

Anchoring the Yield Curve Using Survey Expectations¹

Carlo Altavilla² Raffaella Giacomini³ Giuseppe Ragusa⁴
European Central Bank University College London Luiss University, Rome
CeMMAP

Monday 17th October, 2016

¹We would like to thank the editor, two anonymous referees and M. Bauer, J. Christensen, C. Favero, D. Giannone, A. Justiniano, M. Modugno, G. Rudebush, K. Sheppard, K. Singleton as well as conference and seminar participants to the Society for Economic Dynamics, Cyprus, Econometric Society Summer Meeting, Malaga, European Central Bank, Bank of England, Oxford University, Bank of Italy, Luiss University, Autònoma Barcelona, Chicago Fed, St Louis Fed and San Francisco Fed. Raffaella Giacomini gratefully acknowledges support from the British Academy Mid-career Fellowship. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the European Central Bank.

²E-mail: carlo.altavilla@ecb.int

³University College London, Gower Street WC1E6BT London, UK. Email: r.giacomini@ucl.ac.uk

⁴Luiss University, Viale Romania, 32 - 00197 Rome, Italy. E-mail: gragusa@luiss.it

Abstract

The dynamic behavior of the term structure of interest rates is difficult to replicate with models, and even models with a proven track record of empirical performance have underperformed since the early 2000s. On the other hand, survey expectations can accurately predict yields, but they are typically not available for all maturities and/or forecast horizons. We show how survey expectations can be exploited to improve the accuracy of yield curve forecasts given by a base model. We do so by employing a flexible exponential tilting method that anchors the model forecasts to the survey expectations, and we develop a test to guide the choice of the anchoring points. The method implicitly incorporates into yield curve forecasts any information that survey participants have access to - such as information about the current state of the economy or forward-looking information contained in monetary policy announcements - without the need to explicitly model it. We document that anchoring delivers large and significant gains in forecast accuracy relative to the class of models that are widely adopted by financial and policy institutions for forecasting the term structure of interest rates.

JEL Classification Codes: G1; E4; C5

Keywords: Term Structure Models; Exponential Tilting; Blue Chip Analysts Survey; Forecast Performance; Monetary Policy Forward Guidance; Macroeconomic Factors

1 Introduction

The term structure of interest rates contains crucial information for both policymakers' and investors' decisions. Yet, in spite of a vast and growing literature on yield curve modeling, no single approach has emerged that can accurately describe the dynamic behavior of yields. Two popular classes of yield curve models considered in the literature are no-arbitrage dynamic latent factor models (Duffie and Kan, 1996; Litterman et al., 1991; Dai and Singleton, 2000) and the Dynamic Nelson and Siegel (DNS) model of Diebold and Li, 2006. These models share a similar state-space structure in which the yields depend on three dynamic latent factors (level, slope, and curvature) extracted from the cross-section of yields, but differ in the restrictions they impose on the model's parameters. Although the latter have become the leading method for yield curve forecasting at many policy institutions (BIS, 2005) due to their successful empirical performance (Diebold and Li 2006), one of the findings of this paper is that their performance has deteriorated in recent years. The fact that the three-factor structure is not sufficient to capture the dynamics of yields has been documented before (e.g., Diebold and Rudebusch, 2012; Mönch, 2008), and a general consensus has emerged in the literature that one must look beyond the cross-section of yields to pin down the dynamic behavior of interest rates, for example, by enlarging the model's information set with either observable macroeconomic factors (Diebold et al., 2006; Ang and Piazzesi, 2003; Hördahl et al., 2006; Rudebusch and Wu, 2008; Mönch, 2008; Coroneo et al., 2015) or latent "hidden" factors (Joslin et al., 2010; Duffee, 2011). A thorough review of the literature on the connections between analysis of the term structure and the macroeconomics is Gürkaynak and Wright (2012).

This paper's premise is that latent factor models neglect a key determinant of yield dynamics: expectations about future economic developments. It is a well-documented fact that expectations contained in survey data can accurately forecast key macroeconomic variables, such as GDP, inflation, and yields, especially at short forecast horizons (Stark, 2010; Chun, 2012), and several recent papers have utilized survey data in the analysis of the term structure of interest rates. For example, Chun (2011) uses Blue Chip Financial Analysts (henceforth BC) forecasts as observable factors in a no-arbitrage dynamic latent factor model; Chernov and Mueller (2012) develop a model that incorporates survey expectations and links them to the "hidden factor" of Joslin et al. (2010) and Duffee (2011); Dijk et al. (2014) use (longer-horizon) survey expectations to improve

estimates of some parameters in the DNS model, and Kim and Orphanides (2012) use survey data to overcome some small-sample estimation problems in no-arbitrage dynamic latent factor models.

Rather than incorporating survey data into the model, we employ a formal method that anchors segments of the yield curve forecasts to the corresponding survey expectations and transmits the superior forecasting ability to the rest of the yield curve. In essence, the anchoring constrains the dynamics of some yields to replicate those of the survey expectations and thus implicitly incorporates into the forecasts of the whole yield curve any information that survey participants have access to, without the need to explicitly model it. This can include information about the current state of the economy that survey participants deem relevant for predicting future interest rates and that they potentially extract from large dimensional data sets. In this respect, survey expectations offer the possibility to capture both observable and “hidden” factors that can explain yield curve dynamics (as also argued by Duffee (2011)). Survey expectations can also reflect additional useful information, such as nonlinearities (for example, the zero-lower bound constraint), structural change, and information about the future course of monetary policy that may be difficult to capture with existing backward-looking models. In this paper, we stress in particular the role played by the ability of survey participants to capture the kind of forward-looking information about interest rates that is increasingly contained in monetary policy announcements.

An important question we address is which segments of the yield curve one should anchor, as one typically has access to survey expectations about a subset of points along the yield curve. Our main result is to derive a testable condition such that anchoring delivers an improvement in *density forecast* accuracy for the whole yield curve. In our data, we found that the largest improvement is offered by anchoring using the 3-month survey forecast. As a quick visualization of the effects of anchoring, consider Figure 5, which shows that the method shifts the yield curve forecast toward the actual realization, resulting in sizable accuracy improvements that are particularly visible in regions of the yield curve near the anchoring point.

The theoretical justification of the method is based on exponential tilting (see Robertson et al., 2005 and Giacomini and Ragusa, 2013). The testable condition to guide the choice of anchoring points is a new result.

We apply the method to incorporate Blue Chip financial analysts' monthly expectations about yields into yield curve forecasts based on the DNS model. It is worth emphasizing that, although we take the DNS model as a benchmark due to its popularity in the forecasting literature, the anchoring method is more generally valid and could be applied to any alternative model of the yield curve.

We find that the anchoring procedure results in forecasts that significantly outperform those from the base model. The accuracy gains are sizable, averaging about 30% and up to 52%. The anchored forecast is also the only one that was able to beat the random walk over the period 2000-2011. These results are robust to considering a subsample that ends in 2008, which suggests that the good performance of our method is not solely driven by the fact that survey participants correctly incorporate zero lower bound constraints. This is also the reason why we don't consider in our comparison models that explicitly take the zero lower bound constraint into account (for an interesting recent example, see Christensen and Rudebusch (2015)).

Although these improvements are important on their own, we provide further insight into the economic forces driving the superior performance of the anchored forecasts. We find that the anchored forecasts implicitly incorporate measures of real activity and forward-looking information contained in monetary policy announcements. The ability of the anchoring method to incorporate the information contained in monetary policy announcements, in particular, has two important implications. The first is that the anchoring method is likely to become even more useful as a practical tool for forecasters and central bankers in the future, now that forward guidance has been formally adopted by several central banks around the world, including the Federal Reserve, the Bank of England and the ECB. The second is that any successful attempt to explicitly model the dynamics of yields should acknowledge the value of forward-looking information.

The paper is organized as follows. Section 2 reviews the DNS model. Section 3 describes the anchoring method and derives the test for choosing the anchoring points. Section 4 contains the empirical results and Section 5 concludes. Appendix A describes the data; Appendix B reports the in-sample estimation results of the DNS model and Appendix C considers conditional forecasting.

2 The DNS model and its variants

The DNS model introduced by Diebold and Li (2006) for an m -dimensional vector of yields y_t with typical element $y_t(\tau)$, where τ is the maturity, is given by:

$$y_t(\tau) = \beta_{1t} + \beta_{2t} \left(\frac{1 - e^{-\lambda\tau}}{\lambda\tau} \right) + \beta_{3t} \left(\frac{1 - e^{-\lambda\tau}}{\lambda\tau} - e^{-\lambda\tau} \right) + u_t(\tau), \quad (1)$$

where the dynamic factors β_{1t} , β_{2t} , and β_{3t} are interpreted as the level, slope, and the curvature of the yield curve and λ is a fixed parameter governing the exponential decay rate of the first and second component in (1). Following Diebold and Li (2006), we let $\lambda = 0.069$, the value that maximizes the loading on the curvature factor for the yields with maturity 30 months.

When we observe a series of yields $y_t(\tau_i)$, for a set of n maturities, $\tau_1 < \tau_2 < \dots < \tau_n$, the yield curve can be estimated by cross sectional regressions. From these regressions we obtain a series of estimated factors $\left\{ \hat{\beta}_{1t}, \hat{\beta}_{2t}, \hat{\beta}_{3t} \right\}_{t \leq T}$ which can themselves be modeled dynamically to provide out-of-sample forecasts of the factors and thus of the entire yield curve. The main drawbacks of this two-step approach are, first, that the uncertainty associated with the estimation of the factors is not acknowledged in the second step when the yields are calculated, and, second, the loss of efficiency arising from not exploiting the correlation structure.

To mitigate these problems, we follow Diebold et al. (2006) and exploit the state space representation of the DNS model. To describe this representation, rewrite equation (1) by stacking the yields for the n maturities

$$y_t = \Gamma_y \beta_t + u_t^y, \quad (2)$$

where $y_t = (y_t(\tau_1), \dots, y_t(\tau_n))'$, $\beta_t = (\beta_{1t}, \beta_{2t}, \beta_{3t})'$, and Γ is a $n \times 3$ matrix of loadings whose i -th row is equal to $(1, (1 - e^{-\lambda\tau_i})/\lambda\tau_i, (1 - e^{-\lambda\tau_i})/\lambda\tau_i - e^{-\lambda\tau})$. The disturbance vector is given by $u_t^y = (u_t(\tau_1), \dots, u_t(\tau_n))'$, $u_t^y \sim \mathcal{N}(0, Q)$. The time series process for β_t is a vector autoregression:

$$\beta_{t+1} = \bar{\beta} + \Phi_y \beta_t + \xi_t, \quad \xi_t \sim \mathcal{N}(0, R), \quad (3)$$

where $\bar{\beta}$ is a (3×1) vector and Φ_y is a (3×3) matrix. Equations (2) and (3) constitute the basic version of the model for the yield curve we consider here. We will refer to this model as the yield only model to distinguish it from its extension in which the dynamics of yields

and macroeconomic variables are modeled in terms of common factors, as recently proposed by Coroneo et al. (2015). We call this model the macro augmented DNS model.

Let x_t be a $(p \times 1)$ vector of macroeconomic variables. The macro augmented DNS model consists of a state space model with measurement equation

$$\begin{pmatrix} y_t \\ x_t \end{pmatrix} = \begin{pmatrix} 0 \\ \bar{x} \end{pmatrix} + \begin{bmatrix} \Gamma_y & 0 \\ \Gamma_{xy} & \Gamma_x \end{bmatrix} \begin{pmatrix} \beta_t \\ \theta_t \end{pmatrix} + \begin{pmatrix} u_t^y \\ u_t^x \end{pmatrix}, \quad (4)$$

and transition equation

$$\begin{pmatrix} \beta_{t+1} \\ \theta_{t+1} \end{pmatrix} = \begin{pmatrix} \bar{\beta} \\ \bar{\theta} \end{pmatrix} + \begin{bmatrix} \Phi_y & \Phi_{yx} \\ \Phi_{xy} & \Phi_x \end{bmatrix} \begin{pmatrix} \beta_t \\ \theta_t \end{pmatrix} + \begin{pmatrix} \xi_t^y \\ \xi_t^x \end{pmatrix}. \quad (5)$$

The dimension of Γ_x depends on the number of common factors that determine the dynamics of the macro variables. Several studies have shown that two factors, one real and one nominal, are sufficient to capture the dynamics of a large variety of macroeconomic indicators for the US (Sargent et al., 1977; Giannone et al., 2005). We accordingly choose two factors so that the dimension of Γ_x is $(p \times 2)$.¹

The idiosyncratic disturbances, $u_t \equiv (u_t^y, u_t^x)'$, and the transition equation disturbances, $\xi_t \equiv (\xi_t^y, \xi_t^x)'$, are assumed to be Gaussian

$$\begin{pmatrix} u_t^y \\ u_t^x \end{pmatrix} \sim \mathcal{N} \left(0, \begin{bmatrix} Q_y & 0 \\ 0 & Q_x \end{bmatrix} \right), \quad \begin{pmatrix} \xi_t^y \\ \xi_t^x \end{pmatrix} \sim \mathcal{N} \left(0, \begin{bmatrix} R_y & R_{yx} \\ R_{xy} & R_x \end{bmatrix} \right). \quad (6)$$

Both Q_y and Q_x are diagonal, while the covariance of ξ_t is left unrestricted. The assumption that the variance of the idiosyncratic disturbances is diagonal, which implies that the deviations of yields of various maturities from the yield curve are uncorrelated, is quite standard (Diebold et al., 2006). This assumption also gives computational tractability given the large number of yields and macro factors we use.² The processes u_t and ξ_t are also assumed to be independent, and thus β_t and θ_t are not allowed to react to shocks to y_t and x_t .

¹Coroneo et al. (2015) consider a similar model to ours using different data and a different sample period and also find that two factors capture well the dynamics of macro economic variables for the US.

²Coroneo et al. (2015) relax this restriction by allowing the idiosyncratic errors to follow a first order autoregressive process. However, they restrict the variance of this process to be fixed in an arbitrary way.

We perform estimation by maximum likelihood using the Expectation Maximization (EM) algorithm, which requires at each iteration only one run of the Kalman smoother. As reported in (Doz et al., 2012), this estimation technique is feasible when the number of variables is large, and is robust to non Gaussianity and to cross-sectional correlation in u_t .

Density forecasts of yields at different maturities can be obtained as

$$y_{t+h}|y_t, \dots, y_1, x_1, \dots, x_t \sim \mathcal{N}(\hat{\mu}_{t+h}, \hat{\Sigma}_{t+h}), \quad h = 1, \dots, H, \quad (7)$$

with $\hat{\mu}_{t+h} = \hat{\Gamma}_y E[\beta_{t+h}|y_t, \dots, y_1, x_1, \dots, x_t]$, $\hat{\Sigma}_{t+h} = \hat{\Gamma}_y \text{Var}[\beta_{t+h}|y_t, \dots, y_1, x_1, \dots, x_t] \hat{\Gamma}_y' + \hat{Q}_y$, where the conditional mean and the conditional variance of β_{t+h} are those from the Kalman recursions with missing observations from $t + 1$ to $t + h$.

3 The anchoring method

In this section, we illustrate the anchoring method for incorporating the information contained in survey expectations into an existing model-based forecast. We then discuss how to choose the anchoring points.

The method is presented without reference to a specific forecasting model, as it can be applied to any model that provides a density forecast. We make the simplifying assumption that the sequence of h -step- ahead density forecasts for the vector of yields is normal with (conditional) mean $\hat{\mu}_{t+h}$ and variance $\hat{\Sigma}_{t+h}$,³

$$y_{t+h}|y_t, \dots, y_1, x_1, \dots, x_t \sim \mathcal{N}(\hat{\mu}_{t+h}, \hat{\Sigma}_{t+h}), \quad h = 1, \dots, H, \quad t = 1, \dots, T. \quad (8)$$

At time t , we observe the h -step ahead survey forecast for yields for the first $r < m$ maturities (τ_1, \dots, τ_r) , that we denote as $\tilde{\mu}_{t+h,1:r}$.⁴ Let $y_{t,1:r}$ denote the $r \times 1$ sub-vector of y_t containing yields for maturities $(\tau_1, \tau_2, \dots, \tau_r)$.

We approach the problem of incorporating $\tilde{\mu}_{t+h,1:r}$ into the forecast from an information theoretic point of view, by projecting the density forecast f_t onto the space of densities that have conditional mean equal to the survey forecasts for maturities τ_1, \dots, τ_r . More formally, this

³See Giacomini and Ragusa (2013) for the general case of a nonnormal density forecast.

⁴In the interest of clarity, we consider only the case in which the survey forecasts considered are for maturities $\tau_1, \tau_2, \dots, \tau_r$. It is immediate to extend the results to the case of non-contiguous maturities.

set of densities can be characterized as

$$\tilde{\mathcal{H}}_{t+h} = \left\{ h_t : \int y_{t+h,1:r} h_t(y_{t+h}) dy_{t+h} = \tilde{\mu}_{t+h,1:r} \right\}. \quad (9)$$

It is important to note that no constraints are imposed on the forecasts of yields at longer maturities, $\tau_{r+1}, \dots, \tau_m$. The idea is to select the density in $\tilde{\mathcal{H}}_{t+h}$ that is "closest" to the model-based density forecast f_t , where closeness is measured by the Kullback-Leibler information criterion:

$$h_t^*(y_{t+h}) = \arg \min_{h \in \tilde{\mathcal{H}}_{t+h}} \int \log \left(\frac{h_t(u)}{f_t(u)} \right) h_t(u) du. \quad (10)$$

Minimization problems such as (10) play an important role in statistics and econometrics (Csiszár, 1975; White, 1982; Kitamura and Stutzer, 1997; Newey and Smith, 2004; Ragusa, 2011), and they have been considered in the forecasting literature by Robertson et al. (2005) and Giacomini and Ragusa (2013). Any of the previous references show that the solution is a new multivariate density:

$$h_t^*(y_{t+h}) = \exp \left\{ \zeta_t + \xi_t' [y_{t+h,1:r} - \tilde{\mu}_{t+h,1:r}] \right\} f_t(y_{t+h}), \quad (11)$$

where ζ_t and ξ_t are parameters chosen in such a way that $h_t^*(y_{t+h}) \in \tilde{\mathcal{H}}_{t+h}$. For the special case of a base density that is multivariate normal, Giacomini and Ragusa (2013) show that we have the following analytical expression for $h_t^*(y_{t+h})$:

$$h_t^*(y_{t+h}) = (2\pi)^{-\frac{m}{2}} \left| \hat{\Sigma}_{t+h} \right|^{-\frac{1}{2}} \exp \left\{ -\frac{1}{2} (y_{t+h} - \mu_{t+h}^*)' \hat{\Sigma}_{t+h}^{-1} (y_{t+h} - \mu_{t+h}^*) \right\}, \quad (12)$$

with

$$\mu_{t+h}^* = \begin{pmatrix} \tilde{\mu}_{t+h,1:r} \\ \hat{\mu}_{t+h,r+1:m} - \hat{\Sigma}_{t+h,21} \left(\hat{\Sigma}_{t+h,11} \right)^{-1} (\hat{\mu}_{t+h,1:r} - \tilde{\mu}_{t+h,1:r}) \end{pmatrix} \quad (13)$$

and $\hat{\Sigma}_{t+h,11}$ and $\hat{\Sigma}_{t+h,21}$ are blocks of the partitioned matrix $\hat{\Sigma}_{t+h}$:

$$\hat{\Sigma}_{t+h} = \begin{pmatrix} \hat{\Sigma}_{t+h,11} & \hat{\Sigma}_{t+h,12} \\ \hat{\Sigma}_{t+h,21} & \hat{\Sigma}_{t+h,22} \end{pmatrix}. \quad (14)$$

Thus, the solution to (10) is a normal density with the same variance as the initial forecast density, $\hat{\Sigma}_{t+h}$, but a mean that is equal to the survey forecast for those yields that are directly restricted, and for the remaining yields it is equal to a combination between the model forecast and the discrepancy between the survey and the restricted model forecasts. The effect of anchoring the first r yields to the survey forecasts on the other yields depends on this discrepancy and on $\hat{\Sigma}_{t+h}$. Since forecasts of yields at different maturities are generally positively correlated, this implies that when the model forecast is larger (smaller) than the survey forecast $\mu_{t+h,r+1:N}$ is adjusted downwards (upwards).

3.1 Where to anchor?

A natural question is which survey forecasts one should use to anchor the yield curve, given that survey expectations could be available for a number of yields. Proposition 1 derives testable conditions that provide guidance on where to anchor the yield curve. The sufficient condition in Proposition 1, in particular, has the intuitive interpretation that anchoring delivers accuracy improvements for the whole density forecast if the survey forecast contains additional information relative to the model-based forecast (for a single yield, the condition is equivalent to a normalized version of the null hypothesis of a forecast encompassing test for the survey relative to the model forecast, see e.g., Clark and McCracken (2001, page 90)).

Let $\hat{e}_{t+h,1:r} = y_{t+h,1:r} - \hat{\mu}_{t+h,1:r}$ and $\tilde{e}_{t+h,1:r} = y_{t+h,1:r} - \tilde{\mu}_{t+h,1:r}$ denote the $r \times 1$ vectors containing the model- and survey-based h -step-ahead forecasts for yields with maturities $(\tau_1, \tau_2, \dots, \tau_r)$, respectively.

Proposition 1. *A sufficient condition for the anchored density forecast $h_t^*(y_{t+h})$ to be more accurate than the base forecast $f_t(y_{t+h})$ according to the logarithmic scoring rule, i.e.,*

$$E \left[\log \left(\frac{h_t^*(y_{t+h})}{f_t(y_{t+h})} \right) \right] > 0, \quad (15)$$

is that

$$E[S_t] \geq 0, \quad (16)$$

where $S_t := (\hat{e}_{t+h,1:r} - \tilde{e}_{t+h,1:r})' (\hat{\Sigma}_{t+h,11})^{-1} \tilde{e}_{t+h,1:r}$. A sufficient and necessary condition is that

$$E[S_t + N_t] \geq 0, \quad (17)$$

where $N_t := \frac{1}{2}(\hat{e}_{t+h,1:r} - \tilde{e}_{t+h,1:r})' (\hat{\Sigma}_{t+h,11})^{-1} (\hat{e}_{t+h,1:r} - \tilde{e}_{t+h,1:r})$.

Proof. Since

$$E \left[\log \left(\frac{h_t^*(y_{t+h})}{f_t(y_{t+h})} \right) \right] = E [\zeta_t + \xi_t \tilde{e}_{t+h}(\tau)], \quad (18)$$

it is sufficient to establish that the expectations of both terms of the right-hand side of the previous display are both positive to establish that condition (17) is necessary and sufficient. To this end, we show that $\zeta_t = N_t$ and $E[\xi_t' \tilde{e}_{t+h,1:r}] = E[S_t]$. Since N_t is a.s. positive its expected value is also positive. Thus, positiveness of $E[S_t]$ is sufficient for $E[\zeta_t + \xi_t \tilde{e}_{t+h}(\tau)]$.

Let J be the $(r \times m)$ matrix selecting the element of y_{t+h} and μ_{t+h}^* corresponding to maturity $(\tau_1, \tau_2, \dots, \tau_r)$, i.e., $Jy_{t+h} = y_{t+h,1:r}$. The density $h_t^*(y_{t+h}) = \exp \{ \zeta_t + \xi_t' [Jy_{t+h} - J\mu_{t+h}^*] \} f_t(y_{t+h})$ is given by

$$h_t^*(y_{t+h}) = (2\pi)^{-\frac{m}{2}} \left| \hat{\Sigma}_{t+h} \right|^{-\frac{1}{2}} \times \exp \left\{ -\frac{1}{2}(y_{t+h} - \hat{\mu}_{t+h})' \hat{\Sigma}_{t+h}^{-1} (y_{t+h} - \hat{\mu}_{t+h}) + \zeta_t + \xi_t' [Jy_{t+h} - J\mu_{t+h}^*] \right\}.$$

The term inside the exponential can be written as

$$\begin{aligned} & -\frac{1}{2}(y_{t+h} - \hat{\mu}_{t+h})' \hat{\Sigma}_{t+h}^{-1} (y_{t+h} - \hat{\mu}_{t+h}) + \zeta_t + \xi_t' [Jy_{t+h} - J\mu_{t+h}^*] \\ & = -\frac{1}{2} \left(y_{t+h}' \hat{\Sigma}_{t+h}^{-1} y_{t+h} - y_{t+h}' b - b' y_{t+h} \right) + c, \end{aligned}$$

where $b = J'\xi + \hat{\Sigma}_{t+h}^{-1} \hat{\mu}_{t+h}$ and $c = \zeta_t - \xi_t' J\mu_{t+h}^* - \frac{1}{2} \hat{\mu}_{t+h}' \hat{\Sigma}_{t+h}^{-1} \hat{\mu}_{t+h}$. Completing the squares by adding and subtracting $b' \hat{\Sigma}_{t+h} b$ yields $-\frac{1}{2} \left(y_{t+h}' \hat{\Sigma}_{t+h}^{-1} y_{t+h} - y_{t+h}' b - b' y_{t+h} + b' \hat{\Sigma}_{t+h} b \right) + c + \frac{1}{2} b' \hat{\Sigma}_{t+h} b$. Substituting this expression in the expression for $h_t^*(y_{t+h})$ gives

$$h_t^*(y_{t+h}) = (2\pi)^{-\frac{m}{2}} \left| \hat{\Sigma}_{t+h} \right|^{-\frac{1}{2}} \exp \left(c + \frac{1}{2} b' \hat{\Sigma}_{t+h} b \right) \exp \left\{ -\frac{1}{2} \left(y_{t+h} - \hat{\Sigma}_{t+h} b \right)' \hat{\Sigma}_{t+h}^{-1} \left(y_{t+h} - \hat{\Sigma}_{t+h} b \right) \right\}. \quad (19)$$

Imposing the constraint $\int Jy_{t+h} h^*(y_{t+h}) dy_{t+h} = J\mu_{t+h}^*$ implies that $J\hat{\Sigma}_{t+h} b = J\mu_{t+h}^*$. Solving this last equation for ξ_t gives $\xi_t = \left(J\hat{\Sigma}_{t+h,11} J' \right)^{-1} (J\mu_{t+h}^* - J\hat{\mu}_{t+h})$. To obtain the expression

Table 1. Critical values for testing the condition in Proposition 1 ($k_{\delta,\alpha}$)

α	δ								
	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
0.05	3.176	2.938	2.770	2.624	2.475	2.352	2.248	2.080	1.975
0.10	2.928	2.676	2.482	2.334	2.168	2.030	1.904	1.740	1.600

for ζ_t , note that we must have that $c + \frac{1}{2}b'\hat{\Sigma}_{t+h}b = 0$. Solving this restriction for ζ_t yields, after tedious yet straightforward manipulations, that $\zeta_t = \frac{1}{2}(J\mu_{t+h}^* - J\hat{\mu}_{t+h})' \left(J\hat{\Sigma}_{t+h}J' \right)^{-1} (J\mu_{t+h}^* - J\hat{\mu}_{t+h})$. Noting that $(J\mu_{t+h}^* - J\hat{\mu}_{t+h}) = (J\mu_{t+h}^* - Jy_{t+h} - J\hat{\mu}_{t+h} + Jy_{t+h}) = (\hat{e}_{t+h,1:r} - \tilde{e}_{t+h,1:r})$ and that $\left(J\hat{\Sigma}_{t+h}J' \right)^{-1} = \left(\hat{\Sigma}_{t+h,11} \right)^{-1}$ gives the desired results. \square

The necessary and sufficient condition (17) can be empirically tested using a modification of Giacomini and Rossi (2010)'s fluctuation test, which accounts for the possibility that the expectation might be changing over time. For ease of exposition, in the following we omit the reference to the forecast horizon h , with the understanding that the size of the out-of-sample period T will be different for different forecast horizons. The test takes as primitives two sequences of out-of-sample forecast errors for the survey forecast and for the model-based forecast, \tilde{e}_t and \hat{e}_t for $t = 1, \dots, T$. A test of the null hypothesis that anchoring gives a more accurate forecasts according to the logarithmic score can be obtained by letting $\Delta L_t = S_t + N_t$ in Giacomini and Rossi (2010)'s Fluctuation test. The test is implemented by choosing a fraction δ of the total out-of-sample size T and computing a sequence of standardized rolling means of ΔL_t :

$$F_{t,\delta} = \hat{\sigma}^{-1}(\delta T)^{-1/2} \sum_{j=t-\delta T+1}^t \Delta L_t, \quad t = \delta T, \dots, T, \quad (20)$$

where $\hat{\sigma}$ is an HAC estimator of the standard deviation of ΔL_t computed over the rolling window, typically with truncation lag $h-1$, where h is the forecast horizon. The null hypothesis $H_0 : E[\Delta L_t] \geq 0$ is rejected when $\min_{t \leq T} F_{t,\delta} < -k_{\delta,\alpha}$, where the critical value $k_{\delta,\alpha}$ is given in Table 1.

The statistics is just the empirical version of the logarithmic scoring rule evaluated on a rolling window of observations. As such, its magnitude sheds light on the relative performance of the anchored model versus the baseline model: positive values of $F_{t,\delta}$ occur in windows where anchoring improves the accuracy of the density forecast.

4 Empirical results

4.1 The deteriorating performance of the DNS model and its variants

In this section, we document how the forecasting performance of the DNS model has deteriorated in the years after those considered by Diebold and Li (2006), who found that the model performed well in the sample from 1985-2000. This has been noted before, for example, by Diebold and Rudebusch (2012) and Mönch (2008). We complement their results by showing that augmenting the DNS model to incorporate information extracted from macroeconomic data.

For the yields-only model, we estimate the DNS models using the series of U.S. zero-coupon yields constructed in Gürkaynak et al. (2007).⁵ We consider average-of-the-month data from January 1985 to December 2011 on yields with the following maturities expressed in months: 6, 9, 12, 15, 18, 21, 24, 30, 36, 48, 60, 72, 84, 96, 108, 120. We augment the yield data with the monthly time series of the 3-month Treasury constant maturity rate from the FRED data set (code GS3M), which corresponds to the rate forecasted by the BC analysts.⁶ In total we have a panel of 324 monthly observations on 17 yields.

We estimate the DNS models using an out-of-sample recursive scheme and consider forecast horizons of 3-, 6-, 9- and 12-months ahead. The first estimation period uses data from 1985:1 to 1999:12, and we evaluate the forecasts over the out-of-sample period 2000:1 to 2011:12. We compare the mean squared forecast error (MSFE) of each variant of the DNS model to that of a random walk benchmark, which forecasts the yields as $\hat{\mu}_{t+h} = y_t$. The MSFE for the forecast of a yield of maturity τ at horizon h is given by:

$$MSFE_h(\tau) = \frac{1}{T} \sum_t (\hat{\mu}_{t+h}(\tau) - y_{t+h}(\tau))^2, \quad (21)$$

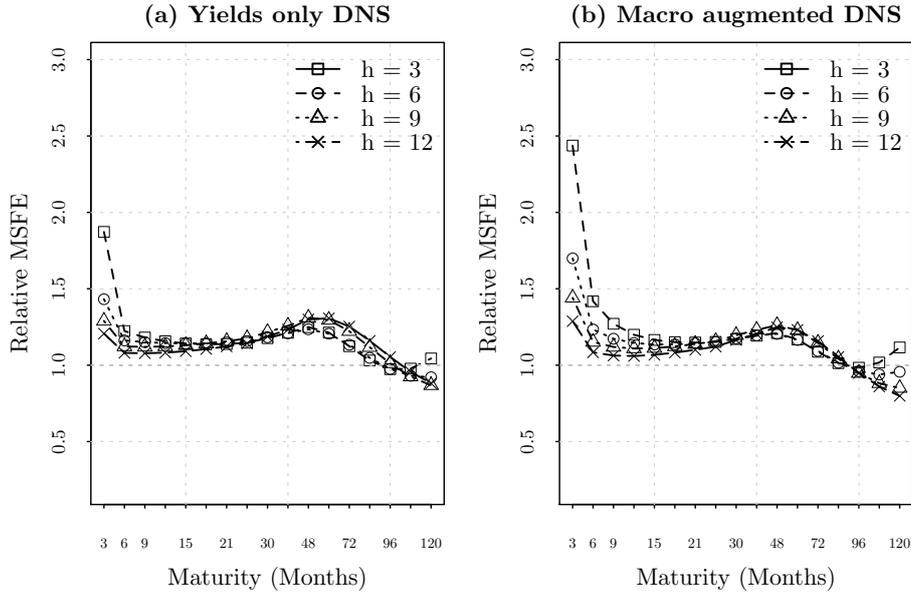
where T is the size of the out-of-sample portion of the sample, which in our case is $T = 144 - h$.

Figure 1 shows that the random walk outperforms both versions of the DNS model. This is generally true for all maturities and all forecast horizons, with a particularly poor performance for maturities around five years. The only exception appears to be the 10-year yield, for which

⁵A detailed description of the data is given in Appendix A. We also performed a similar exercise using the Fama-Bliss data (from CRSP), which are only available for one- to five-year maturities, and obtained similar conclusions, which we do not report in the paper.

⁶We also conducted the analysis using end-of-the-month data and the 3-month yield from the Gürkaynak et al. (2007) data set and obtained qualitatively similar results, which we do not report in the paper.

Figure 1. Relative MSFE for the yield only and macro augmented DNS against the random walk



Notes: The figure reports the ratios of the MSFE for the yield only and macro augmented DNS against the random walk for different maturities and forecast horizons. Values larger than 1 indicate that the random walk outperforms the model.

the model performs as well as the random walk except for the the three-month horizon. An implication is that in our sample incorporating macroeconomic information into the DNS model does not improve its performance. Actually, if anything, the performance of the richer model is poorer especially at shorter maturities.

We should point out that the poor out-of-sample performance of the DNS model in recent years stands in contrast to its good in-sample performance, which we document in the Appendix.

4.2 Testing the condition in Proposition 1

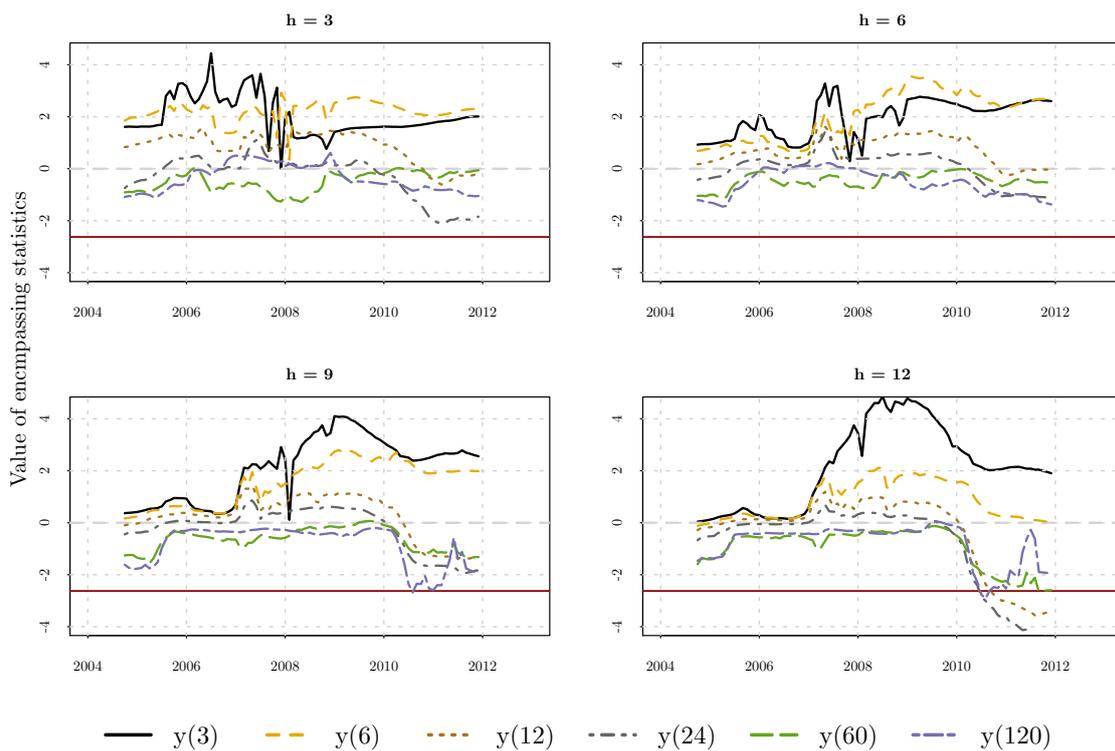
We test the condition in Proposition 1 in order to understand which survey forecasts would deliver the best improvements when used as anchoring points. The BC survey forecasts of yields are available for maturities of 3, 6, 12, 24, 60, and 120 months and for forecast horizons of 3, 6, 9, and 12 months.

Figure 2 reports the results for the test of the condition in Proposition 1, separately considering each maturity and each horizon. The null hypothesis is rejected when the sequence of test statistics crosses the horizontal (red) solid line, which represents the critical value. Although the null hypothesis is not rejected for several survey forecasts, the 3-month yield survey forecast

is the one offering the largest improvements, and it is also the yield for which the value of the test statistics is always positive at all horizons. This suggests that, if one were to anchor to a single survey forecast, the 3-month yield forecast would be the best in terms of density forecast performance. The 6-month yield survey forecast also delivers improvements in density forecast performance, but they are smaller than when using the 3-month. A natural question to ask is whether anchoring using both 3-month and 6-month survey forecasts would give additional improvements. In Figure 3, we thus report the test statistics for the test of the condition in Proposition 1 when anchoring using both survey forecasts. While the null hypothesis is again not rejected, the statistic is much smaller than the one using only the 3-month, suggesting that part of the information contained in the 6-month survey forecast overlaps with that of the 3-month, so that when they are used in combination there is practically no advantage.

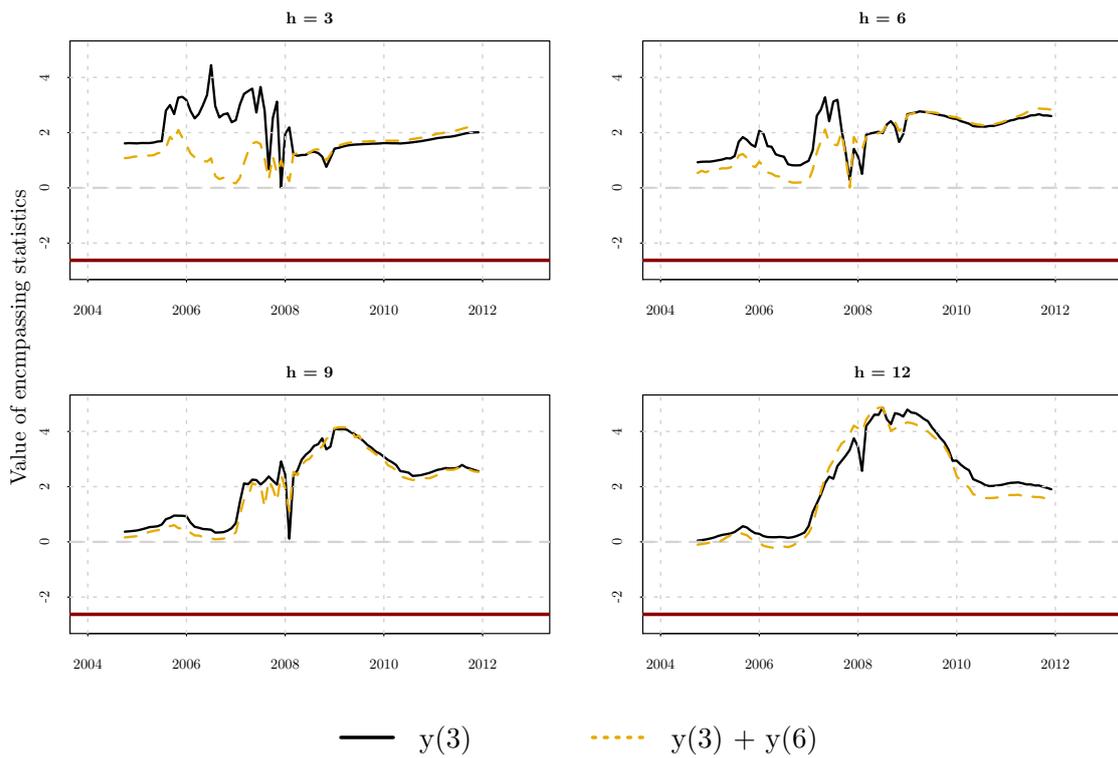
One of the possible explanations for why the survey forecast of the 3-month yield offers informational advantages over the model-based forecast is that this rate closely reacts to macroeconomic news, such as monetary policy decisions. This information gap between surveys and models is likely to be particularly large when the economic environment is changing quickly, making it more difficult for an econometric model to incorporate the new information. This conjecture is corroborated by considering how model-based and survey-based forecasts respond to monetary policy announcements that contain explicit reference to the likely future path of the short-term rate. There have been several instances of monetary policy statements containing forward-looking information of this kind in recent years, especially since the Federal Reserve began adopting forward guidance as a policy measure. Figure 4 below shows how the survey- and model-based forecasts reacted to one particular episode of forward guidance. The figure reports the 1- to 4-quarter-ahead forecasts of the 3-month yield given by the model and the surveys before and after the FOMC Statement of August 9, 2011, which stated that the “Committee currently anticipates that economic conditions [...] are likely to warrant exceptionally low levels for the federal funds rate at least through mid-2013”. The figure clearly shows that before the announcement both the model and the survey participants predicted a rate increase for the following year. However, after the announcement the surveys immediately incorporated the information about the policy decision to keep the rate fixed, whereas the model continued to predict a rate hike for several months afterwards, an increase that didn’t materialize. The

Figure 2. Testing the condition in Proposition 1



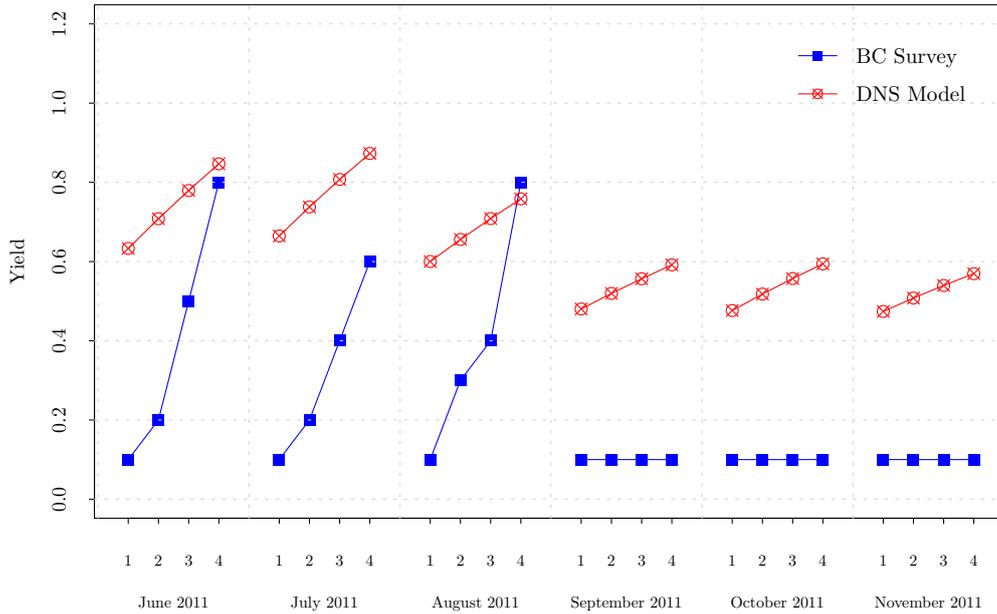
Notes: The figure reports the sequence of test statistics for testing the condition in Proposition 1, for a selection of maturities and horizons for which survey forecasts are available. The null hypothesis is rejected when the sequence of test statistics crosses the horizontal solid line, which represents the critical value (which equals -2.62 for test statistics computed over an estimation window that uses 40% of the out-of-sample observations and for a 5% significance level).

Figure 3. Testing the condition in Proposition 1



Notes: The figure reports the sequence of test statistics for testing the condition in Proposition 1, when anchoring using both the 3-month and the 6-month survey forecast. The null hypothesis is rejected when the sequence of test statistics crosses the horizontal solid line, which represents the critical value (which equals -2.62 for test statistics computed over an estimation window that uses 40% of the out-of-sample observations and for a 5% significance level).

Figure 4. The informational advantage of surveys over models



Note: The figure reports the 1- to 4-quarter-ahead forecasts of the 3-month yield given by the DNS model and the BC survey before and after the FOMC Statement of August 9, 2011.

ability to quickly incorporate this information gave the survey forecast a clear advantage, and this is likely to have occurred on several other occasions during the period that we considered, which was characterized by several episodes of forward guidance.

It is worth emphasizing that the informational advantage of the survey expectations is not due to a misalignment of the information sets on which the survey and the model forecasts are based. As we explain in Appendix B, we were careful in matching the timing of the two forecasts. Appendix B also explains how we transformed the quarterly BC forecasts into monthly forecasts.

4.3 Anchoring improves accuracy

In this section, we apply the anchoring method described in Section 3 to the DNS model using as anchor the 3-month yield BC forecast, which was the survey forecast offering the largest improvements in density forecast accuracy, according to the test of the condition in Proposition 1. Here we assess the out-of-sample point-forecast performance of the individual yield forecasts relative to the DNS forecasts and to the random walk benchmark.

Table 2 reports relative MSFE for the anchored forecasts against either the forecasts from the base DNS model or the random walk benchmark, for each maturity and forecast horizon. Table

Table 2. Relative MSFEs of anchored forecasts

Maturity	Anchored vs DNS				Anchored vs RW			
	h=3	h=6	h=9	h=12	h=3	h=6	h=9	h=12
3	0.42***	0.58*	0.70	0.83	0.79***	0.84*	0.90	1.00
6	0.58***	0.66***	0.74*	0.84	0.71***	0.77***	0.83*	0.91
9	0.66**	0.70**	0.77	0.87	0.78**	0.81**	0.86	0.93
12	0.72*	0.74*	0.79	0.89	0.84*	0.85*	0.89	0.96
15	0.77	0.77	0.81	0.91	0.88	0.88	0.92	0.99
18	0.80	0.80	0.83	0.93	0.90	0.91	0.95	1.02
21	0.81	0.82	0.85	0.94	0.93	0.94	0.99	1.06
24	0.82	0.84	0.87	0.96	0.95	0.97	1.02	1.09
30	0.84	0.86	0.90	0.98	0.98	1.02	1.09	1.17
36	0.84	0.89	0.92	1.01	1.02	1.07	1.16	1.24
48	0.86	0.92	0.96	1.04	1.07	1.13	1.26	1.35
60	0.88	0.95	0.99	1.07	1.07	1.14	1.28	1.40
72	0.92	0.98	1.02	1.10	1.03	1.11	1.25	1.37
84	0.97	1.02	1.06	1.12	1.00	1.07	1.18	1.30
96	1.04	1.06	1.10	1.15	1.01	1.04	1.11	1.21
108	1.10	1.11	1.14	1.17	1.08	1.03	1.05	1.12
120	1.15	1.15	1.17	1.19	1.20	1.06	1.02	1.05

Notes: The table reports the ratios of MSFE for the models considered. The asterisk indicates significance according to the Diebold and Mariano (1995) test of equal accuracy against the alternative that the anchored forecast is more accurate ('***' at the 1%, '**' at the 5%, and '*' at the 10%). The test was implemented using an HAC estimator with $h - 1$ as truncation parameter.

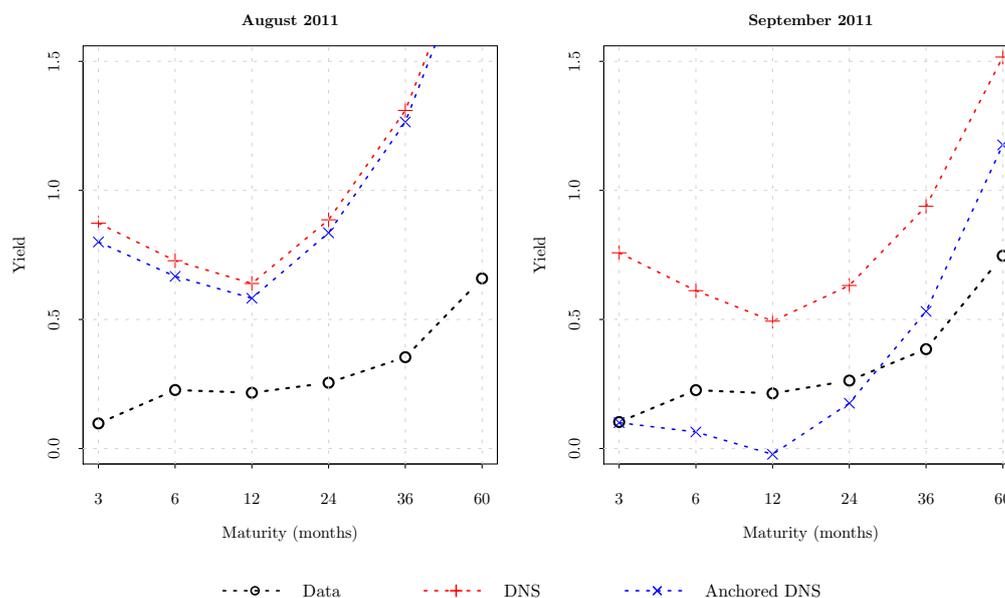
3 reports the same results for a restricted sample that ends in 2008 and thus excludes the zero lower bound period. The asterisks indicate that the Diebold and Mariano (1995) test rejects the null of equal forecast accuracy at 10% against the alternative that the anchored forecast is more accurate. The table shows that the anchored forecasts significantly and strongly outperform the DNS forecasts for almost all maturities and forecast horizons, with typical forecast accuracy gains of about 30% and up to 52%. The only exception is for a few long maturities and short forecast horizons, for which the anchored forecasts and the DNS forecasts perform equally well. The table also shows that the anchored forecasts outperform the random walk, and significantly for maturities up to 15 months and forecast horizons up to 6 months ahead, in a sample in which the DNS model and its variants consistently failed, as shown in Figure 1. These conclusions are robust to excluding the zero lower bound period, suggesting that the superior performance of our method is not solely driven by the ability of the survey expectations to incorporate the zero lower bound constraint.

Table 3. Relative MSFEs of anchored forecasts. Restricted sample: 1980:01-2008:12.

Maturity	Anchored vs DNS				Anchored vs RW			
	h=3	h=6	h=9	h=12	h=3	h=6	h=9	h=12
3	0.47***	0.62**	0.72	0.83	0.77***	0.82**	0.88	0.96
6	0.61***	0.68***	0.75*	0.84	0.70***	0.76***	0.81*	0.88
9	0.67***	0.72**	0.78	0.86	0.76***	0.80**	0.84	0.90
12	0.71**	0.75*	0.80	0.88	0.81**	0.83*	0.87	0.92
15	0.75*	0.78	0.82	0.90	0.84*	0.86	0.89	0.94
18	0.78	0.80	0.84	0.92	0.87	0.89	0.92	0.96
21	0.80	0.82	0.86	0.93	0.89	0.91	0.94	0.99
24	0.82	0.84	0.88	0.95	0.91	0.93	0.97	1.01
30	0.84	0.87	0.91	0.98	0.95	0.98	1.03	1.06
36	0.85	0.90	0.94	1.01	0.99	1.02	1.08	1.11
48	0.87	0.93	0.99	1.05	1.04	1.08	1.16	1.20
60	0.89	0.96	1.03	1.09	1.04	1.10	1.20	1.25
72	0.92	0.99	1.08	1.13	1.01	1.09	1.20	1.26
84	0.95	1.02	1.12	1.17	0.98	1.05	1.18	1.24
96	1.00	1.06	1.17	1.21	0.96	1.02	1.14	1.20
108	1.05	1.10	1.21	1.24	1.00	1.01	1.11	1.16
120	1.09	1.13	1.25	1.26	1.09	1.03	1.10	1.13

Notes: The table reports the ratios of MSFE for the models considered. The asterisk indicates significance according to the Diebold and Mariano (1995) test of equal accuracy against the alternative that the anchored forecast is more accurate ('***' at the 1%, '**' at the 5%, and '*' at the 10%). The test was implemented using an HAC estimator with $h - 1$ as truncation parameter.

Figure 5. DNS and Anchored forecasts before and after a monetary policy announcement



Notes: The figure shows the 12-month-ahead yield curve forecast implied by the DNS model and the corresponding anchored forecast made before and after the FOMC Statement of August 9, 2011, together with the actual yield curve realization.

Figure 5 reports the yield curve implied by the DNS and anchored forecast before and after the policy announcement of August 9, 2011, that was discussed in Figure 4. The figure shows that before the announcement the DNS and anchored yield curve forecasts similarly over-predicted the actual yield curve, whereas after the announcement the anchored forecast quickly incorporates the information contained in the FOMC statement, resulting in a sizable downward shift of the yield curve towards the actual realization. The DNS forecasts, instead, continue to largely over-predict the actual yield curve. This showcases the ability of the anchoring method to swiftly incorporate the informational advantage that surveys have about short yields and transmit it to the rest of the curve.

To understand whether the relative (point-forecast) performance of the anchored forecast and the base model changes over time, in Figure 5 we report rolling MSFEs for the two forecasts computed over rolling windows of four years, for a selection of maturities (rows) and of forecast horizons (columns). The figure shows that the anchored forecast is almost always more accurate, except for the very early years of the sample and for the 10-year maturity. Finally, note that

the superior performance of the anchored forecast relative to the baseline model is not confined to the performance of the point forecast, as one implication of Proposition 1 is that, when the testable condition holds, the entire density forecast of the anchored model is more accurate than the density forecast implied by the baseline model. The results of the fluctuation test reported in Figure 2, therefore, implicitly show that anchoring using the 3-month yield provides more accurate density forecasts over time than using the base model.

5 Conclusions

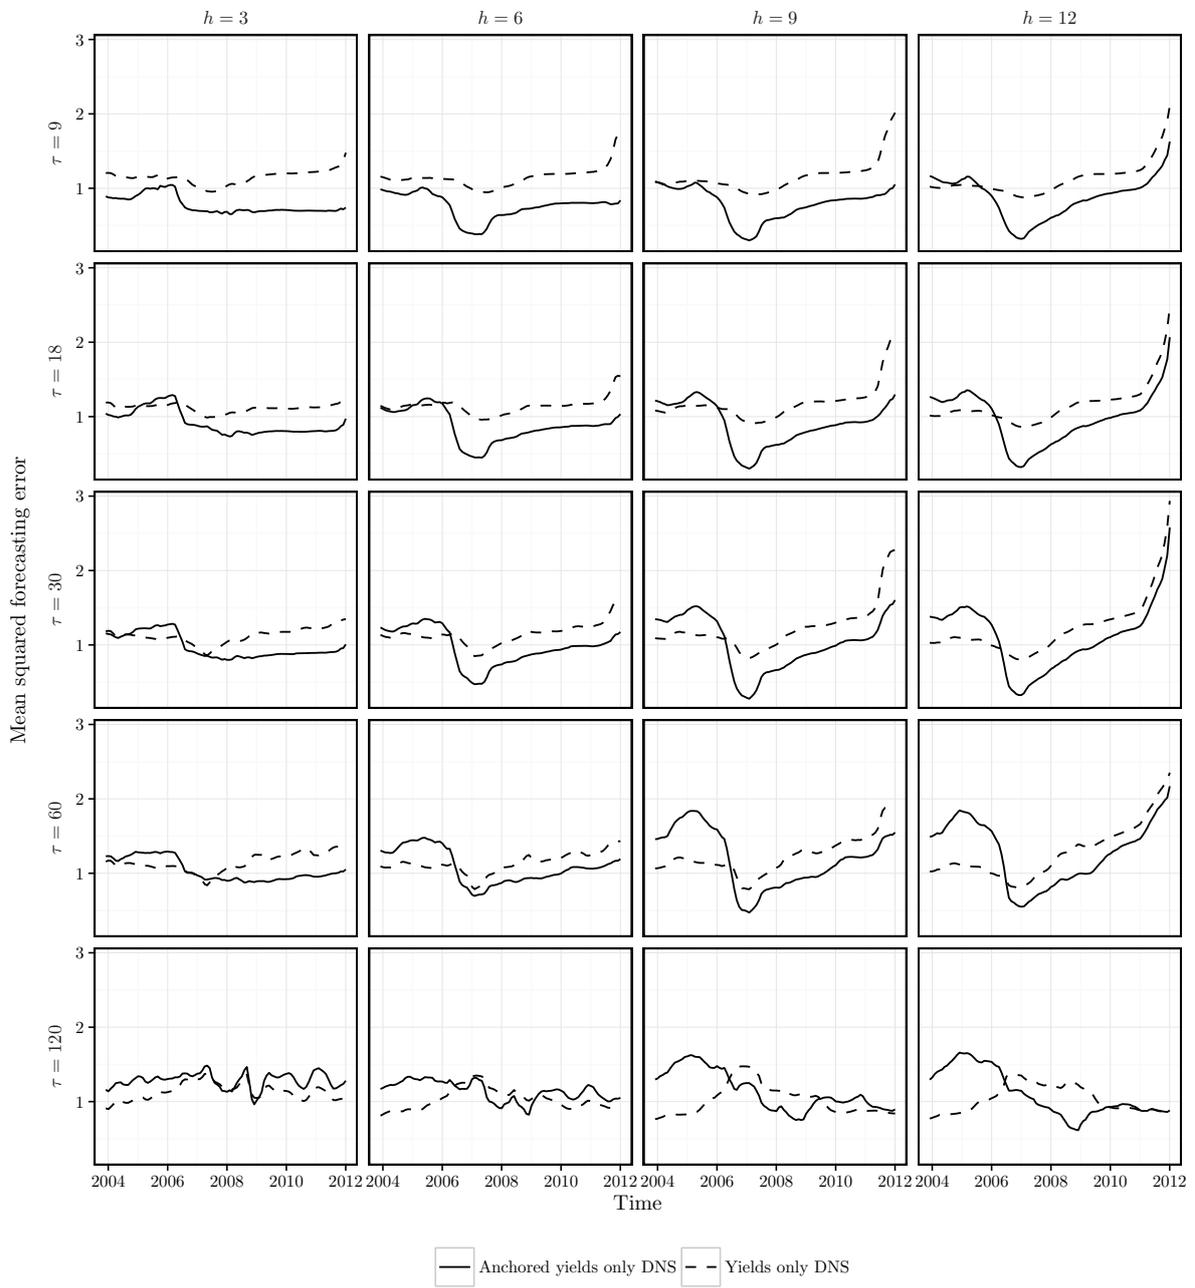
We proposed a formal and computationally simple anchoring method for incorporating survey expectations into a model-based forecast of the yield curve. The method constrains the dynamics of some yields to replicate those of the survey expectations and implicitly incorporates into the forecasts of the whole yield curve any information that survey participants use without the need to explicitly model it. The method is generally applicable to a situation where one has model-based forecasts for a set of variables and extra-model forecasts for a subset of these variables. The choice of the anchors is also general, for example they could come from another model (e.g., a random walk). All it matters for whether they will deliver (density-forecast) accuracy improvements is that they satisfy the condition of Proposition 1.

We applied the method to the Dynamic Nelson and Siegel model of Diebold and Li (2006) anchored to the 3-month yield forecast from the Blue Chip Financial Analysts survey and found large and significant improvements not only in density forecast accuracy relative to the base model, but also in out-of-sample point-forecast accuracy, with typical gains of about 30% and up to 52%. Remarkably, the anchored forecast also outperformed out-of-sample a random walk benchmark, at least for short maturities and forecast horizons.

We provide an interpretation for the accuracy gains of the anchored forecasts and relate them to their ability to capture information about real economic activity as well as forward-looking information contained in monetary policy announcements. This is likely to make the method even more relevant in the future, given that several central banks such as the Federal Reserve, the European Central Bank, and the Bank of England are increasingly adopting forward guidance as a nonstandard monetary policy measure.

Finally, our method offers a way to formally incorporate into yield curve forecasts “hidden”

Figure 6. Time-varying performance of anchored and DNS forecasts



Note: MSFE computed over a rolling four-year window

or “unspanned” factors that go beyond the information contained in the cross-section of yields, and suggests that any successful attempt to explicitly model the dynamics of yields should acknowledge the value of forward-looking information.

There are alternative ways to incorporate survey information into yield curve models that have been considered in the literature. Advantages of the method considered in this paper are that (i) it shows that using survey forecasts as anchors improves the overall density forecasting performance of model-based forecasts, and (ii) the selection of what survey information to use is based on a formal criterion. Advantages of alternative methods are that (i) the forecasting power of surveys does not have to be directly for a yield but may also be at the level of latent factors, and (ii) it is not necessary to use a 2-step approach of first selecting a survey forecast to anchor to, and then using this anchor in the forecast.

References

- ANG, A. AND M. PIAZZESI (2003): “A no-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables,” *Journal of Monetary economics*, 50, 745–787.
- BAÑBURA, M., D. GIANNONE, AND M. LENZA (2015): “Conditional forecasts and scenario analysis with vector autoregressions for large cross-sections,” *International Journal of Forecasting*, 31, 739–756.
- BIS (2005): “Zero-Coupon Yield Curves: Technical Documentation,” *BIS papers*, 25.
- CHERNOV, M. AND P. MUELLER (2012): “The term structure of inflation expectations,” *Journal of financial economics*.
- CHRISTENSEN, J. H. AND G. D. RUDEBUSCH (2015): “Estimating Shadow-Rate Term Structure Models with Near-Zero Yields,” *Journal of Financial Econometrics*, 13, 226–259.
- CHUN, A. L. (2011): “Expectations, bond yields, and monetary policy,” *Review of Financial Studies*, 24, 208–247.
- (2012): “Forecasting Interest Rates and Inflation: Blue Chip Clairvoyants or Econometrics?” in *EFA 2009 Bergen Meetings Paper*.
- CLARK, T. E. AND M. W. MCCrackEN (2001): “Tests of equal forecast accuracy and encompassing for nested models,” *Journal of econometrics*, 105, 85–110.
- CORONEO, L., D. GIANNONE, AND M. MODUGNO (2015): “Unspanned macroeconomic factors in the yield curve,” *Journal of Business & Economic Statistics*.
- CSISZÁR, I. (1975): “I-divergence geometry of probability distributions and minimization problems,” *The Annals of Probability*, 146–158.
- DAI, Q. AND K. J. SINGLETON (2000): “Specification analysis of affine term structure models,” *The Journal of Finance*, 55, 1943–1978.
- DIEBOLD, F. X. AND C. LI (2006): “Forecasting the term structure of government bond yields,” *Journal of econometrics*, 130, 337–364.

- DIEBOLD, F. X. AND R. S. MARIANO (1995): “Comparing Predictive Accuracy,” *Journal of Business & Economic Statistics*, 13, 253–265.
- DIEBOLD, F. X. AND G. D. RUDEBUSCH (2012): *Yield Curve Modeling and Forecasting: The Dynamic Nelson-Siegel Approach*, Princeton University Press.
- DIEBOLD, F. X., G. D. RUDEBUSCH, AND S. BORAGAN ARUOBA (2006): “The macroeconomy and the yield curve: a dynamic latent factor approach,” *Journal of econometrics*, 131, 309–338.
- DIJK, D., S. J. KOOPMAN, M. WEL, AND J. H. WRIGHT (2014): “Forecasting interest rates with shifting endpoints,” *Journal of Applied Econometrics*, 29, 693–712.
- DOZ, C., D. GIANNONE, AND L. REICHLIN (2012): “A Quasi-Maximum Likelihood Approach for Large, Approximate Dynamic Factor Models,” *Review of Economics & Statistics*, 94, 1014–1024.
- DUFFEE, G. R. (2011): “Information in (and not in) the term structure,” *Review of Financial Studies*, 24, 2895–2934.
- DUFFIE, D. AND R. KAN (1996): “A yield-factor model of interest rates,” *Mathematical finance*, 6, 379–406.
- DURBIN, J. AND S. J. KOOPMAN (2002): “A simple and efficient simulation smoother for state space time series analysis,” *Biometrika*, 89, 603–616.
- GIACOMINI, R. AND G. RAGUSA (2013): “Forecasting with judgment,” *mimeo*.
- GIACOMINI, R. AND B. ROSSI (2010): “Forecast comparisons in unstable environments,” *Journal of Applied Econometrics*, 25, 595–620.
- GIANNONE, D., L. REICHLIN, AND L. SALA (2005): “Monetary policy in real time,” in *NBER Macroeconomics Annual 2004, Volume 19*, MIT Press, 161–224.
- GÜRKAYNAK, R. S., B. SACK, AND J. H. WRIGHT (2007): “The US Treasury yield curve: 1961 to the present,” *Journal of Monetary Economics*, 54, 2291–2304.
- GÜRKAYNAK, R. S. AND J. H. WRIGHT (2012): “Macroeconomics and the term structure,” *Journal of Economic Literature*, 50, 331–367.

- HÖRDAHL, P., O. TRISTANI, AND D. VESTIN (2006): “A joint econometric model of macroeconomic and term-structure dynamics,” *Journal of Econometrics*, 131, 405–444.
- JOSLIN, S., M. PRIEBSCHE, AND K. SINGLETON (2010): “Risk premiums in dynamic term structure models with unspanned macro risks,” *Manuscript, Stanford University*.
- KIM, D. H. AND A. ORPHANIDES (2012): “Term structure estimation with survey data on interest rate forecasts,” *Journal of Financial and Quantitative Analysis*, 47, 241–272.
- KITAMURA, Y. AND M. STUTZER (1997): “An information-theoretic alternative to generalized method of moments estimation,” *Econometrica: Journal of the Econometric Society*, 861–874.
- LITTERMAN, R. B., J. SCHEINKMAN, AND L. WEISS (1991): “Volatility and the yield curve,” *The Journal of Fixed Income*, 1, 49–53.
- MÖNCH, E. (2008): “Forecasting the yield curve in a data-rich environment: A no-arbitrage factor-augmented VAR approach,” *Journal of Econometrics*, 146, 26–43.
- NEWKEY, W. K. AND R. J. SMITH (2004): “Higher order properties of GMM and generalized empirical likelihood estimators,” *Econometrica*, 72, 219–255.
- RAGUSA, G. (2011): “Minimum divergence, generalized empirical likelihoods, and higher order expansions,” *Econometric Reviews*, 30, 406–456.
- ROBERTSON, J. C., E. W. TALLMAN, AND C. H. WHITEMAN (2005): “Forecasting Using Relative Entropy,” *Journal of Money, Credit and Banking*, 37, 383–401.
- RUDEBUSCH, G. D. AND T. WU (2008): “A Macro-Finance Model of the Term Structure, Monetary Policy and the Economy*,” *The Economic Journal*, 118, 906–926.
- SARGENT, T. J., C. A. SIMS, ET AL. (1977): “Business cycle modeling without pretending to have too much a priori economic theory,” *New methods in business cycle research*, 1, 145–168.
- STARK, T. (2010): “Realistic evaluation of real-time forecasts in the Survey of Professional Forecasters,” *Federal Reserve Bank of Philadelphia Research Rap, Special Report*.
- WHITE, H. (1982): “Maximum Likelihood Estimation of Misspecified Models,” *Econometrica*, 50, 1–25.

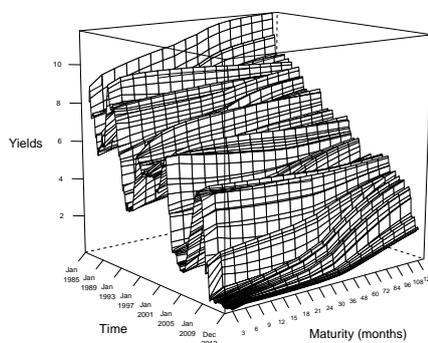
A Data

We describe the yields data, the macroeconomic data and the survey data that we used.

The yields data are pooled from two sources. The three-month yield is from the Fed's H-15 release. For longer maturities, we use zero-coupon yields constructed in Gürkaynak et al. (2007).⁷ We do not use the three-month yield from this dataset because the BC explicitly asks participants to predict this particular rate. We focus on average-of-the-month data from January 1985 to December 2011. We consider yields of the following 17 maturities (in months): 3, 6, 9, 12, 15, 18, 21, 24, 30, 36, 48, 60, 72, 84, 96, 108, 120. This choice provides us with a panel of 324 monthly observations on 17 different yields. Descriptive statistics of the sample are given in Table 4 and a plot of the data is in Figure 7.

The macroeconomic variables used in the macro augmented DNS model consist of monthly observations on 23 U.S. macroeconomic time series from 1985:1 through 2011:12. Table 5 lists the variables and the data transformations.

Figure 7. Bond yields data in three dimensions.



Notes: The figure plots average-of-the-month U.S. Treasury bill and bond yields at maturities ranging from 6 months to 10 years. The three-month yield is taken from the Fed's H-15 release. For longer maturities, we use zero-coupon yields constructed in Gürkaynak et al. (2007). The sample period is January 1985 through December 2012.

Survey data are from Blue Chip Financial Forecasts (BC). The interest rates that the

⁷This dataset is publicly available on the website of the Federal Reserve Board. The data can be obtained at the address: <http://www.federalreserve.gov/pubs/feds/2006/>.

Table 4. Yield data, 1985:1-2011:12. Descriptive statistics

<i>Maturity</i> (months)	<i>Mean</i>	<i>SD</i>	<i>Min</i>	<i>Max</i>	<i>MAE</i>	<i>RMSE</i>	$\hat{\rho}(1)$	$\hat{\rho}(12)$	$\hat{\rho}(30)$
3	-0.109	0.119	-0.772	0.208	0.085	0.119	0.78	0.268	-0.038
6	0.024	0.056	-0.097	0.378	0.037	0.056	0.67	0.098	-0.031
9	0.040	0.056	-0.139	0.335	0.038	0.055	0.75	0.238	-0.069
12	0.046	0.050	-0.129	0.272	0.036	0.050	0.79	0.293	-0.063
15	0.044	0.040	-0.081	0.200	0.030	0.040	0.82	0.330	-0.025
18	0.037	0.029	-0.034	0.123	0.022	0.029	0.85	0.359	0.061
21	0.027	0.020	-0.025	0.083	0.016	0.020	0.83	0.348	0.192
24	0.014	0.017	-0.051	0.052	0.013	0.017	0.74	0.242	0.165
30	-0.013	0.028	-0.145	0.087	0.018	0.028	0.75	0.236	-0.082
36	-0.037	0.040	-0.231	0.089	0.030	0.040	0.80	0.302	-0.061
48	-0.068	0.055	-0.297	0.064	0.042	0.054	0.85	0.371	0.048
60	-0.073	0.053	-0.253	0.067	0.040	0.053	0.87	0.397	0.131
72	-0.057	0.040	-0.173	0.049	0.030	0.040	0.87	0.402	0.191
84	-0.027	0.020	-0.081	0.023	0.015	0.020	0.84	0.371	0.223
96	0.010	0.015	-0.027	0.098	0.011	0.015	0.81	0.325	-0.046
108	0.050	0.037	-0.047	0.190	0.029	0.037	0.86	0.387	0.103
120	0.089	0.063	-0.074	0.271	0.048	0.063	0.86	0.389	0.141

Notes: The data on yields used in the paper are pooled from two sources. The end-of-month three-month yield is taken from the Fed's H-15 release. For longer maturities, we use zero-coupon yields constructed in Gürkaynak et al. (2007).

Table 5. Macroeconomic variables

	Variable	Acronym	Tran
1	Personal Income	PI	$\Delta \ln$
2	University of Michigan Inflation Expectation	UMIE	Δlv
3	Producer Price Index: All Commodities	PPI	$\Delta 2 \ln$
4	Consumer Price Index for All Urban Consumers: All Items	CPI	$\Delta 2 \ln$
5	Personal Consumption Expenditures: Chain-type Price Index	PCE	$\Delta 2 \ln$
6	All Employees: Total non-farm	Emp	$\Delta \ln$
7	4-Week Moving Average of Initial Claims	IC	$\Delta \ln$
8	Moody's Seasoned Aaa Corporate Bond Yield	AAA	Δlv
9	Moody's Seasoned Baa Corporate Bond Yield	BAA	Δlv
10	Industrial Production Index	IP	$\Delta \ln$
11	Capacity Utilization: Total Industry	CU	Δlv
12	Civilian Labor Force	LF	$\Delta \ln$
13	Civilian Unemployment Rate	UR	Δlv
14	Average Weekly Hours of Production and Nonsupervisory Employees: Manufacturing	AWH	lv
15	Housing Starts: Total: New Privately Owned Housing Units Started	HS	\ln
16	ISM Manufacturing: PMI Composite Index	PMI	lv
17	M1 Money Stock	M1	$\Delta 2 \ln$
18	M2 Money Stock	M2	$\Delta 2 \ln$
19	Total Consumer Credit Owned and Securitized, Outstanding	CC	$\Delta 2 \ln$
20	S&P 500 Stock Price Index	SP500	$\Delta \ln$
21	Average Hourly Earnings of Production and Nonsupervisory Employees: Total Private	AHE	$\Delta 2 \ln$
22	Excess Reserves of Depository Institutions	RDI	$\Delta 2 \ln$
23	ISM Manufacturing: Employment Index	EI	lv

BC analysts forecast are the quarterly average of constant maturity Treasury yields as defined by the Federal Reserve Statistical Release H.15.

There are two important issues to consider when comparing the accuracy of survey data to that of model-based forecasts: the alignment of information sets and the conversion of quarterly forecasts into monthly forecasts. Here we explain how we dealt with both issues.

First, in order to guarantee a fair comparison between model and survey forecasts, we made sure that the information sets on which the forecasts are based are aligned. In particular, one must pay attention to the timing of the survey in relation to the type of data used. For example, if the model is estimated using end-of-the-month data and the survey is released around the 15th of the month (which is the typical release date for the quarterly Survey of Professional Forecasters conducted by the Federal Reserve Bank of Philadelphia), the informational advantage of the survey might be simply related to the fact that when panelists submit their forecasts they have access to more information than the model. This is however less of a concern when using BC survey data, because the survey is published on the 1st day of each month, and one can thus make the assumption that the survey forecasts have access to the same information set than a model estimated using end-of-the-month data (or possibly a smaller information set if the survey participants communicate their forecasts a few days before the BC releases them).

The second issue is related to the discrepancy between the frequency of the survey (monthly) and the frequency of the target variables' predictions (quarterly). In fact, this discrepancy introduces a time-variation in the information set that has to be taken into account when analyzing the accuracy of the survey predictions at a monthly frequency. Table 6 describes how we extracted monthly forecasts from the quarterly BC survey expectations. Recall that the BC analysts are asked to predict the average value of a target variable over the current and the following quarters. For this reason, we use average-of-the-month data (although using end-of-the-month data does not significantly alter our results). In practice, we use the expectation made for the current quarter as the 3-months-ahead prediction. This means that for the survey released on the 1st of January the implied 3-month-ahead forecast is given by the expected value for the current

Table 6. Extraction of monthly forecasts from quarterly BC forecasts

Time of forecast release	Forecast horizon ($h =$ months)			
	$h = 3$	$h = 6$	$h = 9$	$h = 12$
January 1 st	BC_{Jan}^{Now}	BC_{Jan}^{1Q}	BC_{Jan}^{2Q}	BC_{Jan}^{3Q}
February 1 st	BC_{Feb}^{1Q}	BC_{Feb}^{2Q}	BC_{Feb}^{3Q}	BC_{Feb}^{4Q}
March 1 st	BC_{Mar}^{1Q}	BC_{Mar}^{2Q}	BC_{Mar}^{3Q}	BC_{Mar}^{4Q}
April 1 st	BC_{Apr}^{Now}	BC_{Apr}^{1Q}	BC_{Apr}^{2Q}	BC_{Apr}^{3Q}

quarter, i.e. the nowcast, contained in the survey (BC_{Now} , Jan in the table). Similarly, the 6- 9- and 12-month-ahead forecasts made on the 1st of January are retrieved from the 1-, 2-, and 3-quarter-ahead forecasts contained in the survey ($BC_{Jan,1Q}$, $BC_{Jan,2Q}$ $BC_{Jan,3Q}$ in the table).

B Conditional forecasting

To obtain the conditional forecasts of Table 7 we follow Bańbura et al. (2015) which in turn adopt the strategy for forecasting with ragged edge data sets using a Kalman filter methodology. Here the ragged edges arise as the survey forecasts are used as future observations. The methodology relies on drawing the states conditionally on the observations and the survey forecasts using the simulation smoother of Durbin and Koopman (2002). More precisely, given the parameter estimates, the simulation smoother draws

$$(\beta_{t|T}, \theta_{t|T}) \sim p(\cdot | y_t, x_t, bc_{T|T+h}, \bar{\Psi}, t = 1, \dots, T), \quad (22)$$

where bc_{T+h} denotes the Blue Chip forecast at horizon h given the information set at time T , and $\bar{\Psi}$ denote the estimated parameters adjusted to take into account the ragged edges that result from introducing the survey forecasts in the state space. The conditional forecast for the yield can then be produced by evaluating the measurement equation at different simulations of $(\beta_{t|T}, \theta_{t|T})$, and then approximating the expected value of the yield by averaging over the resulting draws.

Table 7. Relative MSFEs of conditional forecasts (yield only model)

Maturity	Conditional vs DNS				Conditional vs RW			
	h=3	h=6	h=9	h=12	h=3	h=6	h=9	h=12
3	0.42***	0.58*	0.70	0.83	0.79***	0.84*	0.90	1.00
6	0.56***	0.65**	0.73*	0.82	0.68***	0.75**	0.82*	0.89
9	0.65**	0.70*	0.76	0.85	0.77**	0.81*	0.85	0.91
12	0.72*	0.75	0.79	0.87	0.83*	0.85	0.89	0.94
15	0.77	0.78	0.82	0.89	0.88	0.89	0.92	0.97
18	0.80	0.81	0.84	0.91	0.91	0.93	0.96	1.01
21	0.83	0.84	0.86	0.93	0.94	0.96	1.00	1.04
24	0.84	0.86	0.88	0.94	0.96	0.99	1.03	1.08
30	0.85	0.89	0.91	0.97	1.00	1.05	1.11	1.15
36	0.86	0.91	0.94	0.99	1.04	1.10	1.18	1.22
48	0.87	0.95	0.98	1.03	1.09	1.17	1.28	1.33
60	0.91	0.98	1.01	1.05	1.10	1.18	1.31	1.38
72	0.95	1.01	1.05	1.08	1.07	1.15	1.28	1.35
84	1.01	1.05	1.09	1.10	1.04	1.11	1.22	1.28
96	1.07	1.10	1.13	1.13	1.04	1.08	1.14	1.19
108	1.14	1.15	1.17	1.15	1.12	1.07	1.08	1.10
120	1.19	1.19	1.22	1.18	1.24	1.10	1.06	1.04

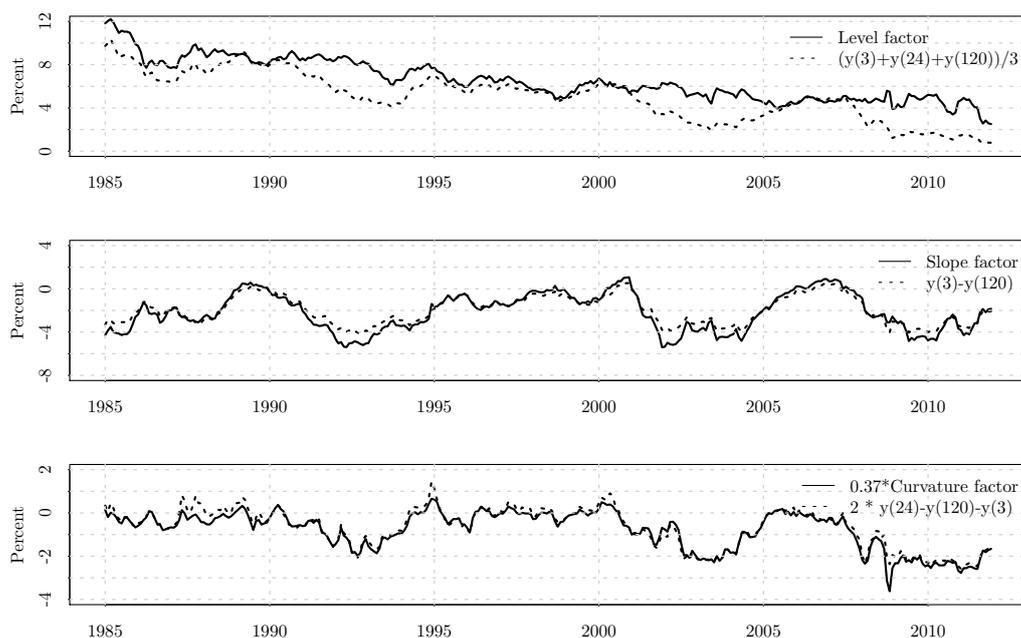
Notes: The table reports the ratios of MSFE of the conditional forecasts to the random walk forecasts and to the yields only DNS model. The asterisk indicates significance according to the Diebold and Mariano (1995) test of equal accuracy against the alternative that the anchored forecast is more accurate ('***' at the 1%, '**' at the 5%, and '*' at the 10%). The Diebold and Mariano test was implemented using an HAC estimator with $h - 1$ truncation parameter.

C In sample fit of DNS model

Here we report results regarding the estimation and the in-sample fit of the state-space version of the DNS model. For brevity we only report the yields only model. Fit statistics for the macro model are qualitatively similar.

Figure 8 shows the estimated time series of the three factors and their empirical counterparts. The first graph plots the level factor $\hat{\beta}_{1,t}$ against the average of short-, medium- and long-term yields, $(y_t(3) + y_t(24) + y_t(120))/3$. The middle panel plots $\hat{\beta}_{2,t}$ against the empirical slope of the yield curve $y_t(3) - y_t(120)$. Finally, the bottom panel shows the behavior of $0.37\hat{\beta}_{3,t}$ and the empirical curvature proxy $2y_t(24) - y_t(3) - y_t(120)$. The curvature factor closely matches the dynamics of its empirical counterpart: the difference between the two series has a mean of 12 basis points and a standard deviation of 19 basis points. Also, the slope factor matches very closely the empirical proxy for the slope with a correlation of 0.98. The level factor shows

Figure 8. Estimated DNS factors and empirical counterparts



Note: The first factor β_{1t} controls the yield curve level, as it can be verified that $\lim_{\tau \rightarrow \infty} y_t(\tau) = \beta_{1t}$. The second factor β_{2t} is related to the yield curve slope, defined as the difference between the 10-year and three-month yields. The third factor β_{3t} governs the curvature of the yield curve, defined as twice the two-year yield minus the sum of the 10-year and three-month yields.

instead a marked departure from the empirical counterpart. Importantly, this departure is most noticeable in the period 2000-2011. In particular, from January 1985 to December 2001, the correlation between $\hat{\beta}_{1,t}$ and $(y_t(3) + y_t(24) + y_t(120))/3$ is 0.81, from February 2002 to December 2011 the correlation drops to 0.29. The mean and standard deviations of the difference increase from 119 basis points to 196 basis points and from 93 basis points to 124 basis points, respectively.

Table 8 presents summary statistics of the residuals for each of the 17 maturities considered. Both the mean and the standard deviations of the residuals are small for all maturities. The largest means are observed at the shortest and longest maturities, respectively 21 and 24 basis points, respectively. For the other maturities, the residual means oscillate between 1.4 and 9 basis points. Similarly, the residuals' standard deviation is fairly stable, with largest values at the beginning and at the end of the curve.

Table 8. In-sample fit of the DNS model

Maturity (months)	Mean	SD	Min	Max	MAE	RMSE	$\hat{\rho}(1)$	$\hat{\rho}(12)$	$\hat{\rho}(30)$
3	-2.1e-01	2.4e-01	-1.6e+00	4.7e-01	1.0e-01	2.4e-01	0.86	0.29	-0.0385
6	-4.9e-02	6.3e-02	-2.7e-01	2.2e-01	6.4e-03	6.3e-02	0.86	0.29	-0.0448
9	-1.7e-02	1.9e-02	-7.7e-02	5.7e-02	6.8e-04	1.9e-02	0.87	0.33	-0.0377
12	-4.3e-08	4.8e-08	-2.2e-07	1.3e-07	4.1e-15	4.7e-08	0.44	0.25	0.1146
15	6.7e-03	6.2e-03	-1.2e-02	2.9e-02	8.4e-05	6.2e-03	0.89	0.39	-0.0064
18	5.9e-03	5.2e-03	-7.7e-03	2.5e-02	6.1e-05	5.2e-03	0.90	0.41	0.0126
21	1.8e-09	6.0e-08	-2.8e-07	2.8e-07	3.6e-15	6.0e-08	0.41	0.11	0.0080
24	-8.9e-03	7.3e-03	-3.6e-02	5.8e-03	1.3e-04	7.3e-03	0.91	0.45	0.0493
30	-3.0e-02	2.4e-02	-1.2e-01	2.0e-02	1.5e-03	2.4e-02	0.91	0.47	0.0803
36	-5.0e-02	3.8e-02	-1.9e-01	3.5e-02	3.9e-03	3.8e-02	0.92	0.48	0.1054
48	-7.1e-02	5.3e-02	-2.6e-01	5.4e-02	7.9e-03	5.3e-02	0.92	0.49	0.1425
60	-6.6e-02	4.8e-02	-2.3e-01	5.3e-02	6.7e-03	4.8e-02	0.92	0.49	0.1689
72	-3.9e-02	2.8e-02	-1.3e-01	3.3e-02	2.4e-03	2.8e-02	0.92	0.48	0.1891
84	1.0e-08	2.0e-08	-3.2e-08	1.6e-07	5.1e-16	2.0e-08	0.32	0.19	-0.0160
96	4.6e-02	3.2e-02	-4.0e-02	1.4e-01	3.1e-03	3.2e-02	0.92	0.46	0.2182
108	9.3e-02	6.4e-02	-8.1e-02	3.0e-01	1.3e-02	6.4e-02	0.91	0.45	0.2282
120	1.4e-01	9.4e-02	-1.2e-01	4.4e-01	2.8e-02	9.4e-02	0.91	0.44	0.2351

Note: In-sample fit statistic of the two step DNS model. The descriptive statistics refer to the residuals of the equation for yields at the corresponding maturities. The last three columns present residual sample autocorrelations at lag 1, 12, and 30, respectively.