International Journal of Forecasting 33 (2017) 121–131

Contents lists available at ScienceDirect

International Journal of Forecasting

journal homepage: www.elsevier.com/locate/ijforecast

Forecasting the Brazilian yield curve using forward-looking variables

Fausto Vieira[a,*], Marcelo Fernandes[b,a], Fernando Chague[c]

[a] Sao Paulo School of Economics, Fundação Getulio Vargas (FGV), Brazil
[b] Queen Mary University of London, United Kingdom
[c] Department of Economics, University of São Paulo, Brazil

ARTICLE INFO

Keywords:
Bonds
Factor-augmented VAR
Forecasting
Term structure
Yield curve

ABSTRACT

This paper proposes a forecasting model that combines a factor augmented VAR (FAVAR) methodology with the Nelson and Siegel (NS) parametrization of the yield curve in order to predict the Brazilian term structure of interest rates. Importantly, we extract the principal components for the FAVAR from a large data set containing a range of forward-looking macroeconomic and financial variables. Our forecasting model improves on the predictive accuracy of extant models in the literature significantly, particularly at short-term horizons. For instance, the mean absolute forecast errors are 15–40% lower than those of the random walk benchmark on predictions at the three-month horizon. The out-of-sample analysis shows that the inclusion of forward-looking indicators is the key to improving the predictive ability of the model.

1. Introduction

The yield curve of treasury bonds plays a central role in both pricing financial assets and shaping market expectations. As such, accurate forecasts of the yield curve are of great importance for the Treasury, central bankers and market participants in general. Unfortunately, none of the extant models in the literature are able to outperform the random walk benchmark consistently at short horizons, while at the same time providing good forecasts at longer horizons.

This paper proposes a forecasting strategy for the yield curve that achieves this. We provide out-of-sample evidence that our forecasting model improves on the random walk benchmark at short horizons (as early as one month ahead) while at the same time providing more accurate forecasts than the extant models at longer horizons. The key ingredient of our strategy is its reliance on a comprehensive data set of macroeconomic and financial variables that are mostly forward-looking. Specifically, we proceed in three steps. In the first, we estimate the entire yield curve using the Nelson and Siegel (1987, NS) parametrization of the yield curve. The NS parametrization successfully summarizes the variation in the yield curve using the level, slope and curvature factors. In the second stage, we predict the future paths of these factors by estimating a factor augmented VAR (FAVAR) model using a comprehensive dataset of macroeconomic and financial variables. Finally, we form forecasts of the yield curve for each maturity at different horizons, using the predicted evolution of the level, slope and curvature factors.

We ensure real time forecasts by repeating these three steps at each prediction point. As our forecasting model combines a Nelson–Siegel decomposition of the yield curve with a FAVAR specification, we denote it by NS-FAVAR. Our forecasts of the yield curve beat the random walk

* Corresponding author.
E-mail address: ytcfausto@yahoo.com.br (F. Vieira).

http://dx.doi.org/10.1016/j.ijforecast.2016.08.001
0169-2070/

benchmark as early as at the one-month horizon. This represents a significant improvement, given that the available models produce meaningful predictions only from the six-month horizon (see Diebold & Li, 2006; and Moench, 2008). At the one-month horizon, our model forecast errors are 5% lower than those of the random walk benchmark, whereas at longer horizons, they are 20%–40% lower than those of the random walk benchmark.

The use of a comprehensive dataset that contains a wide array of forward-looking macroeconomic and financial variables is critical to the superior short-horizon performance of our forecasting strategy. In this respect, Brazilian economic data sets provide a surprisingly rich array of variables. As a consequence of a high-inflationary past, Brazilian market participants consume a variety of price indexes and price expectations indexes, some of which are available at the weekly or even daily frequencies. Moreover, a large number of macroeconomic and financial expectations time series are available readily, in an effort to increase the transparency of the monetary policy and guide market expectations. Our dataset contains 142 macroeconomic and financial variables at the weekly frequency, of which 40% are forward-looking indicators. Examples of macroeconomic forward-looking variables that make an important contribution to our forecasts are the market expectations of GDP growth, of the federal government balance sheet and of the debt-to-GDP ratio.

Our forecasting strategy builds on the work of Diebold, Rudebusch, and Aruoba (2006). However, instead of including only a few macroeconomic variables, we use a comprehensive data set of 103 macroeconomic variables and 39 financial indicators. We deal with this large number of conditioning variables by implementing Bernanke, Boivin, and Eliasz (2005) FAVAR econometric model. The FAVAR model restricts attention to the dynamics of a few principal components that summarize the variation in the data set. We show that conditioning on a broader information set, with many forward-looking macro-financial indicators, is key to improving the predictability.

This is not the first paper to improve yield curve forecasts at shorter horizons. de Pooter, Ravazzolo, and van Dijk (2010) study models with and without arbitrage restrictions that use macroeconomic information. They find that autoregressive models with macroeconomic predictors achieve superior performances at shorter horizons, but fail to improve on the random walk benchmark at longer horizons. Exterkate, Van Dijk, Heij, Patrick, and Groenen (2013) discuss the importance of large data sets for improving yield curve forecasts at short horizons. The authors show that factor augmented Nelson and Siegel models are able to improve the short-term forecasting during volatile periods, but cannot improve on simpler models in periods of low volatility.

In addition, we are also not the first to advocate for the use of forward-looking variables. Altavilla, Giacomini, and Constantini (2014) and Altavilla, Giacomini, and Ragusa (2014) use market and survey expectations to produce lower short-term forecasting errors of the short-term yields at the three- and six-month horizons. However, they are unable to either improve forecasts at longer horizons or ameliorate longer-term yield predictions. In contrast, van Dijk, Koopman, van der Wel, and Wright (2014) improve the forecasting performance for long maturities and at longer horizons by allowing shifting endpoints in the yield curve factors, but their forecasts are weak at shorter horizons.

In summary, we contribute by ameliorating term structure forecasts for virtually every maturity, even at short horizons. We argue that the key is to condition on the information set spanned by a few principal components from a wide array of mainly forward-looking macro-financial indicators.

We organize the remainder of this paper as follows. Section 2 reviews the Nelson–Siegel approach to the modelling of the term structure of interest rates, and describes our forecasting strategy. Section 3 describes the data set, while Section 4 discusses the out-of-sample results of our forecasting strategy. Section 5 contains several robustness exercises, and Section 6 offers some concluding remarks.

2. The forecasting strategy

Our forecasting strategy has three steps. In the first, we estimate the entire yield curve using the Nelson and Siegel (NS) parametrization of the yield curve. In the second, we predict the evolution of the level, slope and curvature factors using a FAVAR approach. Finally, we back out yield forecasts for each maturity at different horizons using the predicted future path of the NS factors.

As such, our forecasting strategy is very similar to that of Diebold et al.'s (2006) VAR model for the level, slope and curvature factors, with the main difference being that they employ only a few macroeconomic variables, whilst we condition on a much broader information set. We do so by following Stock and Watson's (2002b) idea of conditioning on a small number of principal components from a wide array of macroeconomic and financial variables. In particular, we employ a FAVAR model for the level, slope, curvature factors of the yield curve, and for the principal components from a data set of 142 macroeconomic and financial variables.

The Nelson–Siegel decomposition of the yield curve posits that we may approximate the yield with maturity n by

$$\hat{y}_t^{(n)} = \hat{\beta}_{1t} + \hat{\beta}_{2t}\left(\frac{1-e^{-\lambda n}}{\lambda n}\right) + \hat{\beta}_{3t}\left(\frac{1-e^{-\lambda n}}{\lambda n} - e^{-\lambda n}\right), \tag{1}$$

where the betas may vary over time, capturing changes in the level, slope and curvature of the term structure, respectively. The NS decomposition allows one to form predictions of the entire yield curve by simply predicting the dynamics of the level, slope and curvature factors. Following Stock and Watson (2002b), we extract the principal components of a comprehensive data set of 142 macroeconomic and financial predictors at the weekly frequency, to proxy for the broad economic conditions.[1]

[1] Although the principal component analysis formally requires independent and identically distributed observations, Doz, Giannone, and Reichlin (2012) and Stock and Watson (2002a) show that its performance is similar to that of full maximum likelihood estimation for a large panel in the context of static and dynamic factor models, respectively.

We write a FAVAR model as follows. Denote the Nelson–Siegel factors as $Z_t = (\beta_{1t}, \beta_{2t}, \beta_{3t})'$ and the $(k \times 1)$-vector F_t of the principal components, augmented with the SELIC interest rate (the target interest rate of the Brazilian Central Bank). Also, let c denote a $(k+3) \times 1$ vector of constants, $\Phi(L)$ a $(k+3) \times (k+3)$ first-order autoregressive matrix, and ω_t a vector of reduced form shocks. The FAVAR then reads

$$\begin{pmatrix} F_t \\ Z_t \end{pmatrix} = c + \Phi(L) \begin{pmatrix} F_t \\ Z_t \end{pmatrix} + \omega_t. \quad (2)$$

Bernanke et al. (2005) propose two ways of estimating a FAVAR model: (i) two-step estimation (principal components plus VAR estimation), or (ii) a Bayesian method based on Gibbs sampling. They show that the two methods produce similar results, though the two-step estimation is computationally simpler and yields results that are more plausible. Accordingly, we estimate the FAVAR model using the two-step estimation procedure. We first extract the level, slope and curvature factors of the yield curve, as well as the k principal components from our large data set of conditioning variables. We then estimate the coefficients in Eq. (2) in order to form predictions of the evolution of the NS factors as follows:

$$\hat{\beta}_{i,t} = \hat{c}_i + \sum_{j=1}^{3} \hat{\varphi}_{i,k+j} \hat{\beta}_{j,t-1} + \sum_{j=1}^{k} \hat{\varphi}_{i,j} F_{j,t-1}. \quad (3)$$

Finally, we compute the maximum likelihood forecasts of the yield curve h months ahead, given the future values of the level, slope and curvature factors $(\beta_{1,t+h}, \beta_{2,t+h}, \beta_{3,t+h})$, using only information from up to time t:

$$\hat{y}^{(n)}_{t+h|t} = \hat{\beta}_{1,t+h|t} + \hat{\beta}_{2,t+h|t} \left(\frac{1 - e^{-\lambda n}}{\lambda n} \right) + \hat{\beta}_{3,t+h|t} \left(\frac{1 - e^{-\lambda n}}{\lambda n} - e^{-\lambda n} \right). \quad (4)$$

More specifically, we consider forecasting horizons of $h =$ 1, 3, 6, 9 and 12 months, which translate into 4, 13, 26, 39 and 52 weeks, respectively, at our frequency of analysis.

3. Data set

Brazilian economic data sets are relatively comprehensive when it comes to inflation measures and market expectations. As a consequence of a high-inflationary past, Brazilian market participants make use of a variety of price indexes and price expectations indexes, some of which are available at the weekly or even daily frequencies. Moreover, a large number of macroeconomic and financial expectations time series are available readily, in an effort to increase the transparency of the monetary policy and guide market expectations. Market expectations are monitored via the weekly release of the Focus report, which includes market forecasts of daily indicators of activity, inflation and external and fiscal accounts, for the current month or year through to projections for five years ahead. This set of high frequency indicators contains relevant forward-looking information about the Brazilian economy.

Overall, this means that the Brazilian data provide lots of useful high-frequency information about future movements in the yield curve. In particular, we focus on a data set with 142 weekly indicators from the first week of March 2007 to the last week of December 2014. We consider data only from March 2007 because, while the Brazilian Treasury started to issue longer-term bonds at the end of 2006, liquidity did not pick up until 2007. As a robustness check, we also run a similar forecasting exercise for a longer sample starting in 2002, but restricting our attention to shorter-term interest rates.

We consider a multitude of data sources. Real activity is the largest group, comprising 27% of the database. These variables are mainly from the Central Bank of Brazil, except for a few daily activity indicators (e.g., electric energy consumption and credit variables). All of the indicators released by the Central Bank of Brazil (namely GDP, GDP services, industrial production and external accounts) concern market expectations over a certain horizon, e.g., the current month, next year, or next five years. All of the expectations data come from the weekly Focus report that the Central Bank of Brazil releases every Monday. In addition to mean and median forecasts, the Focus database also includes information about the standard deviations of the short- and medium-term forecasts of the inflation, activity, fiscal and balance payments series. This amount to the second largest group of data, with a share of 23% of the overall database.

Inflation-related variables computed from commodity, producer and consumer price indices constitute 20% of the database. They relate to price changes in the last month, as well as expected variation in the current month or some predetermined period (e.g. next 12 months or in five years). The producer prices are from CEASA, a distribution center for crops, fruits and vegetables, and other cooperatives. We gather commodity prices from Bloomberg, whereas we collect consumer prices at the weekly frequency from FIPE (São Paulo only). Fiscal series make up 6% of the database, collecting indicators from the Focus report such as the net sovereign debt, primary and nominal budget balance. Altogether, 56% of the inflation, real activity and fiscal time series that we consider are forward-looking indicators, thereby providing more timely information about the Brazilian outlook in both the short and long term.

Finally, we extract financial and risk indicators from Bloomberg. They correspond to 15% and 9% of the database, respectively, and include real-time indicators of the Brazilian economy, such as the five-year Brazil CDS, the local stock market index, and the currency contracts outstanding, as well as of the global economy, such as the US financial index, Latin America EMBI and the fed funds rate.

When extracting principal components from this broad range of variables, we first ensure that every time series is stationary by taking first differences, if necessary. We construct diffusion indices in two different manners. First, we extract the first two principal components of the full set of indicators in the database, as per Exterkate et al. (2013). Table 1 displays the variables that have the highest correlations with each principal component. The

Table 1
Principal components from the panel of 142 macro-financial indicators, full sample.

Principal components analysis	Correlation
Factor 1	
Latin America EMBI	0.901
Fed Funds rate	0.826
2-year treasury rate	0.813
3-month Libor	0.721
Expected trade balance annual change for the next 12 months	0.762
Factor 2	
Standard deviation of the 12-month industrial production forecast	−0.695
Standard deviation of the 12-month GDP growth forecast	−0.647
5-year US breakeven	0.659
Electric energy consumption: annual change	0.530
Expected consumer price inflation for the next month	0.501

This table reports the variables with the highest correlations with each of the principal components extracted from the panel of 142 macroeconomic and financial indicators.

Table 2
Principal components from the forward-looking indicators, full sample.

	Correlation
Factor 1	
Latin America EMBI	0.946
Import growth for the next 12 months	−0.709
5-year US breakeven	−0.700
Trade balance annual change for the next 12 months	−0.630
3 months risk reversal USD/BRL	0.600
Factor 2	
Standard deviation of the 12-month industrial production forecast	0.727
Standard deviation of the 3- to 5-year primary budget balance forecast	0.707
Standard deviation of the 12-month service-sector GDP growth forecast	0.698
Standard deviation of the 12-month government debt forecast	0.686
Service-sector GDP growth in 3–5 years	0.631

This table lists the variables with the highest correlations with each of the first two principal components of the panel of forward-looking macroeconomic and financial variables.

first component explains more than 33% of the overall variation, and correlates mostly with the Emerging Market Bond Index for Latin America, the US yield curve and the Brazilian external account. The second relates chiefly to the uncertainty of forecasting variables and inflation indicators.

As an alternative, we also consider extracting principal components only from forward-looking indicators. Table 2 reveals that the first component explains almost 40% of the forward-looking subset. It is correlated mostly with external indicators, such as the external sector (import growth and the trade balance annual change) and asset pricing (bonds and Brazilian Real risk reversal for three months[2]). The second principal component, as in the database as a whole, relates mostly to economy forecasting uncertainty, in that it involves mainly the standard deviation of analysts' forecasts.

Unlike the US Treasury emissions, the Brazilian Treasury issues bonds with specific expiration dates. For example, in January 2016, the Treasury issued a fixed rate bond with a maturity of 11 years, expiring in January 2027 (NTN-F 27). This feature of the Brazilian term structure of government bonds makes the Nelson–Siegel decomposition particularly interesting, as it allows us to back out a fixed maturity yield curve. We estimate weekly level, slope and curvature factors given by the betas in Eq. (1), but keep λ constant. We fix the value of λ at which the mean absolute difference between the actual and estimated yields is smallest for the training period between 2007 and 2011. This yields a much higher value for λ, at 0.195, than Diebold and Li's (2006) chosen value of 0.0609 for the term structure in the US.

4. Empirical analysis

4.1. Preliminary results

Bernanke and Boivin (2003) showed that central bankers benefit from considering a wide range of data when making decisions about interest rates. They concluded this by showing that dimension-reduction techniques, such as Stock and Watson's (2002b) diffusion indices, typically improve the forecasts of economy and inflation indicators, with clear benefits for the estimation of the central bank's reaction function. We now show that the same applies to Brazilian central bankers.

Table 3 shows the results of regressing the SELIC interest rate on the overall principal components, as well as on the forward-looking dataset principal component. The principal components are jointly significant, even if the loading on the second principal component is not

[2] Risk reversal is a difference in 25-delta volatility between puts and calls on out-of-money options on the Brazilian currency.

Table 3
Policy rules based on factors.

	PCA(all)	PCA(fwrd)
Constant	10.3413	11.3537
	(0.0874)	(0.0963)
First principal component	0.0599	0.2240
	(0.0881)	(0.0687)
Second principal component	−0.1753	−0.5071
	(0.0959)	(0.0928)
R^2	0.083	0.126

The table documents factor-based rules for the target interest rate of the Central Bank of Brazil. We regress the target interest rate on the first and second principal components of the macroeconomic and financial variables that we consider. We report two sets of coefficient estimates: PCA(all) uses the complete panel of 142 indicators, whereas PCA(fwrd) focuses only on forward-looking indicators. We also display robust standard errors in parentheses.

statistically different from zero. As expected, the estimates indicate that a greater uncertainty about the future and external deterioration lead to higher interest rates.

We assess whether the reaction function of the Central Bank of Brazil responds to a wider array of indicators by adapting Bernanke and Boivin's (2003) augmented Taylor rule to the weekly frequency as follows:

$$R_t = \rho R_{t-1} + (1-\rho)\Big[\beta_1\left(CPI_{12m} - CPI_{5y}\right) + \beta_2\left(g_{12m} - g_{5y}\right) + \beta_3 \hat{R}_t\Big], \quad (5)$$

where $\hat{R}_t = \hat{c} + \sum_{i=1}^{n} \hat{a}_i F_{it}$, so as to link the target SELIC rate R_t to the diffusion indices F_{it} (namely, the two principal components from a set of indicators) explicitly. We proxy the inflation and growth gaps by the difference between the expected inflation and GDP growth over the next 12 months and the expected inflation and GDP growth in the long run (as measured by market expectations over the next five years in the Focus database). Table 4 shows that the information on the principal components is indeed useful for policymakers in their decision making.

Table 5 shows that the yield curve also responds in a statistically significant manner to the variation in the principal components, even when controlling for the SELIC rate. The first principal component has a positive effect on the yields, with the magnitude appearing to increase with the maturity. This confirms the importance of the external channel transmission for the Brazilian yield curve. The second principal component, which is correlated mainly with uncertainty and real activity growth, has a negative effect on shorter maturities only.

Table 6 reports some descriptive statistics for residuals of the NS-FAVAR models. We find that the FAVAR models do a very good job of fitting the level, slope and curvature of the Brazilian yield curve. The mean absolute errors are not only small, at about 10 bps, but also very stable across maturities. Our findings corroborate the results of Faria and Almeida (2014) and Moench (2008), in that the mean absolute errors increase with the maturity. The largest error for almost every maturity is in the second half of 2008. The only exception is the two-year yield, for which the largest error occurred when the Brazilian Central Bank surprised the market by dropping the SELIC rate to its lowest historical value at the end of 2010.

Overall, we find that the principal components convey important information. The next section examines whether this good in-sample performance also translates into superior forecasts.

4.2. Out-of-sample analysis

This section assesses the forecasting performance of our NS-FAVAR model relative to the extant models in the literature. We evaluate the relative importance of the forward-looking variables by comparing the forecasting abilities of NS-FAVAR(all), which extracts principal components from the full database, and NS-FAVAR(fwrd), which considers forward-looking variables only.

We contemplate a number of alternative forecasting models. As usual, we employ a random walk without drift (RW) as a benchmark.[3] Joslin, Singleton, and Zhu (2011) show that the random walk is actually a very challenging benchmark to beat at shorter forecasting horizons. In addition, we also consider a simple autoregressive model (AR), Diebold and Li's (2006) AR model for the level, slope and curvature factors (DL-AR), Diebold et al.'s (2006) dynamic VAR model (DNS),[4] and Moench's (2008) affine FAVAR using the overall principal components as driving factors for the short rate (A-FAVAR). For each model, we choose the lag structure that minimizes the Bayesian information criterion (BIC). This results in first order specifications for every model except the A-FAVAR in all periods.

The A-FAVAR model employs the overall principal components as driving factors for the short rate. In particular, like Moench (2008), we assume that

$$\begin{pmatrix} F_t \\ r_t \end{pmatrix} = \mu + \Phi(L)\begin{pmatrix} F_t \\ r_t \end{pmatrix} + \omega_t, \quad (6)$$

where r_t denotes the short rate and ω_t is a vector of white noise with covariance matrix Ω. After estimating the parameters in Eq. (6), we impose no-arbitrage considerations by minimizing the market prices of risk (λ_0, λ_1) in

$$A_n = A_{n-1} + B'_{n-1}(\mu - \Omega\lambda_0) + \frac{1}{2}B'_{n-1}\Omega B_{n-1}$$
$$B_n = B'_{n-1}(\Phi - \Omega\lambda_1) - \delta',$$

for some initial conditions (A_0, B_0, δ). Next, we obtain the future values of the n-year zero rate using the affine nature of the model: $\hat{y}^{(n)}_{t+h|t} = \frac{-A_n}{n} - \frac{B'_n}{n}\hat{r}_{t+h|t}$.

We estimate each forecasting model using data from the first week of March 2007 to the last week of December 2011. The forecasting performances are then assessed

[3] The out-of-sample results for the random walk with a drift are considerably worse, and thus, we do not report them here, though they are obviously available from the authors upon request.

[4] Diebold et al. (2006) estimate their VAR model using a Kalman filter. In contrast, we estimate the DNS model in two stages. We begin by extracting the Nelson–Siegel factors, then estimate a VAR model by maximum likelihood estimation, which makes the results directly comparable to those of the other forecasting methods. Nonetheless, it is worth noting that the use of a Kalman filter yields very similar forecasting performances.

Table 4
Augmented Taylor rules.

	(A)	(B)	(C)
Past target interest rate	0.983	0.950	0.955
	(0.003)	(0.011)	(0.010)
CPI forecast for the next 12 months	24.34	10.867	10.783
	(2.804)	(2.422)	(2.536)
GDP growth forecast for the next 12 months	15.924	4.282	4.609
	(3.257)	(1.120)	(1.131)
Predicted target interest rate based on PCA(all)		0.582	
		(0.101)	
Predicted target interest rate based on PCA(fwrd)			0.552
			(0.099)
R^2	0.968	0.969	0.969

The table reports the regression results for Eq. (5). Column (A) displays the coefficient estimates for the traditional Taylor rule, whereas columns (B) and (C) show the estimates for augmented Taylor rules that include the target interest rate predicted by the factor models in Table 3. We report robust standard errors in parentheses.

Table 5
Estimation of yields and principal components.

		Constant	PC1	PC2	SELIC	R^2
1-year	PCA(all)	0.100	0.010	−0.007	0.421	0.46
		(0.003)	(0.003)	(0.002)	(0.069)	
	PCA(fwrd)	0.112	0.006	−0.010	0.249	0.44
		(0.003)	(0.002)	(0.001)	(0.049)	
3-year	PCA(all)	0.109	0.011	0.002	0.377	0.53
		(0.002)	(0.002)	(0.002)	(0.053)	
	PCA(fwrd)	0.117	0.008	−0.004	0.230	0.44
		(0.002)	(0.002)	(0.001)	(0.037)	
5-year	PCA(all)	0.112	0.011	0.000	0.354	0.56
		(0.002)	(0.002)	(0.002)	(0.049)	
	PCA(fwrd)	0.118	0.008	−0.002	0.215	0.47
		(0.002)	(0.002)	(0.001)	(0.035)	
7-year	PCA(all)	0.113	0.011	0.000	0.344	0.57
		(0.002)	(0.002)	(0.002)	(0.047)	
	PCA(fwrd)	0.119	0.009	−0.001	0.208	0.49
		(0.002)	(0.002)	(0.001)	(0.034)	
10-year	PCA(all)	0.114	0.011	−0.001	0.337	0.57
		(0.002)	(0.002)	(0.002)	(0.046)	
	PCA(fwrd)	0.119	0.009	0.000	0.203	0.50
		(0.002)	(0.002)	(0.001)	(0.033)	

The table reports the estimation results for regressing interest rate yields on the first and second principal components of the macroeconomic and financial variables that we consider, as well as on the SELIC rate. We report two sets of coefficient estimates: PCA(all) uses the complete panel of 142 indicators, whereas PCA(fwrd) focuses only on forward-looking indicators. We also display robust standard errors in parentheses.

Table 6
Descriptive statistics for the in-sample absolute errors.

	Maturity	Mean	Std deviation	Maximum	Date
NS-FAVAR(all)	1 year	0.09	0.14	1.13	24/06/2008
	2 years	0.07	0.09	0.58	28/12/2010
	5 years	0.08	0.09	0.60	21/10/2008
	7 years	0.10	0.10	0.86	21/10/2008
	10 years	0.12	0.12	1.13	21/10/2008
NS-FAVAR(fwrd)	1 year	0.09	0.14	1.14	24/06/2008
	2 years	0.07	0.09	0.59	28/12/2010
	5 years	0.09	0.09	0.63	21/10/2008
	7 years	0.10	0.10	0.84	21/10/2008
	10 years	0.12	0.12	1.11	21/10/2008

The table reports the sample mean, standard deviation and maximum values of the in-sample absolute errors (in percentage points) of the NS-FAVAR(all) and NS-FAVAR(fwd) for each maturity. We also display the date at which we observe the largest error in magnitude.

for the remaining 156 weeks, up to the last week of December 2014. We compute h-month ahead predictions by iterating the forecasts in real time, by re-estimating the principal components, the Nelson–Siegel loadings and the

Table 7
Mean absolute forecast errors relative to the random walk model.

	RW	AR	A-FAVAR	DL-AR	DNS	NS-FAVAR(all)	NS-FAVAR(fwrd)
1 month ahead							
12	0.395	0.840**	4.790	0.836**	0.932	0.806*	0.967
36	0.330	1.146	5.846	1.078	1.218	1.078	0.975*
60	0.376**	1.086	4.203	0.990	1.109	1.008	0.970*
84	0.424	0.996	3.674	0.901*	1.000	0.928**	0.955**
120	0.475	0.917	3.318	0.955	0.971	0.854*	0.945
3 months ahead							
12	0.766	1.161	2.983	1.015	1.014	1.000	0.901*
36	0.757	1.214	3.047	1.010	1.029	0.962	0.844*
60	0.769	1.239	2.312	1.009	1.021	0.920	0.847*
84	0.778	1.245	2.316	1.009	1.017	0.900	0.872*
120	0.785	1.250	2.378	1.008	1.015	0.890*	0.902**
6 months ahead							
12	1.253	1.343	2.054	1.001	0.982	1.107	0.894*
36	1.255	1.293	2.021	1.001	1.000	1.000	0.900*
60	1.277	1.247	1.442	1.001	1.015	0.954	0.883*
84	1.290	1.224	1.534	1.001	0.999	0.933	0.874*
120	1.301	1.207	1.614	1.001	0.984	0.919	0.868*
9 months ahead							
12	1.748	1.333	1.486	0.995	0.943	1.087	0.844*
36	1.578	1.223	1.565	0.995	0.964	0.995	0.890*
60	1.599	1.123	1.145	0.997	0.962	0.920	0.831*
84	1.621	1.081	1.285	0.998	0.960	0.887	0.797*
120	1.640	1.051	1.365	1.000	0.958	0.864	0.773*
12 months ahead							
12	1.519	1.133	1.539	0.995	0.870	1.011	0.759*
36	1.268	0.926	1.681	0.994	0.883	0.848	0.740*
60	1.268	0.837	1.275	0.993	0.876	0.754	0.664*
84	1.284	0.804	1.444	0.993	0.871	0.713	0.629*
120	1.297	0.784	1.555	0.994	0.870	0.689**	0.611*

The column RW displays the mean absolute forecast error (in percentage points) of the random walk benchmark, while the other columns report the mean absolute forecast error of each model relative to the random walk. We estimate every model using weekly data from March 2007 to December 2011, and then produce h-month-ahead iterated forecasts, with $h = 1, 3, 6, 9$ and 12, for the period from January 2012 to December 2014. NS-FAVAR(all) refers to the NS-FAVAR model with the principal components of the complete panel of macroeconomic and financial variables. NS-FAVAR(fwrd) considers the principal components based on only the forward-looking indicators.

* We identify the models that are superior at the 10% significance level.

** We identify the models that are superior at the 25% significance level.

model parameters each time we add another week to the estimation window, up to December 2014.[5]

Table 7 reports the mean absolute forecast errors that we obtain for each model across the different maturities and horizons. In contrast to the findings of Moench (2008), the A-FAVAR model does not compare well with the random walk for any horizon or maturity. Similarly, the DL-AR forecasts are reasonably good only at the one-month horizon, though the out-of-sample results are no better than the RW benchmark for longer horizons. This evidence is in line with de Pooter et al.'s (2010) finding that a VAR specification for the Nelson–Siegel factors performs well only for short maturities and horizons. Finally, the DNS forecasts improve on the RW forecasts in the medium-run, reducing the mean absolute forecast errors, for instance, by up to 13% at the one-year horizon.

Both of the NS-FAVAR models perform very well, improving the forecasts by up to 15 bps at the three-month horizon, and by 15–50 bps at horizons of longer than six months. In particular, NS-FAVAR(fwrd) shows the best performance for any horizon over a month, irrespective of the maturity. Indeed, it fares very well, especially for the yields with medium and longer maturities, with decreasing relative mean absolute forecast errors. This suggests that forward-looking indicators are key to explaining the short- and medium-run movements in the yield curve. As a matter of fact, the AR and DL-AR models only outperform NS-FAVAR(fwrd) at the shortest horizon of one month and for the shorter maturities. In turn, NS-FAVAR(all) not only achieves lower mean absolute forecast errors than RW for the longer maturities, but also improves on the AR, DL-AR and A-FAVAR forecasts for every maturity (see also Fig. 1).

These results are very promising. Diebold and Li (2006) and Moench (2008) show that their models only provide better forecasts of the US yield curve than the random walk benchmark at longer horizons, say, six months or more. Altavilla, Giacomini, and Constantini (2014) and

[5] See Marcellino, Stock, and Watson (2006) for an excellent discussion of the relative advantages and drawbacks of direct and iterated AR forecasts.

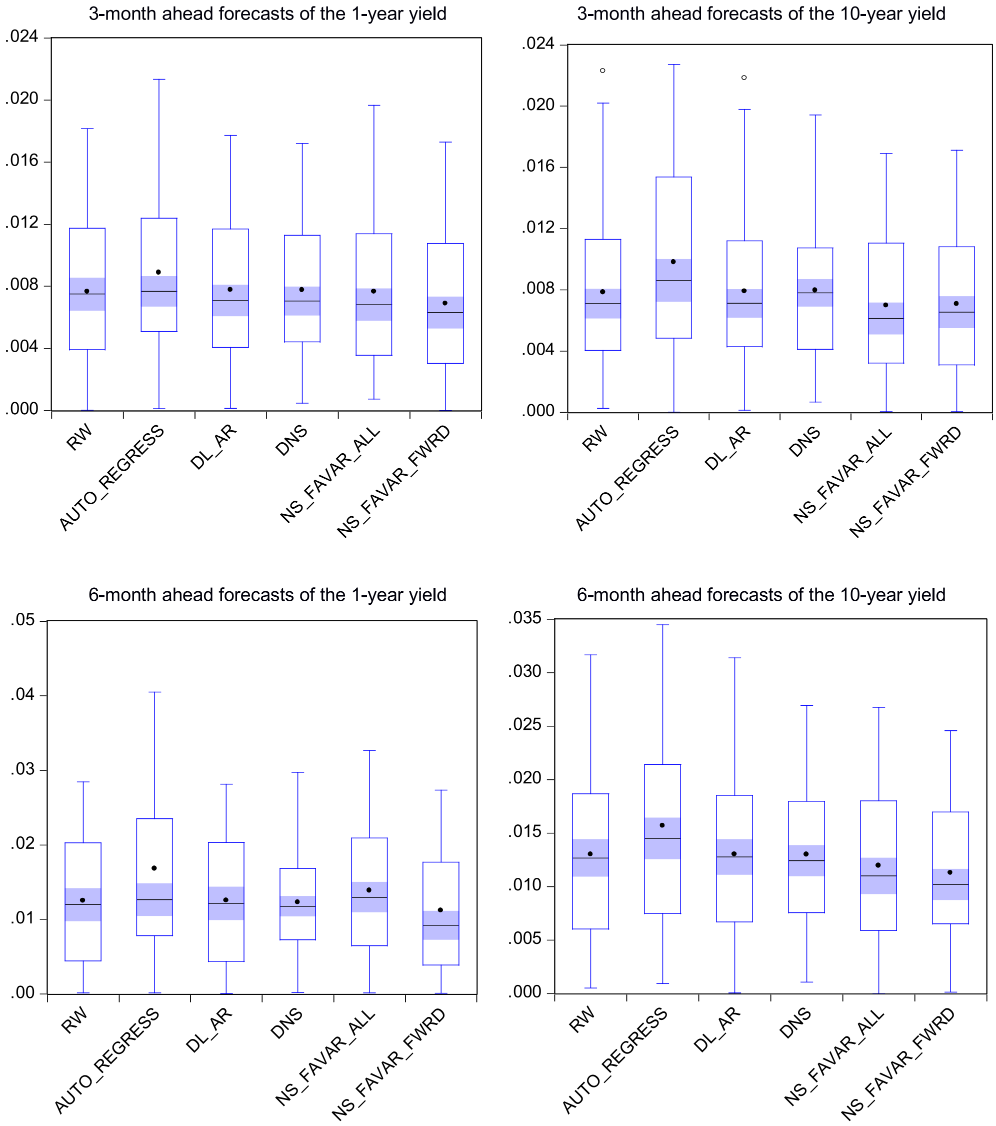

Fig. 1. Box plots of the forecast errors for the 1- and 10-year yields at the 3- and 6-month horizons.

Altavilla, Giacomini, and Ragusa (2014) and Exterkate et al. (2013) are able to beat the random walk benchmark only at short horizons, not longer horizons, though not for every maturity. In stark contrast, our NS-FAVAR models produce improved term structure forecasts for every maturity even at shorter horizons.

Next, we test whether these improvements are indeed statistically significant. To this end, we run a model confidence set (MCS) analysis, as per Hansen, Lunde, and Nason (2011a). This procedure determines the number of superior models within a collection of alterative specifications for a given confidence level. This number obviously depends on how informative the data are. If there is a lot of information in the data, the MCS analysis will select only a few, if not a single model. The main advantage of the MCS is that it is not about comparing predictive abilities with a single benchmark, or any other pairwise comparison. Instead, it treats the performance of each model in a symmetric way, only attempting to identify which models achieve the best out-of-sample predictive powers.[6]

Table 7 identifies (with stars) the superior models for different horizons and maturities, based on a block-bootstrap implementation of the MCS procedure, with blocks of 12 observations. We find that the NS-FAVAR

[6] The MCS procedure determines the number of superior models through a sequence of tests of the null hypothesis of equal predictive ability. The test statistic depends on the loss function of interest. In particular, we employ the conventional mean squared forecast error. The algorithm starts from a set M0 of forecasting models and tests whether they have equal predictive abilities. If the test rejects the null, we eliminate the model with the poorest forecasting performance. We then repeat this procedure until we can no longer reject the null of equal predictive ability. Hansen et al. (2011a) suggest the use of Gonçalves and White's (2005) moving-block bootstrap implementation to control the confidence level. See Hansen, Lunde, and Nason (2011b) for more details.

Table 8
Relative mean absolute forecast errors at the monthly frequency.

	RW	AR	A-FAVAR	DL-AR	DNS	NS-FAVAR(all)	NS-FAVAR(fwrd)
1 month ahead							
12	0.306*	1.436	1.858	1.796	1.155	1.014**	1.045
36	0.321*	1.125	1.857	1.797	1.109	1.157	1.145
60	0.330	0.957*	1.862	1.727	1.062	1.274	1.236
84	0.346	0.967*	1.946	1.718	1.043	1.369	1.328
120	0.358	0.964*	2.192	1.746	1.045	1.439	1.411
3 months ahead							
12	0.724	1.115	1.348	1.253	1.113	0.802*	0.801*
36	0.710	1.339	1.413	1.339	1.177	0.939*	0.962**
60	0.697*	1.346	1.530	1.346	1.121	1.032**	0.985*
84	0.700**	1.347	1.800	1.347	1.087	1.099	0.967*
120	0.703	1.364	2.052	1.364	1.068	1.162	0.948*
6 months ahead							
12	1.278**	0.986	1.076	1.091	1.144	1.026	0.942*
36	1.264**	1.121	1.114	1.121	1.130	1.009**	0.948*
60	1.276	1.135	1.176	1.135	1.069	0.924*	0.975
84	1.286	1.147	1.284	1.147	1.042	0.877*	0.955
120	1.295	1.159	1.388	1.159	1.023	0.838*	0.938
9 months ahead							
12	1.822**	0.991	0.985	1.028	1.047	1.055	0.961*
36	1.669*	1.076	0.998	1.076	1.097	1.058	1.010**
60	1.665	1.061	1.047	1.061	1.039	0.958**	0.941*
84	1.669	1.062	1.120	1.062	1.004	0.897*	0.896*
120	1.675	1.063	1.183	1.063	0.977	0.853*	0.859*
12 months ahead							
12	2.349	1.000	0.998	0.991	0.836	0.955	0.800*
36	2.137	1.036	0.995	1.036	0.863	0.902	0.839*
60	2.071	1.010	1.034	1.010	0.846	0.822	0.780*
84	2.044	0.988	1.107	0.988	0.831	0.770	0.740*
120	2.023	0.972	1.174	0.972	0.817	0.729	0.704*

The column RW displays the mean absolute forecast error (in percentage points) of the random walk benchmark, while the other columns report the mean absolute forecast error of each model relative to the random walk. We estimate every model using monthly data from March 2007 to December 2011, and then produce h-month-ahead iterated forecasts, with $h = 1, 3, 6, 9$ and 12, for the period from January 2012 to December 2014. NS-FAVAR(all) refers to the NS-FAVAR model with the principal components of the complete panel of macroeconomic and financial variables. NS-FAVAR(fwrd) considers the principal components based on only the forward-looking indicators.
* We identify the models that are superior at the 10% significance level.
** We identify the models that are superior at the 25% significance level.

models are among the best models at the 10% significance level for almost every maturity and horizon. In particular, the NS-FAVAR(fwrd) forecasts are usually superior for every maturity at any horizon longer than one month ahead. Its closer competitor is the NS-FAVAR(all), which has a decent performance for any horizon longer than a month. It turns out that the random walk is not actually such a challenging benchmark here. Finally, we fail to uncover any evidence of superior forecasting performance at the usual significance levels for the AR, A-FAVAR, DL-AR and DNS models at horizons of longer than a month.

5. Robustness to different data frequencies and spans

This section reports the results of two robustness checks. First, we repeat the analysis at the monthly frequency in order to make our results more comparable to previous findings in the literature. Second, we increase the length of the out-of-sample period by considering an alternative sample that starts in 2002. As a result of the longer time span, though, we have to drop longer-term yields, as well as some of the variables that we used to extract diffusion indices.

For the monthly analysis, we consider the same data as in Section 4, but looking only at the last week of each month. Table 8 reveals that the NS-FAVAR models perform the best for virtually every yield for any horizon exceeding one month. At the three-month horizon, the NS-FAVAR(fwrd) model shines for every maturity, apart from the three-year yield. At the six- and nine-month horizons, the NS-FAVAR models compete head to head. While the NS-FAVAR(fwrd) has the best performance for the short end of the yield curve, the NS-FAVAR(all) produces lower mean absolute forecast errors for the longer-term yields. At longer horizons, the forecasting performance of the NS-FAVAR(fwrd) model is impressive, reducing the mean absolute forecast errors by 20%–30% relative to the RW benchmark. As expected, the extant models in the literature (A-FAVAR, DL-AR, and DNS) are only superior to the simpler AR and RW alternatives at longer horizons.

Turning to our increase of the time span, recall that longer-term bonds exist only from 2006, with liquidity

Table 9
Mean absolute forecast errors relative to the random walk model for the shorter-term yields.

	RW	AR	A-FAVAR	DL-AR	DNS	NS-FAVAR(all)	NS-FAVAR(fwrd)
1 month ahead							
1	0.246	1.054	1.285	1.510	0.743^{*}	0.982	0.896
2	0.250	1.044	1.660	1.121	0.745^{*}	0.924	0.856
3	0.256	1.037	2.164	1.227	0.836^{*}	0.983	0.918
6	0.287^{*}	1.024	3.257	1.419	1.074	1.125	1.108
12	0.334^{*}	1.039	3.762	1.424	1.297	1.146	1.188
3 months ahead							
1	0.760	1.163	1.295	1.140	0.701^{*}	0.883	0.813
2	0.770	1.131	1.374	1.032	0.808^{*}	0.918	0.869
3	0.780	1.105	1.503	1.064	0.943	0.934	0.902^{*}
6	0.812^{*}	1.059	1.796	1.120	1.126	0.993^{*}	1.015^{**}
12	0.858	1.040	2.033	1.110	1.354	0.928^{*}	0.981
6 months ahead							
1	1.451	1.357	1.171	1.056	0.833	0.812	0.768^{*}
2	1.455	1.295	1.211	1.020	0.919	0.829	0.810^{*}
3	1.457	1.249	1.277	1.035	1.000	0.843^{*}	0.843^{*}
6	1.485	1.139	1.370	1.038	1.196	0.838^{*}	0.865
12	1.523	1.066	1.462	0.999	1.484	0.786^{*}	0.818
9 months ahead							
1	2.051	1.400	1.113	1.032	0.983	0.719	0.699^{*}
2	2.046	1.339	1.148	1.026	1.046	0.739	0.724^{*}
3	2.047	1.290	1.196	1.038	1.110	0.746^{**}	0.740^{*}
6	2.066	1.171	1.240	1.013	1.294	0.743^{*}	0.765
12	2.093	1.072	1.308	0.956	1.612	0.746^{*}	0.787
12 months ahead							
1	2.471	1.410	1.131	1.002	1.178	0.676	0.654^{*}
2	2.463	1.363	1.162	1.015	1.222	0.701	0.692^{*}
3	2.464	1.314	1.199	1.026	1.272	0.717^{*}	0.724^{**}
6	2.465	1.190	1.245	1.007	1.449	0.747^{*}	0.778
12	2.476	1.079	1.331	0.946	1.777	0.773^{*}	0.814

The column RW displays the mean absolute forecast error (in percentage points) of the random walk benchmark, while the other columns report the mean absolute forecast error of each model relative to the random walk. We estimate every model using weekly data from March 2002 to December 2004, and then produce h-month-ahead iterated forecasts, with $h = 1, 3, 6, 9$ and 12, for the period from January 2005 to December 2014. NS-FAVAR(all) refers to the NS-FAVAR model with the principal components of the complete panel of macroeconomic and financial variables. NS-FAVAR(fwrd) considers the principal components based on only the forward-looking indicators.
* We identify the models that are superior at the 10% significance level.
** We identify the models that are superior at the 25% significance level.

picking up only in 2007. Thus, we restrict our attention to shorter-term yields, with maturities of up to 12 months, as per Faria and Almeida (2014) and Vicente and Tabak (2008). This allows us to increase the time span significantly, starting the sample period in 2002 rather than 2007. However, we have to drop some of the macroeconomic indicators that we employ to extract the diffusion indices, as the list of variables in the Appendix shows that we have information since 2002 for only 113 of the 142 macroeconomic and financial indicators that we consider. We estimate the model initially using data from January 2002 to December 2004, and then assess forecasting performances using data from January 2005 to December 2014.

Table 9 reports the out-of-sample results for the 1- to 12-month yields. The NS-FAVAR models remain dominant, comparing favorably with the alternative forecasting models. Although the NS-FAVAR(fwrd) shows a higher predictive ability than the random walk for almost every yield at any forecast horizon, it outperforms the other forecasting models only for the short-end of the yield curve, at horizons of longer than three months. In turn, the NS-FAVAR(all) works best for medium-term maturities at horizons of greater than one month. Perhaps surprisingly, the dynamic Nelson–Siegel model has the smallest mean absolute forecast errors for the shorter-term yields at the one- and three-month horizons. However, the DNS performance deteriorates considerably for longer horizons, obtaining the worst results for the 12-month-ahead forecasts. In turn, it is worth noting that the RW forecasts are only significantly better than the other forecasts for six- and 12-month yields at the one-month horizon.

6. Conclusion

This paper proposes that future values of yields at different maturities be forecast by means of a FAVAR model for the level, slope and curvature of the yield curve. In particular, we estimate an augmented VAR model for a system that includes not only the Nelson–Siegel factors of the Brazilian yield curve, but also the principal components of a large number of macroeconomic and financial indicators. We show that our forecasting approach outperforms the extant models in the literature, including the random walk benchmark, even at shorter horizons. Further analysis reveals that the use of forward-looking state variables is vital for producing better forecasts.

We defer the assessment of external validity to future research. In particular, we plan to examine whether similar forecast improvements can be observed using US data, as there is no reason to believe that our findings will automatically carry through. First, the term structure of interest rates in Brazil may have very specific dynamics. Second, surprisingly, it is easier to gather a large number of forward-looking indicators in Brazil than in the US. This may limit the predictive ability of the NS-FAVAR model, given that market expectations about the economic and financial outlooks are very informative.

Acknowledgments

We are grateful to seminar participants at the Meetings of the Brazilian Finance Society, and to Caio Almeida for sharing the code used to estimate Moench's (2008) model. Fernandes is thankful for financial support from FAPESP (2013/22930-0) and CNPq (302272/2014-3). The usual disclaimers apply.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at http://dx.doi.org/10.1016/j.ijforecast.2016.08.001.

References

Altavilla, C., Giacomini, R., & Constantini, R. (2014). Bond returns and market expectations. *Journal of Financial Econometrics*, *12*, 708–729.

Altavilla, C., Giacomini, R., & Ragusa, G. (2014). Anchoring the yield curve using survey expectations. Working Paper 1632, European Central Bank.

Bernanke, B., & Boivin, J. (2003). Monetary policy in a data-rich environment. *Journal of Monetary Economics*, *50*, 525–546.

Bernanke, B., Boivin, J., & Eliasz, P. (2005). Measuring the effects of monetary policy: A factor-augmented vector autoregressive (FAVAR) approach. *Quarterly Journal of Economics*, *120*, 387–422.

de Pooter, M., Ravazzolo, F., & van Dijk, D. (2010). Term structure forecasting using macro factors and forecast combination. Federal Reserve of the Board of Governors International Finance Discussion Paper 2010–993.

Diebold, F., & Li, C. (2006). Forecasting the term structure of government bond yields. *Journal of Econometrics*, *130*, 337–364.

Diebold, F., Rudebusch, G., & Aruoba, S. (2006). The macroeconomy and the yield curve: A dynamic latent factor approach. *Journal of Econometrics*, *131*, 309–338.

Doz, C., Giannone, D., & Reichlin, L. (2012). A quasi-maximum likelihood approach for large approximate dynamic factor models. *Review of Economics and Statistics*, *94*, 1014–1024.

Exterkate, P., Van Dijk, D., Heij, C., Patrick, J. F., & Groenen, P. J. F. (2013). Forecasting the yield curve in a data-rich environment using the factor-augmented Nelson-Siegel model. *Journal of Forecasting*, *32*, 193–214.

Faria, A., & Almeida, C. (2014). Forecasting the Brazilian term structure using macroeconomic factors. *Brazilian Review of Econometrics*, *34*, 45–77.

Gonçalves, S., & White, H. (2005). Bootstrap standard error estimates for linear regression. *Journal of the American Statistical Association*, *100*, 970–979.

Hansen, P. R., Lunde, A., & Nason, J. M. (2011a). The model confidence set. *Econometrica*, *79*, 453–497.

Hansen, P. R., Lunde, A., & Nason, J. M. (2011b). Supplement to the model confidence set. *Econometrica Supplementary Material*, *79*.

Joslin, S., Singleton, K., & Zhu, H. (2011). A new perspective on Gaussian dynamic term structure models. *Review of Financial Studies*, *24*, 926–970.

Marcellino, M., Stock, J., & Watson, M. (2006). A comparison of direct and iterated AR methods for forecasting macroeconomic series h-steps ahead. *Journal of Econometrics*, *135*, 499–526.

Moench, E. (2008). Forecasting the yield curve in a data-rich environment: A no-arbitrage factor-augmented VAR approach. *Journal of Econometrics*, *146*, 26–43.

Nelson, C., & Siegel, A. (1987). Parsimonious modeling of yield curves. *Journal of Business*, *60*, 473–489.

Stock, J., & Watson, M. (2002a). Forecasting using principal components from a large number of predictors. *Journal of the American Statistical Association*, *97*, 1167–1179.

Stock, J., & Watson, M. (2002b). Macroeconomic forecasting using diffusion indexes. *Journal of Business and Economic Statistics*, *20*, 147–162.

van Dijk, D., Koopman, S. J., van der Wel, M., & Wright, J. H. (2014). Forecasting interest rates with shifting endpoints. *Journal of Applied Econometrics*, *29*, 693–712.

Vicente, J., & Tabak, B. (2008). Forecasting bond yields in the Brazilian fixed income market. *International Journal of Forecasting*, *24*, 490–497.

Fausto Vieira holds a Master degree in Economics from Sao Paulo School of Economics, FGV and he is a Ph.D. candidate at the same School. Research interests include asset pricing, empirical finance, and financial econometrics.

Marcelo Fernandes holds a Ph.D. in Management Science (Finance) from the Solvay Business School, Université Libre de Bruxelles (1999). He is currently Professor of Finance both at the School of Economics and Finance of Queen Mary, University of London and at the Sao Paulo School of Economics, FGV. Research interests include asset pricing, empirical finance, empirical market microstructure, and financial econometrics.

Fernando Chague holds a Ph.D. in Economics from the University of North Carolina at Chapel Hill (2012). He is currently Assistant Professor of Economics at the School of Economics, Business and Accounting of the University of Sao Paulo. Research interests include asset pricing, empirical finance, and financial econometrics.