may have severe consequences for the correlation between the dependent variable and the covariate.

To exploit the advantages of combining data with different time frequencies — a dependent variable in a lower frequency than the covariate — we use the MIDAS technique. This technique uses data-driven aggregation weights. Monte Carlo experiments show that: (i) mixed-frequency covariate-augmented unit root tests have a better power performance than traditional low-frequency tests; and that (ii) mixed-frequency tests are particularly advantageous when we are in the presence of near-integrated variables. The results are robust to the size of the sample, to the lag specification of the test regression and to different combinations of time frequencies.

Applying the unit root tests — both low- and mixed-frequency — to the US unemployment rate, we found evidence that the unemployment rate is stationary, though highly persistent.

Acknowledgments

I would like to thank Paulo Júlio, José Francisco Maria, Carlos Robalo Marques, João Nicolau and Paulo Rodrigues. The usual disclaimers apply.

Bibliography

- ANDREOU, E., E. GHYSELS AND A. KOURTELLOS, *The Oxford Handbook of Economic Forecasting*, chapter 8, Forecasting with mixed-frequency data (Oxford University Press, 2011), 225–267.
- , “Should macroeconomic forecasters use daily financial data and how?,” *Journal of Business and Economic Statistics* 2 (2013), 240–251.
- ANDREWS, D. W. K., “Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation,” *Econometrica* 59 (May 1991), 817–58.
- BLANCHARD, O. J. AND L. H. SUMMERS, “Hysteresis And The European Unemployment Problem,” in *NBER Macroeconomics Annual 1986, Volume 1* NBER Chapters (National Bureau of Economic Research, Inc, 1986), 15–90.
- , “Beyond the Natural Rate Hypothesis,” *American Economic Review* 78 (May 1988), 182–87.
- BRILLINGER, D., *Time Series: Data Analysis and Theory*, Classics in Applied Mathematics (Society for Industrial and Applied Mathematics, 2001).
- CAPORALE, G. M. AND N. PITTIS, “Unit Root Testing Using Covariates: Some Theory and Evidence,” *Oxford Bulletin of Economics and Statistics* 61 (November 1999), 583–95.
- CHANG, Y. AND J. Y. PARK, “On the asymptotics of ADF tests for unit roots,” *Econometric Reviews* 21 (2002), 431–447.
- CHRISTOPOULOS, D. K. AND M. A. LEÓN-LEDESMA, “Time-series output convergence tests and stationary covariates,” *Economics Letters* 101 (December 2008), 297–299.
- ELLIOTT, G., “Efficient Tests for a Unit Root When the Initial Observation is Drawn From Its Unconditional Distribution,” *International Economic Review* 40 (1999), 767–784.

- ELLIOTT, G. AND M. JANSSON, “Testing for unit roots with stationary covariates,” *Journal of Econometrics* 115 (2003), 75–89.
- ELLIOTT, G., T. J. ROTHENBERG AND J. H. STOCK, “Efficient Tests for an Autoregressive Unit Root,” *Econometrica* 64 (July 1996), 813–36.
- FORONI, C. AND M. MARCELLINO, “A Comparison of Mixed Frequency Approaches for Modelling Euro Area Macroeconomic Variables,” Economics Working Papers ECO2012/07, European University Institute, 2012.
- FORONI, C., M. MARCELLINO AND C. SCHUMACHER, “U-MIDAS: MIDAS regressions with unrestricted lag polynomials,” Discussion Paper Series 1: Economic Studies 2011,35, Deutsche Bundesbank, Research Centre, 2011.
- FOSSATI, S., “Covariate unit root tests with good size and power,” *Computational Statistics & Data Analysis* 56 (2012), 3070–3079.
- FRIEDMAN, M., “The Role of Monetary Policy,” *The American Economic Review* 58 (1968), pp. 1–17.
- GALVAO JR., A. F., “Unit root quantile autoregression testing using covariates,” *Journal of Econometrics* 152 (October 2009), 165–178.
- GALVÃO, A. B., “Changes in predictive ability with mixed frequency data,” *International Journal of Forecasting* 29 (2013), 395–410.
- GHYSELS, E. AND J. I. MILLER, “Testing for Cointegration with Temporally Aggregated and Mixed-frequency Time Series,” Working Papers 1307, Department of Economics, University of Missouri, June 2013.
- GHYSELS, E., P. SANTA-CLARA AND R. VALKANOV, “The MIDAS touch: Mixed data sampling regression models,” CIRANO Working Papers 2004s-20, CIRANO, 2004.
- , “Predicting volatility: getting the most out of return data sampled at different frequencies,” *Journal of Econometrics* 131 (2006), 59–95.
- GHYSELS, E., A. SINKO AND R. VALKANOV, “MIDAS regressions: Further results and new directions,” *Econometric Reviews* 26 (2007), 53–90.
- GRANGER, C. AND P. NEWBOLD, “SPURIOUS REGRESSIONS IN ECONOMETRICS,” *Journal of Econometrics* 2 (1974), 111–120.
- GRANGER, C. AND P. SIKLOS, “Systematic sampling, temporal aggregation, seasonal adjustment, and cointegration theory and evidence,” *Journal of Econometrics* 66 (1995), 357–369.

BIBLIOGRAPHY

- HALDRUP, N. AND M. JANSSON, “Improving Size and Power in Unit Root Testing,” in T. C. Mills and K. Patterson, eds., *Palgrave Handbook of Econometrics* volume 2: Applied Econometrics, chapter 7 (Palgrave Macmillan, 2006), 252–277.
- HANSEN, B. E., “Rethinking the Univariate Approach to Unit Root Testing: Using Covariates to Increase Power,” *Econometric Theory* 11 (October 1995), 1148–1171.
- HAUG, A. A., “Temporal Aggregation and the Power of Cointegration Tests: A Monte Carlo Study,” *Oxford Bulletin of Economics and Statistics* 64 (September 2002), 399–412.
- JANSSON, M., “Consistent Covariance Matrix Estimation For Linear Processes,” *Econometric Theory* 18 (December 2002), 1449–1459.
- , “Stationarity Testing With Covariates,” *Econometric Theory* 20 (February 2004), 56–94.
- JUHL, T. AND Z. XIAO, “Power functions and envelopes for unit root tests,” *Econometric Theory* 19 (4 2003), 240–253.
- LUPI, C., “Unit Root CADF Testing with R,” *Journal of Statistical Software* 32 (October 2009).
- MARCELLINO, M., “Consequences of temporal aggregation in empirical analysis,” *Journal of Business and Economic Statistics* 17 (January 1999), 129–136.
- MARCELLINO, M. AND C. SCHUMACHER, “Factor MIDAS for Nowcasting and Forecasting with Ragged-Edge Data: A Model Comparison for German GDP,” *Oxford Bulletin of Economics and Statistics* 72 (08 2010), 518–550.
- MÜLLER, U. K. AND G. ELLIOTT, “Tests for Unit Roots and the Initial Condition,” *Econometrica* 71 (07 2003), 1269–1286.
- MURPHY, K. AND R. TOPEL, “Estimation and Inference in Two-Step Econometric Models,” *Journal of Business and Economic Statistics* 3 (October 1985), 370–379.
- NG, S. AND P. PERRON, “Unit Root Tests in ARMA Models with Data-Dependent Methods for the Selection of the Truncation Lag,” *Journal of the American Statistical Association* 90 (1995), pp. 268–281.
- , “Lag Length Selection and the Construction of Unit Root Tests with Good Size and Power,” *Econometrica* 69 (2001), 1519–1554.

- PERRON, P. AND S. NG, “Useful Modifications to Some Unit Root Tests with Dependent Errors and Their Local Asymptotic Properties,” *Review of Economic Studies* 63 (July 1996), 435–63.
- PERRON, P. AND Z. QU, “A simple modification to improve the finite sample properties of Ng and Perron’s unit root tests,” *Economics Letters* 94 (January 2007), 12–19.
- PESAVENTO, E., “Near-Optimal Unit Root Tests with Stationary Covariates with Better Finite Sample Size,” Economics Working Papers ECO2006/18, European University Institute, 2006.
- PHELPS, E. S., “Phillips Curves, Expectations of Inflation and Optimal Unemployment over Time,” *Economica* 34 (1967), pp. 254–281.
- , “The Structuralist Theory of Employment,” *The American Economic Review* 85 (1995), pp. 226–231.
- PHILLIPS, P. C. B., “Understanding spurious regressions in econometrics,” *Journal of Econometrics* 33 (December 1986), 311–340.
- , “Towards a unified asymptotic theory for autoregression,” *Biometrika* 74 (1987), 535–547.
- PIERSE, R. G. AND A. J. SNELL, “Temporal aggregation and the power of tests for a unit root,” *Journal of Econometrics* 65 (February 1995), 333–345.
- RODRIGUES, P. M., “Properties of recursive trend-adjusted unit root tests,” *Economics Letters* 91 (June 2006), 413–419.
- SAIKKONEN, P., “Asymptotically Efficient Estimation of Cointegration Regressions,” *Econometric Theory* 7 (March 1991), 1–21.
- SCHWERT, G. W., “Tests for Unit Roots: A Monte Carlo Investigation,” *Journal of Business and Economic Statistics* 7 (1989), 147–159.
- SILVESTRINI, A. AND D. VEREDAS, “Temporal Aggregation Of Univariate And Multivariate Time Series Models: A Survey,” *Journal of Economic Surveys* 22 (07 2008), 458–497.
- STOCK, J. H., *Unit roots, structural breaks and trends*, volume 4 of *Handbook of Econometrics*, chapter 46 (Elsevier, 1986), 2739–2841.
- , *A Class of Tests for Integration and Cointegration*, chapter 6, Cointegration, Causality, and Forecasting: A Festschrift for Clive W. J. Granger (Oxford University Press, 1999), 135–167.

BIBLIOGRAPHY

- SUL, D., P. C. B. PHILLIPS AND C.-Y. CHOI, “Prewhitening Bias in HAC Estimation,” *Oxford Bulletin of Economics and Statistics* 67 (08 2005), 517–546.
- TAYLOR, A. M. R., “Regression-Based Unit Root Tests with Recursive Mean Adjustment for Seasonal and Nonseasonal Time Series,” *Journal of Business and Economic Statistics* 20 (2002), 269–281.
- WESTERLUND, J., “On the Importance of the First Observation in GLS Detrending in Unit Root Testing,” *Oxford Bulletin of Economics and Statistics* 1 (2015), 152–161.
- WORKING, H., “Note on the Correlation of First Differences of Averages in a Random Chain,” *Econometrica* 28 (1960), pp. 916–918.

Appendices

Table B.1: Asymptotic critical values for CADF t -statistics

ρ^2	Standard			Demeaned			Detrended		
	1%	5%	10%	1%	5%	10%	1%	5%	10%
1	-2.57	-1.94	-1.62	-3.43	-2.86	-2.57	-3.96	-3.41	-3.13
0.9	-2.57	-1.94	-1.61	-3.39	-2.81	-2.50	-3.88	-3.33	-3.04
0.8	-2.57	-1.94	-1.6	-3.36	-2.75	-2.46	-3.83	-3.27	-2.97
0.7	-2.55	-1.93	-1.59	-3.30	-2.72	-2.41	-3.76	-3.18	-2.87
0.6	-2.55	-1.90	-1.56	-3.24	-2.64	-2.32	-3.68	-3.10	-2.78
0.5	-2.55	-1.89	-1.54	-3.19	-2.58	-2.25	-3.60	-2.99	-2.67
0.4	-2.55	-1.89	-1.53	-3.14	-2.51	-2.17	-3.49	-2.87	-2.53
0.3	-2.52	-1.85	-1.51	-3.06	-2.40	-2.06	-3.37	-2.73	-2.38
0.2	-2.49	-1.82	-1.46	-2.91	-2.28	-1.92	-3.19	-2.55	-2.20
0.1	-2.46	-1.78	-1.42	-2.78	-2.12	-1.75	-2.97	-2.31	-1.95

Note: Following Hansen (1995). The critical values were calculated from 60.000 draws generated from samples of size 1.000 with *iid* Gaussian innovations.

Table B.2: Asymptotic critical values for the CADF-GLS test

R^2	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Cases 1, 2	-1.948	-1.939	-1.929	-1.918	-1.905	-1.881	-1.864	-1.839	-1.818	-1.773
Case 3	-1.948	-1.909	-1.866	-1.812	-1.760	-1.707	-1.647	-1.579	-1.497	-1.405
Case 4	-2.836	-2.786	-2.738	-2.688	-2.628	-2.568	-2.498	-2.418	-2.343	-2.315
Case 5	-2.835	-2.780	-2.730	-2.664	-2.586	-2.497	-2.401	-2.286	-2.152	-2.017

Note: Following Pesavento (2006). The critical values were computed using 60 000 replications of samples with 1000 observations, drawn using i.i.d. Gaussian innovations. The critical values reported are for tests of size 5%, with \bar{c} equal to -7 for cases 1, 2 and 3 and to -13.5 for cases 4 and 5.

Chapter 3

Unit root tests using mixed-frequency VAR models

ABSTRACT

This paper assesses whether the power performance of covariate-augmented feasible point optimal tests can be improved by exploiting mixed-frequency data. To estimate the test regressions we use the mixed data sampling (MIDAS) approach. In particular, an unconstrained, parsimonious, stacked skip-sampled reduced-form VAR-MIDAS model is estimated using standard econometric techniques. A simulation study reveals that the mixed-frequency test has a better power performance than its low-frequency counterpart. An empirical application to the US unemployment rate illustrates potential uses of this new test.

JEL Classification: C12, C15, C22.

Keywords: Unit root, Hypothesis testing, Mixed-frequency data, Unrestricted VAR-MIDAS.

3.1 Introduction

Unit roots have important implications for econometric tools and economic interpretation, as shown in a seminal article by Granger and Newbold (1974). Hence, testing the data for unit roots prior to modelling has become a standard procedure.

Several tests have been proposed to assess the existence of unit roots (see, for example, Schwert, 1989, Stock, 1986 and Haldrup and Jansson, 2006 for reviews on the topic). The most widely used is the Dickey-Fuller test (DF hereafter); see Dickey and Fuller (1979). However, other tests exist with better power performance. Large power gains can be achieved by extending unit root tests to include correlated stationary covariates. For instance, Hansen (1995) proposed generalis-

ing the DF tests to include stationary covariates — the so-called CADF test — in order to exploit the correlation between the weakly exogenous stationary covariate and the variable of interest.¹ The higher the correlation between the two, the greater the potential power gains.

Elliott and Jansson (2003) extend the univariate results in Elliott et al. (1996), suggesting an alternative to the CADF test. This test (EJ test, hereafter) consists in a feasible point optimal test in the presence of deterministic components. In spite of having a slightly worse size performance, the authors concluded that the EJ test outperforms the CADF test in terms of power.

The aim of this article is to go one step further and include mixed-frequency data in the EJ testing framework. We assess whether this extension improves the size and the power of the test. In practice, the benefits of covariate-augmented unit root tests may be hampered when the variables involved are sampled at different/mixed frequencies. The typical approach of temporally aggregating high-frequency variables to the same, low frequency as the variable of interest can affect the correlation between the variables (see Silvestrini and Veredas (2008) for a survey on temporal aggregation and its implications). Hence, combining a low-frequency dependent variable with a high-frequency covariate may entail power gains.

To deal with mixed-frequency data we use the MI(xed) DA(ta) S(ampling) framework (for a brief overview of the main topics related with MIDAS regressions see, for example, Andreou et al., 2011). The flexible, data-driven and typically parsimonious MIDAS weighting schemes are commonly used in a single equation framework. In this article, for merging the EJ test with mixed-frequencies we use the recent MIDAS vector autoregressive (VAR) approach. Hence, building on Ghysels (2012), we propose an unconstrained, though parsimonious, stacked skip-sampled reduced-form VAR model, which is estimated using standard econometric techniques. The performance of the EJ test, with low- or mixed-frequency data is compared with the one of CADF tests with low-frequency data (Hansen, 1995) and with mixed-frequency data (Chapter 2 and Duarte, 2014).

We show that mixed-frequency data helps improving the power performance of

¹This article focuses on unit root tests with non-stationarity as the null hypothesis. See Jansson (2004) for a unit root test with covariates where the null hypothesis is stationarity.

EJ tests. Moreover, the EJ-family of tests has a better power performance than the ADF-family of tests, either with low or mixed-frequency data. The gains are robust to the sample size and to lag selection.

The remainder of this article is organised as follows. Section 3.2 summarises the EJ covariate-augmented unit root test, while Section 3.3 describes the new mixed-frequency approach to the EJ test proposed in this article. Section 3.4 reports a Monte Carlo study on the power implications of this new approach and Section 3.5 compares the performance of the alternative tests for testing the presence of a unit root in US unemployment rate. Finally, Section 3.6 concludes.

3.2 Feasible point optimal unit root tests

This section presents the feasible point optimal unit root tests available in the literature: the univariate test in Elliott et al. (1996) (Section 3.2.1) and the covariate-augmented EJ test suggested by Elliott and Jansson (2003) (Section 3.2.2). For the sake of simplicity, the notation was developed for the case of a single covariate but can be readily extended for multiple covariates.

3.2.1 The ERS test

Aiming at improving the power of univariate unit root tests, Elliott et al. (1996) proposed the, henceforth, ERS tests. Based on strong assumptions regarding the knowledge of the distribution of the data and making use of the Neyman-Pearson Lemma, Elliott et al. (1996) started by deriving an optimal test against any point alternative. Through this optimal test, it is possible to define the power envelope, which is the upper bound for the power function of any test based on the same likelihood. The authors then dropped some of the initial assumptions in order to obtain a feasible test that exhibited the same large-sample properties as the optimal test.

Following Elliott et al. (1996), assume that the variable of interest, Y_t , is the sum of a deterministic component, $d_{Y,t}$, and a stochastic component, $u_{Y,t}$, such as

$$Y_t = d_{Y,t} + u_{Y,t} \tag{3.1}$$

CHAPTER 3. UNIT ROOT TESTS USING MIXED-FREQUENCY VAR
MODELS

where the deterministic component can be modelled as a linear combination of a set of non-random regressors, such as $d_{Y,t} = z_t' \beta$, with $\beta = [\beta_{Y,0} \ \beta_{X,0}]'$, $z_t' = [1 \ t]$ and t denoting a linear trend. In addition,

$$u_{Y,t} = \alpha u_{Y,t-1} + v_{Y,t} \tag{3.2}$$

where $v_{Y,t}$ is an unobserved stationary zero mean error process with constant variance.

Considering $\alpha = 1 + c/T$ and $\bar{\alpha} = 1 + \bar{c}/T$, the aim of the ERS test is to assess whether $\alpha = 1$, or $c = 0$ (null hypothesis), against the local point alternative that $\alpha = \bar{\alpha} < 1$, or $c = \bar{c} < 0$, with c, \bar{c} fixed. The test statistic takes the form of

$$ERS = [S(\bar{\alpha}) - \bar{\alpha}S(1)]/\hat{\sigma}_v^2 \tag{3.3}$$

where $S(a)$ is the sum of squared residuals from a least squares regression of $Y(a)$ on $z(a)$, with $Y(a) = (Y_1, Y_2 - aY_1, \dots, Y_T - aY_{T-1})$ and $z(a) = (z_1, z_2 - az_1, \dots, z_T - az_{T-1})$. This process of transformation of the variables is also known as GLS demeaning/detrending. For any given \bar{c} , this test rejects the null hypothesis for small values of (3.3), yielding the most powerful test against the point alternative hypothesis, and can be interpreted as a simplified version of a likelihood ratio statistic. The test distribution in the case of no deterministic variables or only including a constant is

$$ERS \Rightarrow \bar{c}^2 \int_0^1 W_c^2 - \bar{c}W_c^2(1) \tag{3.4}$$

where W_c is a Ornstein-Uhlenbeck process, such as $W_c(t) = \int_0^t e^{c(t-s)} dW_0(s)$, which satisfies the stochastic differential equation $dW_c(t) = cW_c(t)dt + dW_0(t)$ with initial condition $W_c(0) = 0$ and W_0 is a standard Brownian motion, defined on $[0, 1]$.²

Furthermore, $\hat{\sigma}_v^2$ is a consistent estimate of σ_v^2 , which is the long-run variance of $v_{Y,t}$ (i.e., a consistent estimate of the spectral density at the zero frequency). Similarly to Hansen (1995), in Elliott et al. (1996) a non-parametric estimator of

²See Elliott et al. (1996) for the case of also including a time trend.

σ_v^2 is used, which is given by

$$\hat{\sigma}_v^2 = \sum_{k=-M}^M w(k/M) \frac{1}{T} \sum_{t=1}^{T-m} \hat{v}_{Y,t} \hat{v}_{Y,t+k} \quad (3.5)$$

where the function $w(\cdot)$ is a kernel weight function, such as the Bartlett or Parzen kernel, M is the bandwidth selected to grow slowly with sample size (Andrews, 1991, Jansson, 2002), and $\hat{v}_{Y,t}$ is an estimate of $v_{Y,t}$ (i.e., the residual from an OLS regression of Y_t on Y_{t-1} and z_t).³

Elliott et al. (1996) show that the feasible test in (3.3) reaches the same asymptotic power functions as the optimal test at a specific point (\bar{c}) and remains very close to the power envelope otherwise, regardless of the deterministic terms included. This performance contrasts with the one of alternative tests, such as DF or ADF tests, whose asymptotic power function matches the power envelope only when no deterministic terms are included. Finite sample simulation results also corroborate these findings. The DF test has relatively good size properties in small-sample Monte Carlo studies. This test also has a good power performance when there are no deterministic terms. In contrast, if deterministic terms have to be estimated, the ERS test clearly performs better than the alternatives in terms of power.

The literature shows that values of \bar{c} associated with an asymptotic power of one half yield tests with power functions tangent to the power envelope at that value, and close to the power envelope over a considerable range of alternative values. Simulation results in Elliott et al. (1996) show that when a trend is included considering \bar{c} equal to -13.5 yields a power function tangent to the power envelope at 0.5. If only a constant is included, then the same happens for $\bar{c} = -7$. Hence, in order to implement the test, Elliott et al. (1996) suggest fixing $\bar{c} = \bar{c}^* = -7$ for the constant case and $\bar{c}^* = -13.5$ for the linear trend case.

³The same authors also consider a parametric estimator of σ_v^2 , following Berk (1974), which is an autoregressive spectral estimator of the long-run variance. This alternative estimator was not used in this article; for more details see Elliott et al. (1996).

3.2.2 EJ test

As mentioned before, the EJ test is an extension of the ERS test in order to include covariates. Assume that the stationary covariate series, X_t , can be expressed as

$$X_t = d_{X,t} + u_{X,t}. \quad (3.6)$$

Hence, consider the VAR model formulation

$$\begin{bmatrix} Y_t \\ X_t \end{bmatrix} = d_t + u_t, \quad (3.7)$$

where $d_t = z_t' \beta$, $\beta = [\beta_{Y,0} \ \beta_{X,0} \ \beta_{Y,1} \ \beta_{X,1}]'$,

$$z_t' = \begin{bmatrix} 1 & 0 & t & 0 \\ 0 & 1 & 0 & t \end{bmatrix} \text{ and } u_t = \begin{bmatrix} u_{Y,t} \\ u_{X,t} \end{bmatrix}. \quad (3.8)$$

Five different combinations of the deterministic variables can be considered:

- Case 1: $\beta_{Y,0} = \beta_{Y,1} = \beta_{X,0} = \beta_{X,1} = 0$;
- Case 2: $\beta_{Y,1} = \beta_{X,0} = \beta_{X,1} = 0$;
- Case 3: $\beta_{Y,1} = \beta_{X,1} = 0$; (3.9)
- Case 4: $\beta_{X,1} = 0$;
- Case 5: No restrictions.

As suggested in Elliott and Jansson (2003), these five different cases can also be represented as restrictions on β , such as $(I_4 - S_i)\beta = 0$, $i = 1, \dots, 5$, where I_4 is a (4×4) identity matrix and S_i are (4×4) matrices as follows

$$S_1 = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}; S_2 = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}; S_3 = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix};$$

$$S_4 = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \text{ and } S_5 = I_4. \quad (3.10)$$

In addition, the stochastic component u_t can be rearranged as

$$A(L) \begin{bmatrix} (1 - \alpha L)u_{Y,t} \\ u_{X,t} \end{bmatrix} = e_t \quad (3.11)$$

where $A(L)$ is a matrix polynomial of order k in the lag operator L and $e_t = [e_{Y,t} \ e_{X,t}]'$. Consider equations (3.7) and (3.11) and define $\Gamma_t(\alpha) = [(1 - \alpha L)Y_t \ X_t]'$, with $\Gamma_1(\alpha) = [Y_1 \ X_1]'$.

In the following analysis we assume that:

Assumption 1: The roots of $|A(z)|$ lie outside the unit circle;

Assumption 2: $u_0, u_{-1}, \dots, u_{-k}$ are $O_p(1)$;

Assumption 3: $E_{t-1}(e_t) = 0$, $E_{t-1}(e_t e_t') = \Sigma$ where Σ is positive definite, and $\sup_t E \|e_t\|^{2+\kappa} < \infty$, for some $\kappa > 0$,

where E_{t-1} denotes the conditional expectation based on the information set covering e_{t-1}, e_{t-2}, \dots . Assumption 1 is a standard stationarity condition. Assumption 2 implies that the initial values are asymptotically negligible and Assumption 3 implies that e_t satisfies a multivariate functional central limit theorem (Phillips, 1987).

The aim of the EJ test is to assess whether $H_0 : \alpha = 1$, meaning that Y_t has a unit root, against the local alternative that $H_1 : \alpha = 1 + c/T = \bar{\alpha} = 1 + \bar{c}/T < 1$, with c, \bar{c} fixed, exploiting the information in X_t . As in Elliott et al. (1996), resorting to the Neyman-Pearson Lemma and then relaxing some strong assumptions, Elliott and Jansson (2003) propose a family of feasible covariate-augmented unit root tests, which attain the power envelope at specific points (\bar{c}). While Hansen's CADF test is nearly efficient only in the absence of deterministic terms, the EJ

CHAPTER 3. UNIT ROOT TESTS USING MIXED-FREQUENCY VAR MODELS

tests are feasible point optimal tests also in the presence of deterministic components.

Consider r indicating the parameter α under the null and alternative hypothesis. Hence, for $r = (1, \bar{\alpha})$, the test statistic is

$$EJ = T(\text{tr}[\tilde{\Sigma}(1)^{-1}\tilde{\Sigma}(\bar{\alpha})] - (1 + \bar{\alpha})) \quad (3.12)$$

where

$$\tilde{\Sigma}(r) = T^{-1} \sum_{t=k+1}^T \tilde{e}_t(r)\tilde{e}_t(r)', \quad (3.13)$$

$\bar{\alpha} = 1 + \bar{c}/T$, $\bar{c} < 0$ and fixed, and $\tilde{e}_t(r)$ are the estimated residuals from the VAR model $\tilde{e}_t(r) = \tilde{A}(L)\tilde{v}_t^i(r)$, for $i = 1, \dots, 5$ corresponding to the alternative deterministic restrictions described in 3.9. The variables in this VAR model result from detrending the original data, both under the null and the alternative hypothesis, i.e. $\tilde{v}_t^i(r) = \Gamma_t(r) - z_t(r)'\tilde{\beta}^i(r)$, where

$$\begin{aligned} \tilde{\beta}^i(r) &= \left[S_i \left(\sum_{t=1}^T z_t(r)\hat{\Omega}^{-1}z_t(r)' \right) S_i \right]^{-1} \left[S_i \sum_{t=1}^T z_t(r)\hat{\Omega}^{-1}\Gamma_t(r) \right], \\ z_t(r)' &= \begin{bmatrix} 1-r & 0 & (1-rL)t & 0 \\ 0 & 1 & 0 & t \end{bmatrix} \text{ for } t > 1 \text{ and} \\ z_1(r)' &= \begin{bmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{bmatrix}. \end{aligned} \quad (3.14)$$

Note that M^- denotes the Moore Penrose inverse of M . Furthermore, the estimate $\hat{\Omega}$ is obtained from the residuals of the VAR model $A(L)\Gamma_t(1) = \text{deterministics} + e_t$, i.e.,

$$\hat{\Omega} = \begin{bmatrix} \hat{\omega}_{YY} & \hat{\omega}_{YX} \\ \hat{\omega}_{YX} & \hat{\omega}_{XX} \end{bmatrix} = \hat{A}(1)^{-1} \left(T^{-1} \sum_{t=k+1}^T \hat{e}_t\hat{e}_t' \right) \hat{A}'(1)^{-1}, \quad (3.15)$$

with $\hat{A}(1) = I_2 + \sum_{j=1}^k \hat{A}_j$ and \hat{A}_j is the estimated matrix element of $A(L)$ corresponding to lag j . Elliott and Jansson (2003) suggest choosing the number

of lags according to theory or some information criteria (e.g., BIC).

From the estimates of $\hat{\Omega}$, one can readily obtain the estimated value of R^2 , i.e.,

$$R^2 = \frac{\omega_{YX}\omega_{XX}^{-1}\omega_{YX}}{\omega_{YY}}, \quad (3.16)$$

which represents the frequency zero correlation between the shocks to stationary X_t and the quasi-differences of Y_t .

For the case of no deterministic variables, the asymptotic test distribution is

$$EJ \Rightarrow \bar{c}^2 \int_0^1 W_{1c}^2 - \bar{c}W_{1c}^2(1) + (\bar{c}^2 - 2\bar{c}c)Q \int_0^1 W_{1c}^2 + 2\bar{c}Q^{1/2} \int_0^1 W_{1c}dW_2 \quad (3.17)$$

where $Q = R^2/(1 - R^2)$, W_{1c} is an Ornstein-Uhlenbeck process, such that $W_{1c}(r) = c \int_0^r e^{c(r-s)}W_1(s)ds + W_1(r)$, which is equivalent to $W_{1c}(r) = \int_0^r e^{c(r-s)}dW_1(s)$ (Phillips, 1987), and W_1 and W_2 are independent univariate standard Brownian motions.⁴ The distribution of the test statistic depends on R^2 . If there is no long-run correlation, then R^2 equals zero and the EJ test resumes to its univariate version, the ERS test. Conversely, as the correlation increases, $R^2 \rightarrow 1$ and the distributions are different. Information is lost when the covariates are ignored. Higher R^2 are associated with more powerful tests.

Similarly to the ERS test, the critical values of the EJ test also depend on the fixed alternative chosen. Although the choice of $\bar{c} = \bar{c}^*$ can be conditional on R^2 , Elliott and Jansson (2003) decided to use the values proposed by Elliott et al. (1996) ($\bar{c}^* = -7$ for cases 1 to 3 and $\bar{c}^* = -13.5$ for cases 4 and 5). Given that as R^2 increases the power of the test improves, these figures reflect the worst case scenarios.

Considering these values for \bar{c}^* , Elliott and Jansson (2003) reported asymptotic critical values for a significance level of 5 per cent and different values of R^2 (from 0 to 0.9); see Table C.1 of the Appendix.

The small-sample evaluation in Elliott and Jansson (2003) showed that when $R^2 > 0$ the power of the unit root test when using covariates is higher than when the relevant covariates are ignored. Furthermore, the EJ test outperforms the

⁴Regarding the deterministic variables, for the other cases see Elliott and Jansson (2003).

CADF test in terms of power, although its performance in terms of size is slightly worse.

3.3 Mixed-frequency EJ unit root test

In this article, we extend the covariate-augmented EJ unit root test to include mixed-frequency data and assess its performance in terms of size and power. The idea behind merging these two strands of the literature — covariate-augmented unit root testing and mixed-frequency models — is that by combining high and low frequency data one may uncover stronger correlations and, thus, may strengthen the power performance of unit root tests.

We use the MIDAS framework to deal with mixed-frequency data. In particular, extending the EJ tests requires using a mixed-frequency VAR approach. We propose a new, simple and unrestricted VAR-MIDAS approach, which will be briefly described in the next section. We then apply this new estimation method to the mixed-frequency EJ unit root test (M-EJ test, henceforth) in Section 3.3.2.

3.3.1 VAR-MIDAS approach

Typically, the mixed-frequency VAR model is defined at the highest sampling frequency. The high-frequency observations of the low-frequency variables are treated as missing observations. The high-frequency VAR is cast in a state-space form. The Kalman filter is used to estimate the model, tackling the missing values and taking into account the mixed-frequency nature of the data. Examples of applications of these models are Zdrozny (1988), Kuzin et al. (2011) and, for factor models, Mariano and Murasawa (2003) and Nunes (2005); see Forni et al. (2013) for a brief survey. Instead of using this kind of frequentist/classical approach, some studies (e.g., Qian, 2010 and Chiu et al., 2011) suggested the use of a bayesian approach.

Note that this strand of the literature delivers results in the high frequency — latent dependent variable, latent shocks and latent impulse response functions come as a by-product. Alternatively, Ghysels (2012) proposed a simpler version of mixed sampling frequency VAR models where latent variables and shocks are

not involved. Inspired by periodic models (Hansen and Sargent, 2013), this VAR model only uses observed data, being based on stacked skip-sampled versions of the high-frequency variables. That is, each of the m high-frequency periods is treated as a separate variable. Although latent variables or shocks are not included, this specification can account for the impact of high-frequency data onto low-frequency data and vice-versa.

Consider a low-frequency variable, Y_t , and a high-frequency variable, x_t , which has a time frequency m times higher. Following Ghysels (2012), the finite order reduced-form VAR representation in a stacked form is

$$\begin{bmatrix} x_t^{(1)} \\ \vdots \\ x_t^{(m)} \\ Y_t \end{bmatrix} = d_t + \sum_{j=1}^J \Psi_j \begin{bmatrix} x_{t-j}^{(1)} \\ \vdots \\ x_{t-j}^{(m)} \\ Y_{t-j} \end{bmatrix} + \varepsilon_t, \quad (3.18)$$

where d_t denotes the deterministic component, Ψ_j are the coefficient matrices associated with the J low-frequency lags, $x_t^{(m)}$ is the skip-sampled version of high-frequency x_t , and ε_t is a standard *iid* error term. Given the stacked form, this VAR representation can only be implemented with fixed m , i.e., with the same number of high-frequency observations for each low-frequency observation throughout the whole sample.

The unconstrained VAR model in (3.18) can be estimated using standard VAR estimation tools as in Hamilton (1994); see, for example, Ghysels and Miller (2013) and Hecq et al. (2013) for an application to cointegration tests, and Götz and Hecq (2014b) for an analysis of instantaneous and Granger causality. However, as in traditional low-frequency VAR models, parameter proliferation can also be an issue in the estimation of mixed-frequency VAR models. With only two variables (Y_t and x_t) the number of parameters to be estimated equals $(m + 1) \times (m + 1) \times J$ plus the coefficients associated with the deterministic variables. For example, if $m = 3$ (mimicking the case of quarterly/monthly variables) and $J = 3$ there will be at least 64 parameters to be estimated. The number of parameters to be estimated can, therefore, easily exceed the typical size of available samples. Although mixed frequencies are at play, notice that the number of available observations,

i.e., sample size, is still determined by the low-frequency variable.

So far, to tackle the parsimony issue, two alternatives have been considered. Firstly, Ghysels (2012) suggested using MIDAS data-driven weighting functions in this equation. The author noted that the mixed-frequency VAR model is a multivariate extension of MIDAS regressions — the last equation can be seen as a MIDAS regression augmented with autoregressive terms, without contemporaneous observations of the high-frequency variable. Moreover, given that stacking the high-frequency data may be nothing more than repeating similar dynamics m times, the Ψ matrices can be restricted accordingly.

Secondly, Ghysels (2012), McCracken et al. (2014) and Götz and Hecq (2014a) also consider a bayesian approach, which easily accommodates the potentially large set of parameters to be estimated.

In this article we want to estimate a linear, non-bayesian, unrestricted VAR-MIDAS model using standard econometric techniques. Our aim is to have a parsimonious specification with a number of parameters that does not depend on m . To avoid parameter proliferation and building on the suggestion of Ghysels (2012), we propose treating all equations in the VAR model as MIDAS regressions. Then, a VAR-MIDAS model can be written as

$$\begin{bmatrix} x_t^{(1)} \\ \vdots \\ x_t^{(m)} \\ Y_t \end{bmatrix} = d_t + \sum_{j=1}^J \begin{bmatrix} \psi_{1,j}^{(1)} & \cdots & \psi_{m,j}^{(1)} & \phi_{1,j} \\ \vdots & & \vdots & \vdots \\ \psi_{1,j}^{(m)} & \cdots & \psi_{m,j}^{(m)} & \phi_{m,j} \\ \psi_{1,j}^{(m+1)} & \cdots & \psi_{m,j}^{(m+1)} & \phi_{m+1,j} \end{bmatrix} \begin{bmatrix} x_{t-j}^{(1)} \\ \vdots \\ x_{t-j}^{(m)} \\ Y_{t-j} \end{bmatrix} + \varepsilon_t, \quad (3.19)$$

where $\Psi^{(i)} = \{\psi_{1,1}^{(i)}, \dots, \psi_{m,1}^{(i)}, \dots, \psi_{1,J}^{(i)}, \dots, \psi_{m,J}^{(i)}\}$ for $i = 1, \dots, m + 1$ are the coefficients associated with the high-frequency variable and its lags, which are captured by an aggregation weighting scheme.

As already pointed out by Ghysels (2012), the last equation in (3.19) resembles a standard MIDAS equation. In this article we suggest that the other equations can also be seen as a univariate version of a MIDAS regression (except for the low-frequency variable), in the sense that the weighting schemes are applied to lags of the dependent variable.⁵

⁵Meanwhile, a similar approach, developed independently, was also proposed by Mikosch and

Several weighting functions can be chosen. Ghysels et al. (2007) considered two alternatives — the exponential Almon lag and the beta polynomial. In both the weights are determined by a few hyperparameters and the functional specification is nonlinear. In addition, there is the aggregation scheme underlying the unrestricted MIDAS regressions (U-MIDAS), used in Marcellino and Schumacher (2010), Forni and Marcellino (2012) and Forni et al. (2011). Although linear, the drawback of this scheme is being prone to dimensionality issues, namely when the difference between the low and the high frequency, i.e. m , is large.

We focus on the traditional Almon lag polynomial. This weighting scheme has features that are particularly appealing for our purpose. Firstly, it is a very flexible scheme, allowing different shapes, such as increasing, decreasing or hump-shape. Secondly, Almon polynomials are linear and can thus be estimated by OLS. Finally, regressions using this polynomial can be tightly parameterised. The Almon aggregation scheme assumes that g (in this case, $g = m \times J$) lag weights can be related to d linearly estimable underlying parameters, with $d < g$, as follows:

$$a_{h,j}^{(i)} = \sum_{k=1}^d \theta_k^{(i)} g^k \quad (3.20)$$

where $\theta_k^{(i)}$ denotes the hyperparameters, $i = 1, \dots, m + 1$, $h = 1, \dots, m + 1$, and $k = 1, \dots, d$. The number of parameters to be estimated in (3.19) equals $(m + 1) \times J$ (the ϕ parameters) plus $(m + 1) \times d$ (the a parameters) plus the coefficients associated with deterministic variables, in case they are included. With $m = 3$, $J = 3$ and $d = 2$ there will be 20 parameters to be estimated (plus the coefficients of the deterministic variables), which is a more reasonable number given the typical size of available samples.

3.3.2 M-EJ test

The new M-EJ test proposed in this paper is similar to the EJ test described in Section 3.2.2. Except for the fact that the covariate is now a high-frequency variable, with a frequency m times higher than the dependent variable. To take

Neuwirth (2015).

into account this change, instead of estimating a traditional VAR model, we use the new VAR-MIDAS approach as described in the previous section. In particular,

$$\begin{bmatrix} x_t^{(1)} \\ \vdots \\ x_t^{(m)} \\ Y_t \end{bmatrix} = d_t + \sum_{j=1}^J \begin{bmatrix} \Psi_j^{(1)}(\theta) & \vdots & \Psi_j^{(m)}(\theta) & \Psi_j^{(m+1)}(\theta) \\ \vdots & \vdots & \vdots & \vdots \\ \phi_{1,j} & \vdots & \phi_{m,j} & \phi_{m+1,j} \end{bmatrix} \begin{bmatrix} x_{t-j}^{(1)} \\ \vdots \\ x_{t-j}^{(m)} \\ Y_{t-j} \end{bmatrix} + \varepsilon_t, \quad (3.21)$$

where $\Psi^{(i)}(\theta)$ with $i = 1, \dots, m+1$ are the MIDAS weighting schemes. The advantages of this VAR-MIDAS are (i) the fact that the estimation uses standard econometric techniques and (ii) the number of parameters does not depend on m . The estimation of this model delivers the estimates of the error series used to calculate the Ω matrix in (3.15) and the R^2 parameter in (3.16), which should now be interpreted as a multiple correlation coefficient. The estimate of the latter parameter will benefit from exploiting the high-frequency information on the covariate.

Note that the covariate used in the mixed-frequency tests — a skip-sampled version of high-frequency x_t — is still a stationary variable. As shown by Granger and Siklos (1995) and Marcellino (1999), among others, zero frequency unit roots are not affected by temporal aggregation.

The new mixed-frequency EJ test simply consists in taking into account m stationary covariates, one for each high-frequency observation of the covariate. Hence, under Assumptions 1 to 3, the test distribution is as in (3.17). Similarly to the ERS and EJ tests, the M-EJ test also depends on \bar{c}^* . As in Elliott and Jansson (2003) and to ensure comparability with the existing literature, the figures used for \bar{c}^* are -7 for cases 1 to 3 and -13.5 for cases 4 and 5.

3.4 Monte Carlo simulation

The finite sample size and power performance of the proposed M-EJ unit root test was investigated by setting up a Monte Carlo simulation exercise. Following Hansen (1995), Elliott et al. (2005) and Galvao Jr. (2009), this exercise considers

the following data generating process (DGP)

$$\begin{aligned} y_t &= d_{y,t} + \alpha y_{t-1} + v_{y,t} \\ x_t &= d_{x,t} + v_{x,t} \end{aligned} \tag{3.22}$$

where y_t is potentially non-stationary, Δx_t is a stationary process, $\alpha = 1 + c/T$ and d_t represents the deterministic terms. Both variables are in the same (high) time frequency.

Regarding the deterministic terms, Case 3 (constant for both variables) and Case 5 (constant and time trend for both variables) were considered; see Elliott and Jansson (2003) and Juhl and Xiao (2003). The error process $v_t = [v_{y,t} \ v_{x,t}]'$ is generated by a VARMA model $A(L)v_t = B(L)\xi_t$, where $A(L) = I_2 - AL$, $B(L) = I_2 + BL$,

$$A = \begin{bmatrix} a_1 & a_2 \\ a_2 & a_1 \end{bmatrix}, \quad B = \begin{bmatrix} b_1 & b_2 \\ b_2 & b_1 \end{bmatrix}, \tag{3.23}$$

and $\xi_t \sim N(0, \Sigma)$, where I_2 is a (2×2) identity matrix and Σ is such that the long-run variance matrix of v_t satisfies

$$\Omega = (I_2 - AL)^{-1}(I_2 + BL)\Sigma(I_2 + BL)'(I_2 - AL)^{-1'} = \begin{bmatrix} 1 & R \\ R & 1 \end{bmatrix} \tag{3.24}$$

and, as before, $R^2 \in [0, 1]$. Various values for R^2 are examined, namely $R^2 = \{0.2, 0.5, 0.8\}$. Note that this DGP does not include seasonal features. We make this assumption in order to exclude the possibility that the unit root at the zero frequency arises from temporal aggregation of a series which, in fact, has a unit root at some seasonal frequencies (Granger and Siklos, 1995). The initial condition was set at zero Müller and Elliott, 2003.⁶

The low-frequency series Y_t and X_t are obtained by aggregating the generated high-frequency data y_t and x_t . The aggregating scheme considered is skip-sampling

⁶A sensitivity analysis to this assumption was performed and the results indicate that the main result — better power performance of mixed-frequency covariate-augmented unit root tests — does not qualitatively change if a different initial value was considered.

(as for stock variables). As noted by Working (1960), if the aggregation scheme is skip-sampling then the aggregated error term is not serially correlated. In contrast, when the aggregation scheme is some kind of averaging Working (1960) showed that serial correlation arises and cannot be dealt with by using a high-order autoregressive polynomial.

We assume $m = 12$, mimicking the combination of annual and monthly data. The sample size is set at $T = 100$ and 10,000 replications are used. Below, we present a sensitivity analysis for a larger T , namely $T = 500$.

As shown in Pierse and Snell (1995), the aggregation process affects the value of α tested in the alternative hypothesis, yielding $\alpha^{aggregate} = \alpha^m$, whether the aggregated variable is a flow or a stock. Hence, apart from the null hypothesis ($\alpha = 1$), we consider three alternative values for α (and, consequently, for c), such that their aggregate counterparts equal 0.95, 0.90 and 0.85.

The number of lags is assumed unknown, replicating what happens in practice. In Elliott et al. (1996) the BIC is used to select the truncation lag, with a lower bound of 3 lags, because even larger size distortions would have resulted if zero was the lower bound. Elliott and Jansson (2003) suggested choosing the lag length of the VAR through theory or a consistent information criterion, such as the BIC. Also for the EJ test, Fossati (2012) suggested using the MAIC proposed by Qu and Perron (2007), which showed that this criterion leads to VAR models that yield cointegration tests with better size properties and little power loss.

In this article we used the AIC for choosing the number of lags.⁷ In simulations not reported but available upon request, we concluded that this criterion led to the best size performance, compared with BIC and MAIC from Qu and Perron (2007). The maximum number of lags is given by the rule $int(12(T/100)^{0.25})$ (Schwert, 1989).

Table 3.4.1 shows the range of simulation designs, as well as the number of lags chosen, under the null hypothesis, for each DGP. The values for the initial R^2 refer to the high-frequency processes. Different values were used for the persistence in the high-frequency autoregressive dynamics of the dependent variable, namely $a_1 = \{0.2, 0.5, 0.8\}$, and for the high-frequency moving average dynamics, $b_1 =$

⁷The number of lags was chosen with the low frequency dataset.

CHAPTER 3. UNIT ROOT TESTS USING MIXED-FREQUENCY VAR
MODELS

Table 3.4.1: Simulation design and median of the lag order selected by AIC

	Simulation design					Median of the lag order			
	Error structure				Initial	$T = 100$		$T = 500$	
	a_1	a_2	b_1	b_2	R^2	Constant	Trend	Constant	Trend
1	0	0	0	0	0.20	0	0	0	1
2	0	0	0	0	0.50	0	0	0	1
3	0	0	0	0	0.80	0	0	0	1
4	0.2	0	0	0	0.20	0	0	0	0
5	0.2	0	0	0	0.50	0	0	0	0
6	0.2	0	0	0	0.80	0	0	0	0
7	0.5	0	0	0	0.20	0	0	0	0
8	0.5	0	0	0	0.50	0	0	1	0
9	0.5	0	0	0	0.80	0	0	1	0
10	0.8	0	0	0	0.20	1	1	1	2
11	0.8	0	0	0	0.50	1	1	1	2
12	0.8	0	0	0	0.80	1	1	1	2
13	0.2	0.2	0	0	0.20	0	0	0	0
14	0.2	0.2	0	0	0.50	0	0	1	0
15	0.2	0.2	0	0	0.80	0	0	1	0
16	0	0	-0.2	0	0.20	0	0	0	1
17	0	0	-0.2	0	0.50	0	0	0	1
18	0	0	-0.2	0	0.80	0	0	0	1
19	0	0	-0.5	0	0.20	0	0	1	2
20	0	0	-0.5	0	0.50	0	0	1	2
21	0	0	-0.5	0	0.80	1	1	1	2
22	0	0	-0.8	0	0.20	1	1	2	6
23	0	0	-0.8	0	0.50	1	1	2	6
24	0	0	-0.8	0	0.80	1	1	2	6

$\{-0.2, -0.5, -0.8\}$.⁸

⁸Only negative figures were considered for the moving average dynamics because they represent the most difficult cases in terms of size distortions, as the reversion to the mean is higher. Regarding autoregressive dynamics, other processes were tested, namely with a higher persistence in the covariate. The results remain qualitatively unchanged. So, for the sake of

For the ERS test, heteroskedasticity and autocorrelation consistent (HAC) estimates of the elements of the long-run variance-covariance matrix Ω are used. This calculation commonly involves the use of pre-whitening filters based on simple autoregressive models. This procedure may induce bias in the estimation of autoregressive coefficients, which is transmitted to the recolouring filter. To mitigate the potential bias associated with these filters, recursive demeaning/detrending procedures were assessed, as in Taylor (2002), Sul et al. (2005) and Rodrigues (2006).

We use finite-sample critical values to implement the unit root tests. For $\alpha = 1$, the observed rejection rates of each test were based on critical values from the limiting distribution obtained for the DGP with the simplest dynamics, i.e., DGP 1, 2 and 3 in Table 3.4.1. For $\alpha < 1$ the size-adjusted power of the tests was based on critical values estimated from the simulated data generated under the null hypothesis ($\alpha = 1$).⁹ Following the suggestion in Elliott and Jansson (2003), the critical values were interpolated for the estimated figures of R^2 .

3.4.1 Baseline results

Table 3.4.2 reports the probability of rejecting the null hypothesis under the unit root case, i.e., the finite sample size for unit root tests considering a nominal size of 5 per cent.¹⁰

For comparison reasons, the CADF-family of unit root tests, with both low- and mixed-frequency data, was also computed. For more details on the implementation of these tests, including the choice of lags, see Chapter 2 or Duarte (2014).

Size distortions tend to be larger when there is a linear trend in the regression. Moreover, the size performance is also negatively affected in the presence

brevity, those results will not be reported, but are available from the author upon request.

⁹As mentioned by Haug (2002), size-unadjusted power is rather misleading, so those results are not reported.

¹⁰The codes were written in Matlab. Some functions were taken from the Econometrics Toolbox by James P. LeSage (<http://www.spatial-econometrics.com>). The procedure to perform the CADF unit root tests was greatly inspired in the code made available by Bruce E. Hansen (http://www.ssc.wisc.edu/~bhansen/progs/et_95.html). The procedure to perform the EJ unit root tests was greatly inspired in the code made available by Michael Jansson. The MIDAS toolbox was inspired in a code kindly provided by Arthur Sinko.

CHAPTER 3. UNIT ROOT TESTS USING MIXED-FREQUENCY VAR
MODELS

Table 3.4.2: Finite sample size for unit root tests considering nominal size of 5 per cent, $T = 100$

DGP		Constant						Trend					
	Initial R^2	CADF	CADF	EJ	M	M	M	CADF	CADF	EJ	M	M	M
			GLS		CADF	CADF	EJ		GLS		CADF	CADF	EJ
1	0.20	5.0	5.0	5.1	5.1	5.0	5.5	5.1	5.1	5.1	5.2	5.2	5.9
2	0.50	5.0	5.0	4.9	5.2	5.2	5.8	4.9	5.0	5.0	5.1	5.0	6.5
3	0.80	5.0	5.0	4.9	4.8	4.9	4.5	4.8	4.8	4.8	4.5	4.9	4.2
4	0.20	4.9	4.8	4.6	5.0	5.1	5.4	4.9	4.9	4.3	5.0	5.0	6.0
5	0.50	4.8	4.9	4.6	5.0	5.0	5.6	4.4	4.6	4.3	4.6	4.5	6.2
6	0.80	4.7	4.8	4.4	4.1	4.7	4.3	4.2	4.5	4.0	3.8	4.0	4.1
7	0.20	4.5	4.3	4.0	4.8	4.8	5.1	4.3	4.1	3.6	4.6	4.2	5.9
8	0.50	4.1	4.1	4.0	4.5	4.6	5.9	3.5	3.8	3.5	4.0	4.0	7.1
9	0.80	3.5	3.8	3.8	3.9	4.6	4.6	3.1	3.2	3.1	3.1	3.5	5.1
10	0.20	3.4	2.3	5.1	4.4	3.6	8.3	1.9	1.4	4.4	3.3	2.5	10.0
11	0.50	2.1	1.8	5.1	4.9	4.2	10.4	0.9	0.7	4.2	3.8	3.3	15.6
12	0.80	1.1	1.4	4.1	4.4	4.3	8.0	0.4	0.2	2.7	3.2	3.4	13.2
13	0.20	4.5	4.6	4.3	4.6	4.4	4.9	4.8	5.1	3.9	4.9	4.7	6.0
14	0.50	4.2	4.5	4.0	4.5	4.5	5.5	4.3	4.5	3.8	4.2	4.2	6.5
15	0.80	3.9	4.4	3.6	3.4	4.3	4.4	3.4	3.9	3.6	3.2	3.5	4.6
16	0.20	5.2	5.2	5.4	5.4	5.3	5.8	5.4	5.2	5.9	5.7	5.6	6.4
17	0.50	5.1	5.2	5.5	5.7	5.8	6.2	5.1	5.2	6.1	5.8	5.7	7.1
18	0.80	5.0	5.2	5.7	6.1	5.8	4.9	5.2	5.1	6.0	5.9	6.4	4.8
19	0.20	5.4	5.4	7.8	6.1	5.6	7.7	5.7	5.8	9.5	6.7	6.7	9.5
20	0.50	5.3	5.2	7.7	7.6	7.0	7.9	5.6	5.4	9.6	8.2	8.0	9.6
21	0.80	5.2	5.0	7.6	9.6	9.0	6.7	5.4	5.3	9.4	10.0	10.4	7.5
22	0.20	5.6	5.8	14.6	6.8	6.2	13.0	7.6	7.7	18.1	8.8	8.3	17.5
23	0.50	5.9	6.0	13.9	8.7	7.9	11.5	7.9	7.8	16.4	10.6	10.5	15.6
24	0.80	5.9	5.9	13.0	9.6	8.7	9.6	7.7	7.6	15.1	10.8	11.6	12.6

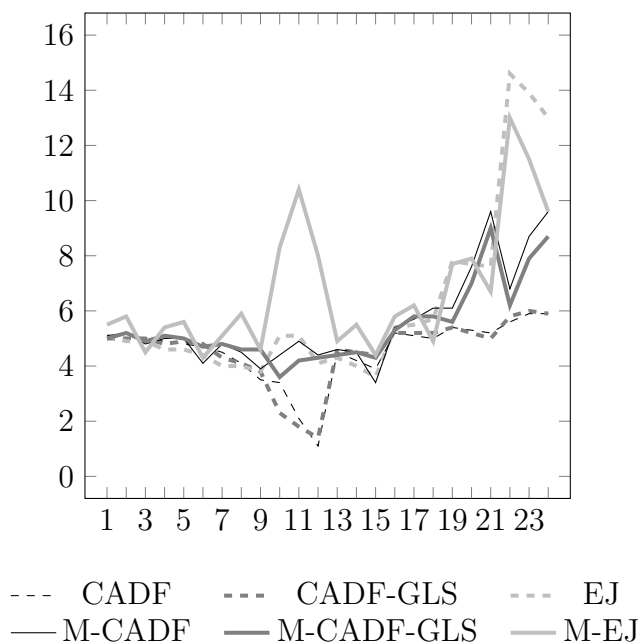


Figure 3.1: Finite sample size, $T = 100$, Constant only

of a negative moving average root, i.e., when b_1 is nonzero, which is a common result in the unit root literature (see, among others, Schwert, 1989). Downward distortions are mainly for stronger and more complex autoregressive dynamics, as also reported in Hansen (1995) and Galvão (2013).

In most cases the size differences between the two sets of tests — mixed- and low-frequency — are not substantial. Figure 3.1 shows the finite sample size of the unit root tests with nominal size of 5 per cent for each DGP.¹¹ When downward distortions exist, they tend to be less marked for mixed-frequency tests. For a DGP with a negative moving average root mixed-frequency tests tend to be more biased.

As in Elliott et al. (1996) and Elliott and Jansson (2003), feasible point optimal tests tend to be worse in terms of size performance than the ADF-family of tests, namely in the presence of large negative moving average roots. The M-EJ test is highly biased in the case of strong autoregressive dynamics. In contrast, it is less biased than the EJ test for the DGP with moving average dynamics, especially in

¹¹The figure shows the figures for the case of only a constant term; results are similar if also a time trend is included.

the presence of large negative moving average roots.

Tables 3.4.3 and 3.4.4 report the empirical rejection frequency of the null hypothesis under the alternative, i.e., the power of the unit root tests, for constant only and including a time trend, respectively. Recall that the power is size-adjusted.

In general, including a time trend in the test regression leads to lower power than when only a constant is included. As expected, the power of the tests increases as more correlated covariates are included in unit root test regressions. Furthermore, feasible point optimal tests tend to have a better power performance than the ADF-family of tests. This result is less clear-cut when a time trend is included, which suggests that the M-EJ test and, to a lesser extent, the EJ test may be penalised in shorter samples for the fact that they require the estimation of a less parsimonious VAR model, with more than one time trend.

Focusing on the impact of using mixed-frequency data, the size-adjusted power differences between the low- and mixed-frequency EJ tests, for the case of constant only, are presented in Figure 3.2. In general, these differences are positive, meaning that the power of mixed-frequency unit root tests is higher than that of low-frequency tests. Results are similar when a time trend is included.

One of the key features of mixed-frequency regressions is that they tend to capture the original correlation between the dependent variable and the covariate. In most cases, this correlation is hampered by time aggregation and, thus, the actual (and estimated) R^2 for the low-frequency tests differs significantly from the initial R^2 of the time disaggregated DGP, as shown in tables 3.4.3, 3.4.4 and 3.4.5. Results point to an overestimation of R^2 in the case of the M-EJ test, especially for lower correlations, which fades away as the sample size increases. A more detailed analysis of the results for a larger sample will be presented in the next section.

Table 3.4.5 shows also the finite-sample and asymptotic critical values for each test.¹² The asymptotic values were obtained from a simulation exercise with 1,500 observations and 60,000 replications. Separate exercises were performed for the low- and mixed-frequency approaches, delivering similar outcomes. Results are

¹²Table 3.4.5 shows the results for the DGP with the simplest dynamics, i.e., DGP 1, 2 and 3 in Table 3.4.1. Moreover, it covers the case of including a constant only. Results are qualitative similar if a time trend is also included and are available from the author upon request.

CHAPTER 3. UNIT ROOT TESTS USING MIXED-FREQUENCY VAR
MODELS

Table 3.4.3: Size-adjusted power of unit root tests, Constant only, $T = 100$

DGP	Initial R^2	Estimated R^2						$\alpha = 0.95$					
		CADF	CADF GLS	EJ	M CADF	M CADF GLS	M EJ	CADF	CADF GLS	EJ	M CADF	M CADF GLS	M EJ
1	0.20	0.05	0.05	0.04	0.22	0.23	0.31	7.8	26.1	28.6	13.6	31.1	34.9
2	0.50	0.07	0.07	0.06	0.45	0.50	0.56	8.1	26.6	30.1	29.1	46.9	54.4
3	0.80	0.09	0.10	0.08	0.66	0.77	0.81	8.1	27.6	30.7	63.8	74.1	82.8
4	0.20	0.05	0.06	0.04	0.22	0.23	0.31	7.8	26.7	29.8	13.6	31.2	35.0
5	0.50	0.08	0.09	0.07	0.45	0.50	0.56	8.1	27.8	31.0	30.1	47.7	54.9
6	0.80	0.12	0.13	0.10	0.67	0.78	0.81	8.2	29.0	31.4	65.9	74.3	82.8
7	0.20	0.07	0.08	0.05	0.21	0.24	0.32	8.1	29.2	29.7	14.0	32.4	35.2
8	0.50	0.12	0.15	0.10	0.44	0.51	0.57	8.7	30.6	31.3	31.1	49.6	53.0
9	0.80	0.18	0.22	0.16	0.66	0.78	0.81	9.2	32.0	33.5	67.0	75.5	79.8
10	0.20	0.13	0.17	0.14	0.20	0.25	0.38	8.8	33.1	29.0	11.7	33.5	25.7
11	0.50	0.26	0.36	0.32	0.42	0.52	0.60	12.5	39.8	39.8	30.6	51.0	41.0
12	0.80	0.42	0.53	0.50	0.65	0.79	0.81	13.5	46.4	51.1	67.9	77.6	62.2
13	0.20	0.05	0.06	0.04	0.20	0.24	0.30	8.1	27.3	31.0	13.9	33.2	35.4
14	0.50	0.09	0.11	0.07	0.44	0.51	0.54	8.5	28.6	31.6	30.3	49.6	53.5
15	0.80	0.14	0.17	0.12	0.66	0.78	0.80	9.0	29.7	33.6	66.2	75.7	80.9
16	0.20	0.04	0.04	0.04	0.21	0.23	0.31	7.8	25.2	28.5	13.2	30.3	34.3
17	0.50	0.06	0.06	0.06	0.44	0.49	0.56	7.9	25.8	29.7	27.2	45.2	54.2
18	0.80	0.08	0.08	0.08	0.66	0.76	0.81	8.0	25.9	30.5	60.2	71.5	82.2
19	0.20	0.04	0.04	0.04	0.19	0.22	0.32	7.8	24.2	27.1	11.3	27.8	31.5
20	0.50	0.05	0.05	0.06	0.41	0.48	0.56	7.7	24.2	27.5	23.0	39.9	49.1
21	0.80	0.06	0.06	0.07	0.64	0.75	0.81	7.7	24.3	27.1	52.6	62.6	75.6
22	0.20	0.05	0.05	0.05	0.16	0.21	0.33	9.0	19.5	21.5	10.9	20.7	23.1
23	0.50	0.08	0.07	0.05	0.36	0.47	0.51	9.2	19.8	21.2	20.5	30.3	35.3
24	0.80	0.10	0.09	0.06	0.66	0.77	0.75	9.0	20.0	20.2	44.1	50.4	60.2

DGP	Initial R^2	$\alpha = 0.90$						$\alpha = 0.85$					
		CADF	CADF GLS	EJ	M CADF	M CADF GLS	M EJ	CADF	CADF GLS	EJ	M CADF	M CADF GLS	M EJ
1	0.20	18.7	48.9	62.7	30.0	57.0	68.2	34.4	62.0	83.7	48.0	69.1	84.8
2	0.50	19.5	50.3	64.1	57.3	74.9	90.1	35.2	63.1	85.2	73.4	84.1	96.3
3	0.80	20.0	51.2	65.8	89.8	92.9	99.0	36.5	64.3	86.0	95.6	96.0	99.2
4	0.20	18.9	50.7	64.0	31.0	57.4	68.0	35.3	63.3	85.4	49.0	69.5	84.8
5	0.50	20.0	52.2	65.7	58.9	75.5	90.1	36.9	65.0	86.7	74.3	84.5	96.4
6	0.80	21.1	53.3	67.4	90.4	93.0	98.9	38.7	66.5	87.7	95.9	96.2	99.1
7	0.20	20.5	54.9	64.4	31.4	59.6	67.9	38.4	67.4	86.2	50.5	71.9	85.8
8	0.50	22.9	57.6	67.5	60.5	77.6	89.0	42.7	70.3	88.1	76.2	86.3	95.9
9	0.80	25.2	60.3	70.6	90.7	93.6	98.4	46.0	73.1	90.2	96.0	96.7	98.7
10	0.20	23.2	64.9	56.1	27.6	63.8	47.6	45.7	80.3	76.5	49.0	79.2	65.2
11	0.50	35.7	73.7	70.8	62.3	81.5	71.4	62.8	86.6	86.3	79.6	91.0	83.8
12	0.80	43.9	80.8	84.4	92.9	96.0	90.5	73.7	90.9	93.5	97.4	98.4	92.9
13	0.20	20.0	51.3	65.4	31.2	59.3	68.3	37.3	64.1	86.3	49.3	71.3	85.9
14	0.50	21.6	53.6	67.4	59.3	77.0	89.4	39.8	66.5	87.6	75.0	85.3	96.6
15	0.80	23.6	56.3	70.1	90.3	93.5	98.8	43.0	69.1	89.7	95.8	96.5	99.0
16	0.20	18.5	47.6	61.5	29.5	55.1	67.4	33.6	60.2	82.3	46.6	67.6	83.7
17	0.50	18.9	48.3	62.7	54.6	72.8	89.4	34.1	61.1	83.5	71.5	82.6	95.8
18	0.80	19.5	48.8	64.2	88.1	91.9	98.8	35.2	62.0	83.9	94.7	95.5	99.1
19	0.20	17.7	44.4	56.2	25.8	50.7	60.7	31.3	57.0	74.5	41.1	62.8	75.1
20	0.50	17.8	44.7	55.5	47.2	66.0	82.6	31.6	57.1	74.2	64.5	76.9	90.7
21	0.80	17.5	44.7	54.2	82.0	86.9	97.1	31.7	57.7	73.5	91.6	92.5	98.3
22	0.20	18.6	33.3	39.7	23.5	36.3	36.2	30.8	43.2	52.0	36.9	46.1	43.8
23	0.50	19.3	34.0	38.7	41.0	49.1	59.6	31.7	44.2	51.2	56.9	59.5	69.6
24	0.80	19.4	34.5	37.6	71.0	71.5	90.3	32.5	44.6	50.5	83.1	79.0	94.5

CHAPTER 3. UNIT ROOT TESTS USING MIXED-FREQUENCY VAR
MODELS

Table 3.4.4: Size-adjusted power of unit root tests, Time trend included, $T = 100$

DGP	Initial R^2	Estimated R^2						$\alpha = 0.95$					
		CADF	CADF GLS	EJ	M CADF	M CADF GLS	M EJ	CADF	CADF GLS	EJ	M CADF	M CADF GLS	M EJ
1	0.20	0.04	0.05	0.04	0.22	0.23	0.31	7.7	11.0	11.6	11.6	15.2	15.2
2	0.50	0.07	0.07	0.06	0.46	0.49	0.56	7.8	11.3	12.1	21.9	25.0	24.6
3	0.80	0.09	0.10	0.09	0.70	0.76	0.81	7.9	11.6	12.3	48.2	49.4	36.3
4	0.20	0.05	0.06	0.04	0.23	0.23	0.31	8.0	11.4	11.7	11.9	15.4	15.7
5	0.50	0.08	0.09	0.07	0.47	0.49	0.56	8.1	11.9	12.1	23.4	26.0	24.6
6	0.80	0.11	0.12	0.10	0.71	0.77	0.81	8.2	12.1	12.3	51.1	51.9	36.5
7	0.20	0.06	0.08	0.05	0.23	0.23	0.32	8.2	12.1	11.9	11.9	16.2	15.1
8	0.50	0.12	0.15	0.11	0.48	0.50	0.57	8.7	13.0	12.4	25.2	27.9	23.4
9	0.80	0.18	0.21	0.16	0.71	0.77	0.82	8.8	13.1	13.0	54.8	54.7	31.4
10	0.20	0.12	0.17	0.14	0.23	0.24	0.38	11.7	15.8	12.4	12.0	16.0	12.2
11	0.50	0.26	0.35	0.32	0.46	0.51	0.61	15.4	19.7	18.3	26.4	30.1	18.2
12	0.80	0.41	0.52	0.51	0.70	0.78	0.82	14.4	22.2	22.6	59.4	57.4	24.3
13	0.20	0.05	0.06	0.04	0.22	0.23	0.30	7.9	11.1	12.3	11.6	15.1	15.0
14	0.50	0.09	0.11	0.07	0.47	0.50	0.55	8.2	11.6	12.0	23.2	26.1	23.5
15	0.80	0.14	0.16	0.12	0.71	0.77	0.80	8.0	11.9	12.1	51.4	52.2	34.7
16	0.20	0.04	0.04	0.04	0.22	0.23	0.31	7.6	11.3	11.6	11.2	14.3	15.2
17	0.50	0.06	0.06	0.06	0.45	0.48	0.56	7.8	11.1	12.1	20.3	23.4	23.9
18	0.80	0.07	0.08	0.08	0.68	0.75	0.81	7.7	11.0	12.6	44.5	45.7	35.7
19	0.20	0.04	0.04	0.04	0.19	0.22	0.32	7.8	10.8	10.8	9.9	12.6	14.6
20	0.50	0.05	0.05	0.06	0.40	0.47	0.56	7.5	10.5	10.9	16.6	19.7	21.8
21	0.80	0.06	0.06	0.08	0.63	0.74	0.81	7.5	10.5	10.4	37.9	37.6	31.0
22	0.20	0.05	0.05	0.05	0.15	0.21	0.34	7.6	9.3	9.7	8.5	10.7	11.7
23	0.50	0.07	0.08	0.06	0.34	0.46	0.52	7.5	9.7	9.4	14.2	16.2	18.5
24	0.80	0.09	0.10	0.06	0.61	0.75	0.76	7.3	9.4	8.6	30.0	31.2	26.7

DGP	Initial R^2	$\alpha = 0.90$						$\alpha = 0.85$					
		CADF	CADF GLS	EJ	M CADF	M CADF GLS	M EJ	CADF	CADF GLS	EJ	M CADF	M CADF GLS	M EJ
1	0.20	14.8	25.9	29.6	24.6	34.0	38.1	27.6	43.5	54.5	41.3	52.8	62.6
2	0.50	15.6	26.7	31.5	46.9	55.2	63.1	28.7	44.6	57.0	66.6	73.7	87.7
3	0.80	15.8	27.5	33.1	82.6	85.2	92.1	29.2	45.9	59.1	92.5	94.2	97.8
4	0.20	15.2	26.6	31.0	25.0	34.8	38.2	28.6	45.1	56.5	42.0	54.1	62.7
5	0.50	16.3	27.7	32.3	48.9	56.9	63.6	30.3	46.7	59.0	68.3	74.8	87.9
6	0.80	16.9	29.1	33.6	84.6	86.7	91.6	31.6	48.3	60.9	93.1	94.8	97.9
7	0.20	16.3	29.1	30.7	26.0	37.0	35.9	31.3	49.2	57.4	44.1	57.4	60.7
8	0.50	18.3	31.6	33.4	52.1	60.5	60.4	34.8	52.8	60.4	71.5	78.6	85.7
9	0.80	19.9	33.8	35.6	86.6	88.2	87.6	38.0	55.7	64.1	94.4	95.6	96.8
10	0.20	22.8	39.3	26.5	24.8	38.7	22.7	41.7	65.0	43.3	43.6	61.8	37.1
11	0.50	34.5	51.0	40.4	55.7	65.6	40.3	60.1	75.9	59.5	76.7	83.8	61.8
12	0.80	39.6	60.5	56.5	90.1	91.4	64.9	70.5	84.9	76.8	96.4	97.5	82.2
13	0.20	15.5	26.3	30.3	24.7	35.4	35.5	29.1	44.5	56.1	42.0	54.4	59.9
14	0.50	16.6	28.6	31.7	49.4	57.4	60.5	31.6	47.4	58.2	69.2	75.9	86.6
15	0.80	18.1	30.7	33.6	85.1	87.3	90.3	34.5	51.4	61.7	93.3	95.1	97.6
16	0.20	14.7	25.6	28.8	23.9	32.5	37.1	27.0	42.4	52.8	40.0	50.8	61.3
17	0.50	15.0	25.8	30.5	44.7	52.9	61.9	27.6	42.8	54.8	64.3	71.3	86.2
18	0.80	15.1	26.0	31.9	80.0	82.6	91.2	27.6	43.5	56.9	91.4	93.0	97.6
19	0.20	14.4	23.3	25.5	20.1	28.8	33.4	25.9	38.6	44.4	34.2	45.2	52.7
20	0.50	14.4	23.3	25.6	36.4	44.3	54.8	25.9	38.7	43.4	55.6	62.8	77.1
21	0.80	14.2	22.9	24.7	71.0	73.2	84.1	25.6	38.6	41.8	86.3	87.0	94.9
22	0.20	14.0	18.8	19.0	17.0	22.2	21.9	23.8	29.7	30.7	28.1	34.4	30.0
23	0.50	14.3	19.1	18.5	29.6	34.1	38.5	24.3	30.6	30.7	46.0	49.1	54.9
24	0.80	14.0	19.2	18.1	59.3	59.5	69.9	24.6	31.0	29.9	75.7	74.5	87.8




Figure 3.2: Differences in size-adjusted power between mixed- and low-frequency EJ tests for $\alpha = 0.95, 0.90$ and 0.85 , Constant only, $T = 100$

also similar to the ones collected from the original articles. The asymptotic critical values presented in the table are interpolated for the values of R^2 . Finite-sample critical values converge to their asymptotic values, in both mixed- and low-frequency tests.

Overall, the results suggest that the information loss resulting from the lower correlation penalises the power performance of the low-frequency tests. Average power gains from exploiting mixed-frequency data in the M-EJ test are quite substantial, reaching 7.9, 18.4 and 22.1 percentage points for $\alpha = 0.85, 0.90$ and 0.95 , respectively, for the case of constant only. The power gains increase as the alternative hypothesis are more demanding, being closer to the unit root, meaning that the losses in power performance are smaller for the mixed-frequency tests than for the low-frequency ones. When a time trend is included, this pattern does not occur, with gains of 22.4, 26.9 and 10.6 percentage points for $\alpha = 0.85, 0.90$ and 0.95 , respectively.

Time aggregation has a milder impact on the correlation between the aggregated data when strong autoregressive dynamics is in place. In this scenario, the

CHAPTER 3. UNIT ROOT TESTS USING MIXED-FREQUENCY VAR
MODELS

Table 3.4.5: Estimated R^2 and critical values of EJ-family unit root tests, Constant only

	Initial R^2					
	Low-frequency test			Mixed-frequency test		
	0.2	0.5	0.8	0.2	0.5	0.8
<hr/>						
$T = 100$						
Estimated R^2	0.04	0.06	0.08	0.31	0.57	0.82
Finite-sample critical values	2.99	3.03	3.07	3.24	4.39	10.08
Asymptotical critical values	3.37	3.38	3.40	3.73	4.91	10.93
<hr/>						
$T = 500$						
Estimated R^2	0.02	0.05	0.07	0.22	0.51	0.80
Finite-sample critical values	3.20	3.19	3.23	3.35	4.24	9.10
Asymptotical critical values	3.35	3.38	3.39	3.57	4.48	9.17
<hr/>						

Notes: All estimates were obtained from the DGP with the simplest dynamics, i.e., DGP 1, 2 and 3 in Table 3.4.1. The estimated R^2 corresponds to the average values over the replications. The asymptotical critical values are interpolated for the estimated R^2 .

performance of the low-frequency tests is less affected.

3.4.2 The case of a larger sample size

Let us consider now a sample with 500 observations. The corresponding alternative hypotheses are $\alpha = 0.99, 0.98$ and 0.97 . Not only is the sample larger but also the alternative hypotheses are closer to the unit root case.

In general, there are less size distortions using this larger sample. As in the case of the sample with 100 observations, when downward distortions exist, they tend to be less marked for mixed-frequency tests (Figure 3.3). Results are qualitatively similar when a time trend is also considered. The exception is the M-EJ test. Despite having less size distortions than in the sample with 100 observations, in this case the M-EJ test tends to be more oversized than the other tests, especially in the presence of large negative moving average roots. This result is likely to reflect the presence of a cumulative bias from estimating $m + 1$ trend parameters.

Power tends to be higher in the sample with 500 observations. Moreover, size-adjusted power of feasible point optimal unit root tests is higher than the power of the comparable ADF-family tests. Focusing on the former type of tests, size-adjusted power of M-EJ tests is higher than the one of EJ tests. This is true for both deterministic cases considered. The power differences between the low- and mixed-frequency EJ tests for the case of including only a constant are presented in Figure 3.4.

The M-EJ test continues to deal better with near-integration if the deterministic variables include only a constant. On average, the power gains from exploiting mixed-frequency data increase as the alternative hypothesis gets closer to the unit root case, ranging from 5.4 to 17.1 and 26.4 percentage points for $\alpha = 0.85, 0.90$ and 0.95 , respectively. In the case of also including a time trend, we get 23.8 to 31.6 and 13.6 percentage points for $\alpha = 0.85, 0.90$ and 0.95 , respectively.

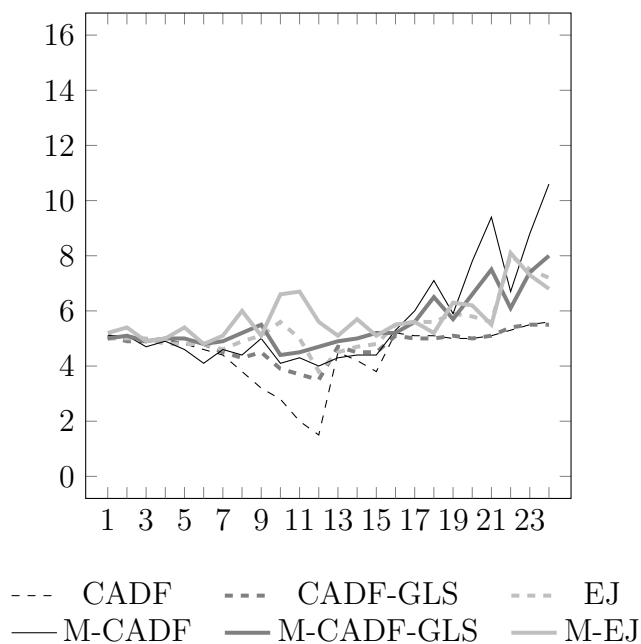


Figure 3.3: Finite sample size, $T = 500$, Constant only




Figure 3.4: Differences in size-adjusted power between mixed- and low-frequency EJ tests for $\alpha = 0.99, 0.98$ and 0.97 , Constant only, $T = 500$

3.4.3 Different lags

In this section we present a sensitivity analysis to the choice of the truncation lag. Figure 3.5 shows the differences between the size-adjusted power of mixed- and low-frequency EJ tests with an ad-hoc number of lags included in the test regression (1, 2, 3 and 4 lags, respectively). Note that the number of lags refers to the low frequency. The same number of lags is used in each replication and for each variable (the dependent variable and the covariate). These results illustrate the case of including a constant in the test regression.¹³ A positive bar means that the size-adjusted power of mixed-frequency tests is higher than the one of low-frequency tests.

Regardless of the choice of lags, mixed-frequency tests tend to outperform the low-frequency ones. This is mainly due to the fact that exploiting mixed-frequency data enables us to capture the (stronger) underlying correlation between the dependent variable and the covariate. By looking at Figure 3.5 (a) we can see that the gains from using mixed-frequency data tend to decrease with the number of lags included. This pattern can be partially explained by small sample effects, given that the differences between the lags sharply decrease in the sample with 500 observations. Nevertheless, the remaining differences are still slightly decreasing in the number of lags, which may be related with the fact that VAR models of the mixed-frequency approach can be less parsimonious than traditional low-frequency VAR models, especially with a high number of lags and a high m .

Overall, even in case of high uncertainty about the choice of the truncation lag, as is typically the case in empirical applications, using the mixed-frequency framework may contribute to reduce the impact of potential misspecification in the power of the unit root tests.

¹³Results for other alternative hypotheses and including a time trend are not qualitatively different and are available from the author upon request.

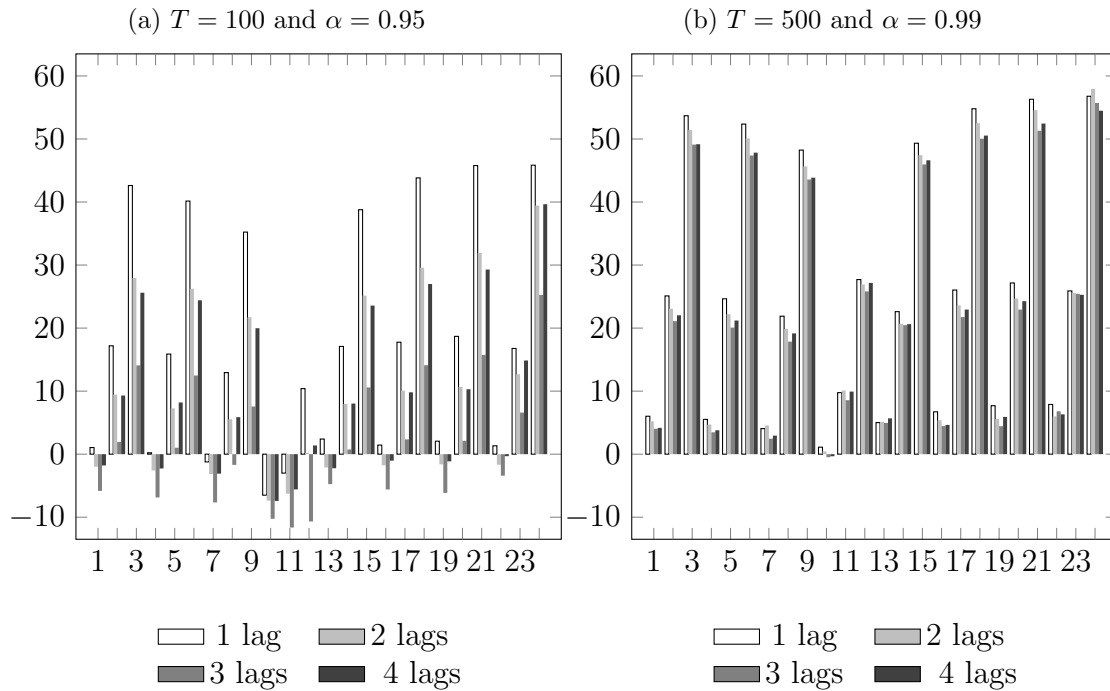


Figure 3.5: Differences in size-adjusted power between mixed- and low-frequency tests with GLS demeaning by number of lags included in the test regression, Constant only

3.5 An application to the US unemployment rate

The following exercise illustrates the use of the above-mentioned unit root tests, both the low- and mixed-frequency versions. We apply the EJ and M-EJ tests to assess whether the US unemployment rate has a unit root.¹⁴ We exploit insights provided by the correlation between the variable of interest and a stationary covariate that is the continued jobless claims. We use the same data as in Chapter 2 or Duarte (2014), to allow the comparison with the ADF-family of unit root tests.

The overall unemployment rate (in logs) is a monthly series taken from the US Bureau of Labor Statistics and is seasonally adjusted. The stationary covariate is

¹⁴The difference between the unemployment rate and the natural rate of unemployment is also often analysed. In this exercise we focused on the level of the unemployment rate. Estimating a natural rate of unemployment is beyond the scope of this article. Simply using an estimate of the natural rate of unemployment collected elsewhere (e.g., the Congressional Budget Office estimates a quarterly natural rate of unemployment, which is made available by the St. Louis Federal Reserve Bank) would lead to bias on hypothesis testing, as shown by Murphy and Topel (1985).

CHAPTER 3. UNIT ROOT TESTS USING MIXED-FREQUENCY VAR MODELS

the continued jobless claims, also expressed in logs, which is a weekly series. The continued jobless claims are released by the US Department of Labor and are also seasonally adjusted. The stationarity of the covariate was confirmed by univariate unit root tests, namely ADF and ADF-GLS tests.¹⁵ The data cover the period from January 1980 to June 2014. For the low-frequency version of the covariate we use the beginning of the period value. Figure 3.6 shows the series used, in quarterly frequency (the common frequency) and in logs.

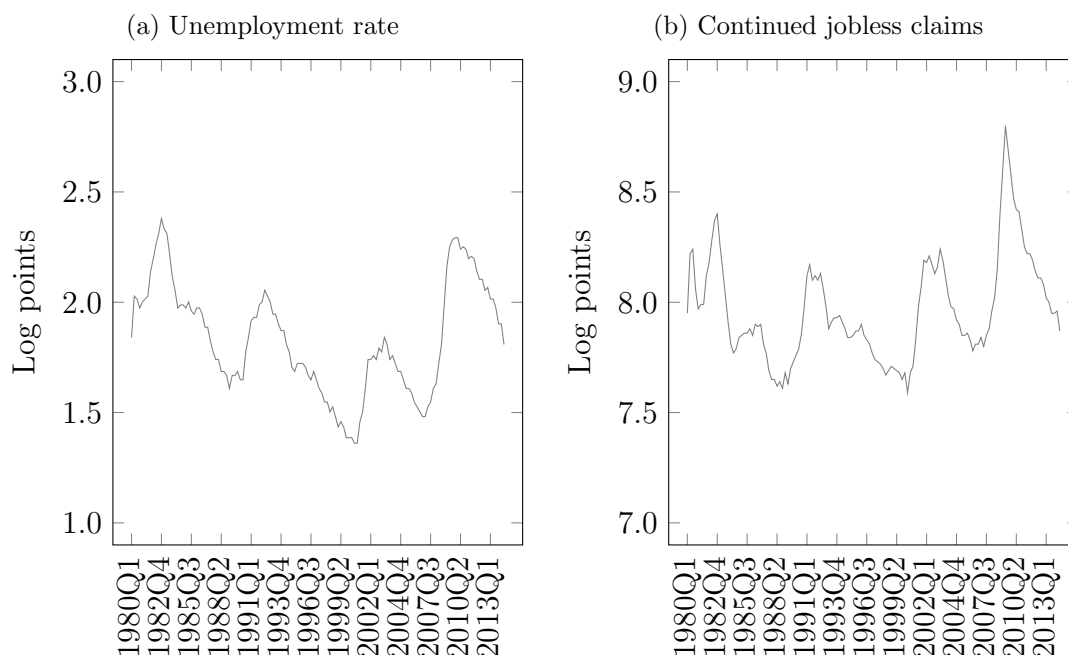


Figure 3.6: Data on US unemployment

The truncation lags were chosen using the AIC (for more details, see Section 3.4). In order to have a fixed m , we considered the first four weeks in each month, i.e., $m = 4$ for each low-frequency observation. Table 3.5.1 presents the results, providing the test statistics and the estimated R^2 for the covariate-augmented tests. We consider two different samples, one from 1980 Q1 to 2014 Q2 (top panel in Table 3.5.1) and another excluding the great recession period, from 1980 Q1 to 2006 Q4 (bottom panel in Table 3.5.1). In addition, Figure 3.6 highlights the relevance of including a time trend in the test regression, namely in the shorter

¹⁵The tests were performed for weekly, monthly and quarterly frequencies and, as expected, the result was always the same.

CHAPTER 3. UNIT ROOT TESTS USING MIXED-FREQUENCY VAR
MODELS

sample.¹⁶ Therefore, for each sample, Table 3.5.1 shows results for only including a constant and for also including a time trend.

Table 3.5.1: Unit root tests for US monthly unemployment rate, using jobless claims as covariate

Sample: 1980Q1 to 2014Q2			
Constant only			
	ERS	EJ	M-EJ
t -statistic	8.854	-1.131*	1.521*
R^2	-	0.81	0.88
Time trend included			
	ERS	EJ	M-EJ
t -statistic	32.699	-5.163*	-0.186*
R^2	-	0.82	0.86
Sample: 1980Q1 to 2006Q2			
Constant only			
	ERS	EJ	M-EJ
t -statistic	16.118	15.197	21.177
R^2	-	0.77	0.85
Time trend included			
	ERS	EJ	M-EJ
t -statistic	32.420	3.180*	5.296*
R^2	-	0.76	0.84

Note: For the low-frequency covariate-augmented unit root tests the frequency of the covariate equals the frequency of the dependent variable. For the mixed-frequency tests, the covariate has a weekly frequency. * significant at a 5 per cent asymptotic level. ** significant at a 1 per cent asymptotic level. For the covariate-augmented GLS family of tests, Pesavento (2006) only presents 5 per cent asymptotic significance levels. The same happens with the EJ test in Elliott and Jansson (2003).

Starting with the shorter sample, all tests which consider only a constant in the test regression point in the same direction, not rejecting the null hypothesis,

¹⁶There is no consensus in the literature about including or not a trend when modelling unemployment, existing an ongoing discussion based on sample-driven and theoretical arguments.

suggesting that the series is $I(1)$. The conclusion from the univariate ERS test remains unchanged when a time trend is included. Low- and mixed-frequency covariate-augmented tests, however, reject the null hypothesis, suggesting that the unemployment rate is trend stationary in that sample.

Using the longer sample, the ERS test continues to suggest that the level of US unemployment is not stationary. In contrast, all covariate-augmented tests reject the null hypothesis, whether or not a time trend is included. Hence, the level of unemployment rate seems to be stationary. Furthermore, the estimates of R^2 are always higher when the mixed-frequency approach is used. The results obtained in this empirical illustration are fairly similar to the ones presented in Chapter 2 or Duarte (2014), for the ADF-family of unit root tests.

3.6 Conclusion

This article contributes to the literature by considering high-frequency covariates in feasible point optimal covariate-augmented unit root tests. Time aggregation does not affect the long-run properties of time series, but can have severe effects on the correlation between variables.

Previous works already showed that taking into account correlated stationary covariates in unit root test regressions contributes to improve the power performance of these tests. This is true for feasible point optimal covariate-augmented unit root tests, as proposed by Elliott and Jansson (2003) — EJ tests — and for the covariate-augmented ADF-family of tests, proposed by Hansen (1995) and Pesavento (2006).

To deal with the mixed-frequency data we use the MI(xed) DA(ta) S(ampling) framework. Given that the EJ tests require estimating VAR models, we propose an unconstrained, though parsimonious, stacked skip-sampled reduced-form VAR-MIDAS model, which can be estimated using standard econometric techniques.

Our simulation exercise shows that mixed-frequency data in general improves the power performance of the EJ tests. Moreover, the gains are robust to the sample size and to lag selection. Mixed-frequency EJ tests (M-EJ) also tend to have better power performance than mixed-frequency CADF tests (both with OLS or GLS demeaning/detrending) proposed by Duarte (2014) (in chapter 2), although with higher size distortions (a similar result to the one presented in the literature for the traditional single frequency tests).

We take the M-EJ test to the data by applying it to the US unemployment rate. Results suggest that the unemployment rate is stationary.

Acknowledgments

I would like to thank Paulo Júlio, Carlos Robalo Marques, João Nicolau and Paulo Rodrigues. The usual disclaimers apply.

Bibliography

- ANDREOU, E., E. GHYSELS AND A. KOURTELLOS, *The Oxford Handbook of Economic Forecasting*, chapter 8, Forecasting with mixed-frequency data (Oxford University Press, 2011), 225–267.
- ANDREWS, D. W. K., “Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation,” *Econometrica* 59 (May 1991), 817–58.
- BERK, K. N., “Consistent Autoregressive Spectral Estimates,” *The Annals of Statistics* 2 (May 1974), 489–502.
- CHIU, C. W. J., B. ERAKER, A. T. FOERSTER, T. B. KIM AND H. D. SEOANE, “Estimating VAR’s sampled at mixed or irregular spaced frequencies: a Bayesian approach,” Research Working Paper RWP 11-11, Federal Reserve Bank of Kansas City, 2011.
- DICKEY, D. A. AND W. A. FULLER, “Distribution of the Estimators for Autoregressive Time Series With a Unit Root,” *Journal of the American Statistical Association* 74 (1979), 427–431.
- DUARTE, C., “Covariate-augmented unit root tests with mixed-frequency data,” Working Papers 7–15, Banco de Portugal, June 2014.
- ELLIOTT, G. AND M. JANSSON, “Testing for unit roots with stationary covariates,” *Journal of Econometrics* 115 (2003), 75–89.
- ELLIOTT, G., M. JANSSON AND E. PESAVENTO, “Optimal Power for Testing Potential Cointegrating Vectors With Known Parameters for Nonstationarity,” *Journal of Business and Economic Statistics* 23 (2005), 34–48.
- ELLIOTT, G., T. J. ROTHENBERG AND J. H. STOCK, “Efficient Tests for an Autoregressive Unit Root,” *Econometrica* 64 (July 1996), 813–36.
- FORONI, C., E. GHYSELS AND M. MARCELLINO, “Mixed-frequency vector autoregressive models,” *Advances in Econometrics* 32 (2013), 247–272.

BIBLIOGRAPHY

- FORONI, C. AND M. MARCELLINO, “A Comparison of Mixed Frequency Approaches for Modelling Euro Area Macroeconomic Variables,” Economics Working Papers ECO2012/07, European University Institute, 2012.
- FORONI, C., M. MARCELLINO AND C. SCHUMACHER, “U-MIDAS: MIDAS regressions with unrestricted lag polynomials,” Discussion Paper Series 1: Economic Studies 2011,35, Deutsche Bundesbank, Research Centre, 2011.
- FOSSATI, S., “Covariate unit root tests with good size and power,” *Computational Statistics & Data Analysis* 56 (2012), 3070–3079.
- GALVAO JR., A. F., “Unit root quantile autoregression testing using covariates,” *Journal of Econometrics* 152 (October 2009), 165–178.
- GALVÃO, A. B., “Changes in predictive ability with mixed frequency data,” *International Journal of Forecasting* 29 (2013), 395–410.
- GHYSELS, E., “Macroeconomics and the Reality of Mixed Frequency Data,” mimeo, University of North Carolina (UNC) at Chapel Hill - Department of Economics, May 2012.
- GHYSELS, E. AND J. I. MILLER, “Testing for Cointegration with Temporally Aggregated and Mixed-frequency Time Series,” Working Papers 1307, Department of Economics, University of Missouri, June 2013.
- GHYSELS, E., A. SINKO AND R. VALKANOV, “MIDAS regressions: Further results and new directions,” *Econometric Reviews* 26 (2007), 53–90.
- GRANGER, C. AND P. NEWBOLD, “SPURIOUS REGRESSIONS IN ECONOMETRICS,” *Journal of Econometrics* 2 (1974), 111–120.
- GRANGER, C. AND P. SIKLOS, “Systematic sampling, temporal aggregation, seasonal adjustment, and cointegration theory and evidence,” *Journal of Econometrics* 66 (1995), 357–369.
- GÖTZ, T. AND A. HECQ, “Testing for Granger causality in large mixed-frequency VARs,” Research Memorandum, Maastricht University, Graduate School of Business and Economics (GSBE) 028, Maastricht University, Graduate School of Business and Economics (GSBE), July 2014a.
- GÖTZ, T. B. AND A. HECQ, “Nowcasting causality in mixed frequency vector autoregressive models,” *Economics Letters* 122 (2014b), 74–78.

- HALDRUP, N. AND M. JANSSON, “Improving Size and Power in Unit Root Testing,” in T. C. Mills and K. Patterson, eds., *Palgrave Handbook of Econometrics* volume 2: Applied Econometrics, chapter 7 (Palgrave Macmillan, 2006), 252–277.
- HAMILTON, J., *Time Series Analysis* (Princeton University Press, 1994).
- HANSEN, B. E., “Rethinking the Univariate Approach to Unit Root Testing: Using Covariates to Increase Power,” *Econometric Theory* 11 (October 1995), 1148–1171.
- HANSEN, L. AND T. J. SARGENT, *Recursive Models of Dynamic Linear Economies* (Princeton University Press, 2013).
- HAUG, A. A., “Temporal Aggregation and the Power of Cointegration Tests: A Monte Carlo Study,” *Oxford Bulletin of Economics and Statistics* 64 (September 2002), 399–412.
- HECQ, A., J. URBAIN AND T. GÖTZ, “Testing for common cycles in non-stationary VARs with varied frequency data,” Research Memorandum 002, Maastricht University, Graduate School of Business and Economics (GSBE), 2013.
- JANSSON, M., “Consistent Covariance Matrix Estimation For Linear Processes,” *Econometric Theory* 18 (December 2002), 1449–1459.
- , “Stationarity Testing With Covariates,” *Econometric Theory* 20 (February 2004), 56–94.
- JUHL, T. AND Z. XIAO, “Power functions and envelopes for unit root tests,” *Econometric Theory* 19 (4 2003), 240–253.
- KUZIN, V., M. MARCELLINO AND C. SCHUMACHER, “MIDAS versus mixed-frequency VAR: Nowcasting GDP in the euro area,” *International Journal of Forecasting* 27 (2011), 529–542.
- MARCELLINO, M., “Consequences of temporal aggregation in empirical analysis,” *Journal of Business and Economic Statistics* 17 (January 1999), 129–136.
- MARCELLINO, M. AND C. SCHUMACHER, “Factor MIDAS for Nowcasting and Forecasting with Ragged-Edge Data: A Model Comparison for German GDP,” *Oxford Bulletin of Economics and Statistics* 72 (08 2010), 518–550.

BIBLIOGRAPHY

- MARIANO, R. S. AND Y. MURASAWA, “A new coincident index of business cycles based on monthly and quarterly series,” *Journal of Applied Econometrics* 18 (2003), 427–443.
- MCCRACKEN, M., M. OWUANG AND T. SEKHPOSYAN, “Real-time forecasting with a large, mixed frequency bayesian VAR,” mimeo, February 2014.
- MIKOSCH, H. AND S. NEUWIRTH, “Real-time forecasting with MIDAS VAR,” BOFIT Discussion Papers 13, Bank of Finland, BOFIT Institute for Economies in Transition, April 2015.
- MÜLLER, U. K. AND G. ELLIOTT, “Tests for Unit Roots and the Initial Condition,” *Econometrica* 71 (07 2003), 1269–1286.
- MURPHY, K. AND R. TOPEL, “Estimation and Inference in Two-Step Econometric Models,” *Journal of Business and Economic Statistics* 3 (October 1985), 370–379.
- NUNES, L. C., “Nowcasting quarterly GDP growth in a monthly coincident indicator model,” *Journal of Forecasting* 24 (2005), 575–592.
- PESAVENTO, E., “Near-Optimal Unit Root Tests with Stationary Covariates with Better Finite Sample Size,” Economics Working Papers ECO2006/18, European University Institute, 2006.
- PHILLIPS, P. C. B., “Towards a unified asymptotic theory for autoregression,” *Biometrika* 74 (1987), 535–547.
- PIERSE, R. G. AND A. J. SNELL, “Temporal aggregation and the power of tests for a unit root,” *Journal of Econometrics* 65 (February 1995), 333–345.
- QIAN, H., “Vector autoregression with varied frequency data,” MPRA Paper 34682, Munich Personal RePEc Archive, 2010.
- QU, Z. AND P. PERRON, “A Modified Information Criterion For Cointegration Tests Based On A Var Approximation,” *Econometric Theory* 23 (August 2007), 638–685.
- RODRIGUES, P. M., “Properties of recursive trend-adjusted unit root tests,” *Economics Letters* 91 (June 2006), 413–419.
- SCHWERT, G. W., “Tests for Unit Roots: A Monte Carlo Investigation,” *Journal of Business and Economic Statistics* 7 (1989), 147–159.

BIBLIOGRAPHY

- SILVESTRINI, A. AND D. VEREDAS, “Temporal Aggregation Of Univariate And Multivariate Time Series Models: A Survey,” *Journal of Economic Surveys* 22 (07 2008), 458–497.
- STOCK, J. H., *Unit roots, structural breaks and trends*, volume 4 of *Handbook of Econometrics*, chapter 46 (Elsevier, 1986), 2739–2841.
- SUL, D., P. C. B. PHILLIPS AND C.-Y. CHOI, “Prewhitening Bias in HAC Estimation,” *Oxford Bulletin of Economics and Statistics* 67 (08 2005), 517–546.
- TAYLOR, A. M. R., “Regression-Based Unit Root Tests with Recursive Mean Adjustment for Seasonal and Nonseasonal Time Series,” *Journal of Business and Economic Statistics* 20 (2002), 269–281.
- WORKING, H., “Note on the Correlation of First Differences of Averages in a Random Chain,” *Econometrica* 28 (1960), pp. 916–918.
- ZADROZNY, P., “Gaussian likelihood of continuous-time ARMAX models when data are stocks and flows at different frequencies,” *Econometric Theory* 4 (April 1988), 108–124.

Appendices

Table C.1: Asymptotic critical values for the EJ test

R^2	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Cases 1, 2	3.34	3.41	3.54	3.76	4.15	4.79	5.88	7.84	12.12	25.69
Case 3	3.34	3.41	3.54	3.70	3.96	4.41	5.12	6.370	9.17	17.99
Case 4	5.70	5.79	5.98	6.38	6.99	7.97	9.63	12.60	19.03	39.62
Case 5	5.70	5.77	6.00	6.40	7.07	8.15	10.00	13.36	20.35	41.87

Note: Taken from Elliott and Jansson (2003). The critical values were computed using 1500 steps as approximations to the Brownian motion terms in the limit theorem representations and 60 000 replications. The critical values reported are for tests of size 5%, with \bar{c} equal to -7 for cases 1, 2 and 3 and to -13.5 for cases 4 and 5.

Conclusion

Techniques that allow to exploit timely releases of high-frequency data play a key role in many areas of economic analysis, research and policymaking. Having started in the financial field, the MIDAS approach has been gaining a widespread attention in other fields.

In the first chapter, I use MIDAS regressions for forecasting, considering a wide set of these regressions. Starting from the discussions existing in the literature, I clarify the issue on how to deal with the autoregressive augmentation of the regressions. The results suggest that the MIDAS approach contributes to increase forecast accuracy. The benefits from this data-driven, and potentially more parsimonious, weighting scheme are higher for forecast horizons up to 1 quarter ahead. Moreover, MIDAS is able to exploit the information content in high-frequency series, regardless of the exact time frequency.

In the second and third chapters, I propose new unit root tests that exploit mixed-frequency information. In particular, I merge two strands of the literature — mixed-frequency techniques, such as the MIDAS approach, and covariate-augmented unit root tests. Taking as main premise that temporal aggregation of the variable of interest is unavoidable, the new tests presented in this thesis clearly have better power performance than the CADF and EJ unit root tests (Chapters 2 and 3, respectively). For the EJ test, I also develop an unconstrained, though parsimonious, stacked skip-sampled reduced-form VAR-MIDAS model, which can be estimated using standard econometric techniques.

The results of a simulation exercise show that taking on board mixed-frequency data contributes to improve the power performance of covariate-augmented unit root test. As an empirical illustration, I apply the covariate-augmented unit root tests — both low- and mixed-frequency — to the US unemployment rate.

Throughout this thesis I always assume that the data do not have seasonal features. Moreover, I also do not allow for structural breaks in the simulation exercises. Dealing explicitly with seasonality and structural breaks are two subjects that are out of the scope of the present thesis, but are interesting avenues for future research.