Anchoring the Yield Curve Using Survey Expectations*

Carlo Altavilla[†]
European Central Bank

Raffaella Giacomini[‡]
University College London
CeMMAP
CEPR

Giuseppe Ragusa[§]
Luiss University, Rome

Thursday 17th October, 2013

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 are accurate predictors of yields, but only for very short maturities. We argue that this is partly due to the ability of survey participants to incorporate information about the current state of the economy as well as forward-looking information such as that contained in monetary policy announcements. We show how the informational advantage of survey expectations about short yields can be exploited to improve the accuracy of yield curve forecasts given by a base model. We do so by employing a flexible projection method that anchors the model forecasts to the survey expectations in segments of the yield curve where the informational advantage exists and transmits the superior forecasting ability to all remaining yields. The method implicitly incorporates into yield curve forecasts any information that survey participants have access to, without the need to explicitly model it. We document that anchoring delivers large and significant gains in forecast accuracy for the whole yield curve, with improvements of up to 52% over the years 2000-2012 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

---

*We would like to thank Domenico Giannone, Carlo Favero, Alejandro Justiniano, Michele Modugno, Kevin Sheppard, Ken 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, Autonoma Barcelona, Chicago 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.

[†]E-mail: carlo.altavilla@ecb.int

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

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

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. The two broad classes of yield curve models 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), which are extracted from the cross-section of yields. Broadly speaking, their differences lie 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), and Coroneo et al. (2013)) or latent "hidden" factors (Joslin et al. (2010) and Duffee (2011)).

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) and 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); Van Dijk et al. (2012) use 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.

In contrast to the existing approaches in the literature, we do not incorporate survey data into the model, as we show that this makes very little difference to the model's performance. Instead, we employ a formal "anchoring" method that anchors segments of the yield curve forecasts to the corresponding survey expectations about yields and transmits the superior forecasting ability to the rest of the 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, the survey expectation offers the possibility to capture both observable and "hidden" factors that can explain yield curve dynamics (as also argued by Duffee (2011)). The survey expectation 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 several points along the yield curve. Moreover, survey expectations are not necessarily accurate, so it is desirable to shed some light on the link between the accuracy of the survey expectations and that of the resulting anchored forecast. Our main result is to show that the anchoring procedure results in an improvement in accuracy for the whole yield curve if the survey expectations one utilizes are informationally efficient relative to the model-based forecasts they replace, which in practice corresponds to a testable encompassing condition. In our data, we found that the informational efficiency condition is satisfied only for the 3-month yield, so in practice we suggest anchoring the short end of the yield curve to the corresponding survey expectation and then adjusting all remaining yields in a formal way, which we make explicit below. As a quick visualization of the effects of anchoring, consider Figure 4, which shows that the method shifts an existing yield curve forecast toward the actual realization, with 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)). Here we establish a link between the presence of an informational advantage of the surveys over model-based forecasts and the accuracy of the anchored forecast.

We conduct a thorough empirical evaluation of the out-of-sample forecasting performance of the anchoring method, which incorporates 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 base model of the yield curve.

We find that the anchoring procedure results in forecasts that uniformly and significantly outperform those produced by several versions of the DNS model, including ones that explicitly incorporate macroeconomic factors or survey forecasts. 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 2000-2012.

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 documents the informational advantage of surveys over variants of the DNS model. Section 3 describes the anchoring method. Section 4 contains the empirical results and Section 5 concludes. Appendix A describes the yield and macroeconomic data; Appendix B reports the in-sample estimation results of the DNS model; and Appendix C discusses the BC survey data.

2 The informational advantage of surveys over models

We first introduce the DNS model and variants of the model that incorporate the information contained in macroeconomic factors or in survey data. We document that none of these models were able to outperform the random walk in recent years. We then show that survey expectations of yields have an information advantage over the model-based forecasts, but only for the very short yield. We conclude by linking the informational advantage of surveys over models to the ability of survey participants to capture forward-looking information such as that contained in monetary policy announcements.

2.1 The DNS model and its variants

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

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

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 calibrated parameter governing the exponential decay rate of the coefficients. As in Diebold and Li (2006), we let $\lambda = 0.069$.

We consider two different specifications for the law of motion of the factors. In the first, the vector of factors β_{t+h}, where h is the forecast horizon, follows the process

$$\beta_{t+h} = C + \Gamma\beta_t + \eta_{t+h}, \tag{2}$$

where C a 3×1 vector of constant, Γ is assumed to be diagonal and $\eta_{t+h} \sim N(0, S)$ with elements independent of each other and S diagonal. Although we do not report the results here, we also considered a non diagonal specification for Γ, but we found that it made little difference to the conclusions.

In the second specification, the evolution of the factors depends on additional observable information X_t:

$$\beta_{t+h} = C + \Gamma \beta_t + \Lambda X_t + \eta_{t+h}. \tag{3}$$

We consider three variants: 1) $X_t = (f_t^{(real)}, f_t^{(nominal)})$ where $f_t^{(real)}$and $f_t^{(nominal)}$ are the first two principal components extracted from the 23 macroeconomic variables listed in Appendix A. We denote them $f_t^{(real)}$and $f_t^{(nominal)}$ based on the fact that the first principal component has a high correlation with real variables (e.g., correlation 0.75 with Industrial Production) and the second has a high correlation with nominal variables; 2) X_t equals the consensus h-step-ahead forecast of inflation from the BC survey; and 3) X_t equals the consensus h-step-ahead forecast of the three-month yield from the BC survey.

Estimation of ((1)) proceeds in two stages. In the first stage, the cross-section of yields is used to estimate β_{1t}, β_{2t}, and β_{3t} at each time period t using ordinary least squares. The outcome of this first step is thus a times series of estimated factors $\hat{\beta}_t = (\hat{\beta}_{1t}, \hat{\beta}_{2t}, \hat{\beta}_{3t})$. In the second stage, the parameters of equation ((2)) (or, alternatively equation ((3))) are estimated regressing each element of $\hat{\beta}_t$ on each element of $\hat{\beta}_{t-h}$ and a constant.

The h-step-ahead conditional mean forecast of the yields at time t, $\hat{\mu}_{t+h}$, is obtained as:

$$\hat{\mu}_{t+h} = Z\hat{\beta}_{t+h}, \quad Z = \begin{pmatrix} 1 & \frac{1-e^{-\lambda\tau_1}}{\lambda\tau_1} & \frac{1-e^{-\lambda\tau_1}}{\lambda\tau_1} - e^{-\lambda\tau_1} \\ 1 & \frac{1-e^{-\lambda\tau_2}}{\lambda\tau_2} & \frac{1-e^{-\lambda\tau_2}}{\lambda\tau_2} - e^{-\lambda\tau_2} \\ \vdots & & \\ 1 & \frac{1-e^{-\lambda\tau_m}}{\lambda\tau_m} & \frac{1-e^{-\lambda\tau_m}}{\lambda\tau_m} - e^{-\lambda\tau_m} \end{pmatrix}. \tag{4}$$

where $\hat{\beta}_{t+h} = \hat{C} + \hat{\Gamma}\hat{\beta}_t$, or, alternatively $\hat{\beta}_{t+h} = \hat{C} + \hat{\Gamma}\hat{\beta}_t + \hat{\Lambda}X_t$. To derive the density forecast of y_{t+h}, which is needed for the anchoring procedure, we assume that the pricing errors are independent over t and are normally distributed:

$$u_t \equiv \begin{bmatrix} u_t(\tau_1) \\ \vdots \\ u_t(\tau_m) \end{bmatrix} \sim N(0, Q), \qquad Q = E[u_t u_t']$$

Under this assumption and under the specifications for β_t given in ((2)) or ((3)), y_{t+h} is conditionally normally distributed

$$y_{t+h} : \left\{ f_t(y_{t+h}) \sim \mathscr{N}(Z\hat{\Gamma}\hat{\beta}_t, \hat{\Sigma}), \quad \hat{\Sigma} = Z\hat{S}Z' + \hat{Q} \right\}. \quad t = 1, \ldots, T.$$

In practice, we recursively estimate Σ from the residuals of ((1)) and from the residuals of ((2)) or ((3)).

2.2 The forecasting 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 or surveys does not solve the problem.

We estimate the DNS models using the series of U.S. zero-coupon yields constructed in Gürkaynak et al. (2007).[1] We consider average-of-the-month data from January 1985 to December 2012 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.[2] In total we have a panel of 324 monthly observations on 17 yields.

We estimate the DNS model and its variants 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 2012:12. 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:

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

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 substantially outperforms all 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 the model performs as well as the random walk at the three-month horizon. This means that incorporating macroeconomic or survey-based information directly into the model does not improve its performance.

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 Appendix B.

2.3 Surveys win at short maturities

As discussed in the introduction, a well-known fact in the forecasting literature is that survey participants, such as those participating in the BC survey that we consider in this paper, often

[1] A 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.

[2] We 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 of DNS variants against the random walk

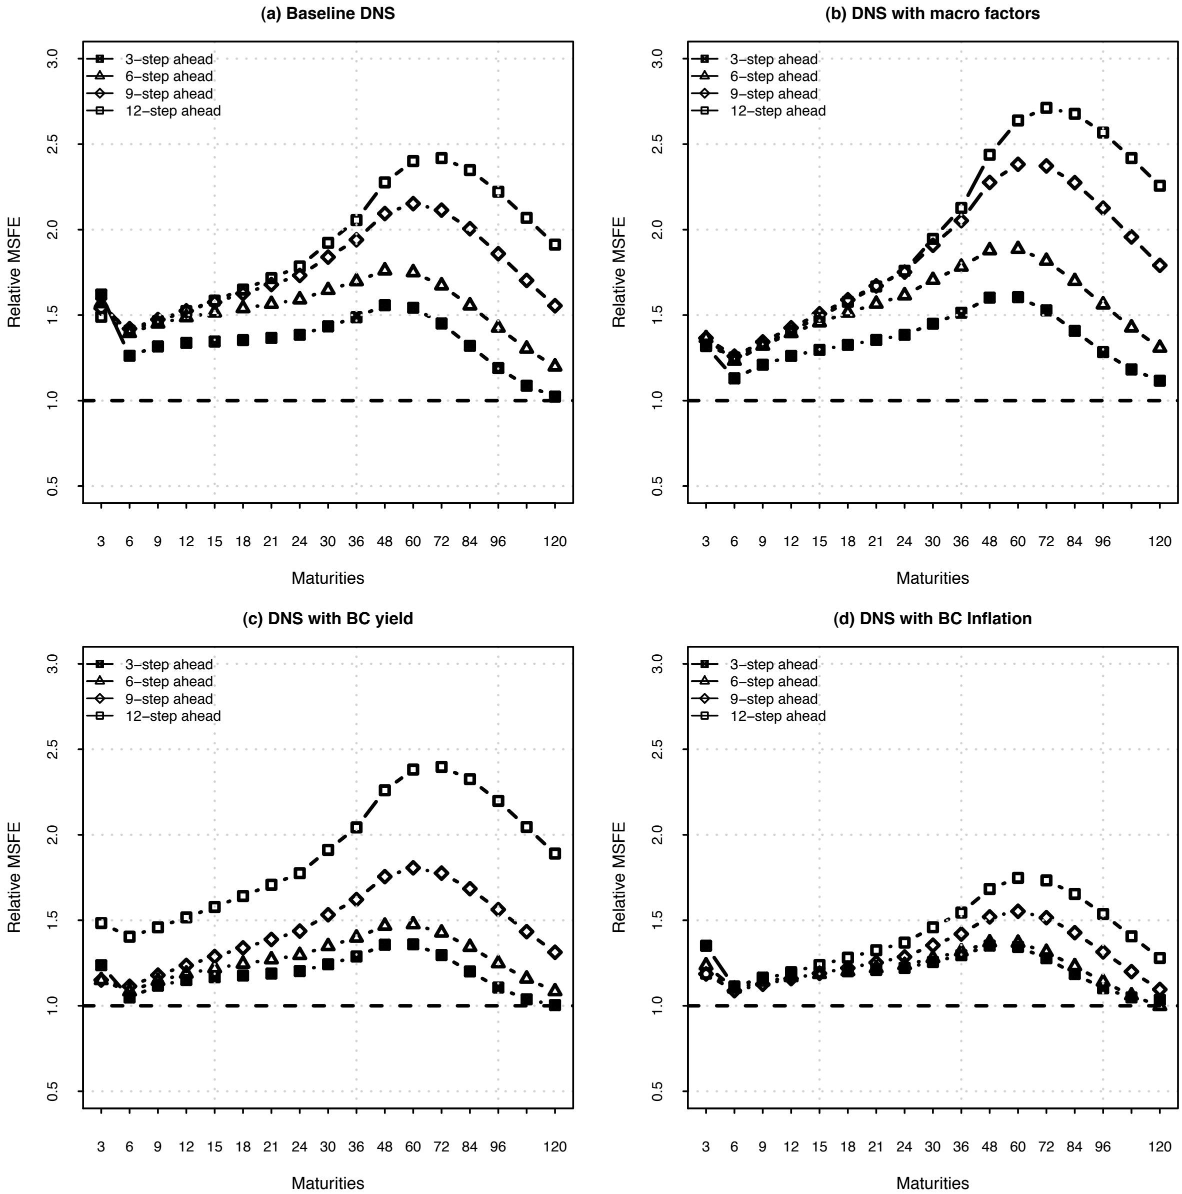

Notes: The figure reports the ratios of the MSFE for each variant of the DNS model against the MSFE of the random walk for different maturities and forecast horizons. Values larger than 1 indicate that the random walk outperforms the model.

Figure 2. Relative performance of BC vs. DNS forecasts over time

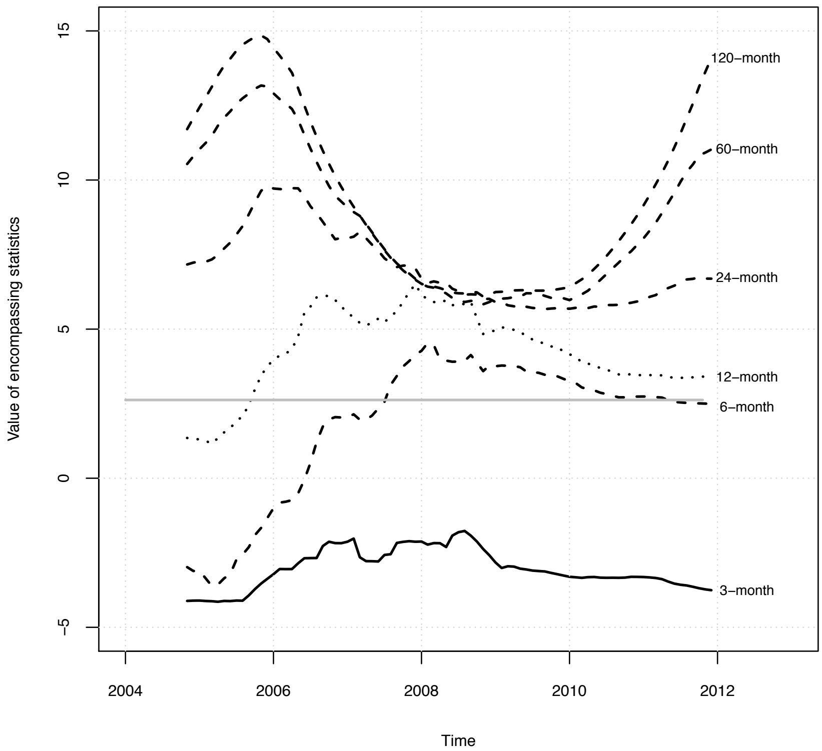

Notes: The figure reports the sequence of test statistics for the time-varying encompassing test described in Section 3.2, testing the null hypothesis that the BC forecast encompasses the DNS forecast, against the alternative hypothesis that it does not. The null hypothesis is rejected when the sequence 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).

produce more accurate forecasts than those based on econometric models. Here we focus in particular on the BC consensus forecasts of yields, which are available for maturities of 3, 6, 12, 24, 60, and 120 months and forecast horizons of 3, 6, 9, and 12 months.

Figure 2 reports the relative forecast performance of the BC forecasts and the DNS forecasts, assessed using the encompassing test described in detail in Section 3.2. The figure shows the sequence of test statistics computed over time testing the null hypothesis that the 3-month-ahead BC forecasts for maturities 3, 6, 12, 24, 60 and 120 months encompasses the corresponding DNS forecast. The null hypothesis is rejected when the sequence of test statistics crosses the horizontal solid line, which represents the critical value. The test clearly rejects the null of encompassing for all but the 3-month maturity, suggesting that the BC forecasts are informationally efficient relative to the DNS forecast only at the very short end of the yield curve.

One of the possible explanations for why the survey forecast of the 3-month yield has strongly and consistently outperformed the model's forecasts since 2000 is that this rate closely reacts to macroeconomic news. 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. In particular, the fact that the informational efficiency of the BC forecasts relative to the model-based forecast is limited to the 3-month yield could be due to the fact that this is the rate that more closely reacts to monetary policy decisions. Indeed, the 3-month Treasury bill rate is usually used as a proxy for the monetary policy rate in many macroeconomic models.

This conjecture is corroborated by looking at 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 3 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 ability to quickly incorporate this information gave the survey forecast a clear accuracy gain, 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.

In closing this section, it is important to emphasize 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.

3 Anchoring the yield curve to survey expectations

The previous section showed that survey participants can have an informational advantage over model-based forecasts, but this is only true for the very short end of the yield curve. This means that, on one hand, survey expectations alone cannot be used to produce accurate forecasts of the entire yield curve and, on the other hand, that model-based forecasts cannot be entirely discarded. In this section, we illustrate an anchoring method for incorporating the information contained in the survey expectations of the short yield into an existing model-based forecast of the yield curve.

Figure 3. The informational advantage of surveys over models

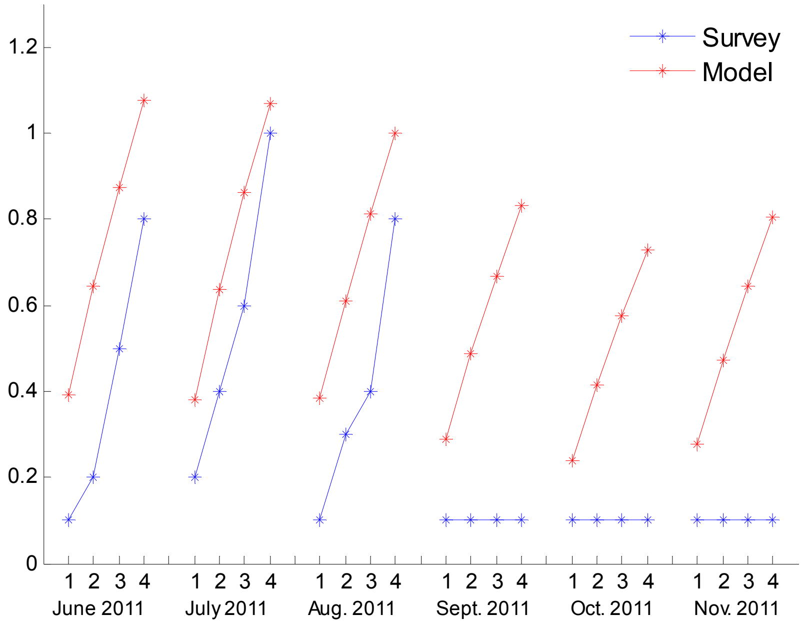

Note: The figure reports the 1- to 4-quarter-ahead forecasts of the 3-month yield given by the DNS model and the BC survey before and after the FOMC Statement of August 9, 2011.

3.1 The anchoring method

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]

$$y_{t+h} : \left\{ f_t(y_{t+h}) \sim \mathscr{N}(\hat{\mu}_{t+h}, \hat{\Sigma}_{t+h}) \right\}. \quad t = 1, \ldots, T.$$

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

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 $\tau_1, \ldots, \tau_r$. More formally, this set of densities can be characterized as

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

It is important to note that no constraints are imposed on the forecasts of yields at longer maturities, $\tau_{r+1}, \ldots, \tau_m$. The idea is to select the density in $\widetilde{\mathscr{H}}_{t+h}$ that is "closest" to the model-based density forecast f_t, where closeness is measured by the Kullback-Leibler information criterion. We seek a solution to the following minimization problem

$$h_t^*(y_{t+h}) = \arg \min_{h \in \widetilde{\mathscr{H}}_{t+h}} \int \log\left(\frac{h_t(u)}{f_t(u)}\right) h_t(u) du. \tag{5}$$

Minimization problems such as ((5)) play an important role in statistics and econometrics (Csiszár (1975); **?**; 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 taking the form

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

where ζ_t and ξ_t are parameters chosen in such a way that $h_t^*(y_{t+h}) \in \widetilde{\mathscr{H}}_{t+h}$. 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 $h_t^*(y_{t+h})$:

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

[4] In the interest of notational clarity, we consider only the case in which the survey forecasts considered are for maturities $\tau_1, \tau_2, \ldots, \tau_r$. It is, however, immediate to extend the results of this section to cases in which survey forecasts of noncontiguous maturities are considered.

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

with

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

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

$$\hat{\Sigma}_{t+h} = \begin{pmatrix} \underset{r\times r}{\hat{\Sigma}_{t+h,11}} & \underset{r\times(m-r)}{\hat{\Sigma}_{t+h,12}} \\ \underset{(m-r)\times r}{\hat{\Sigma}_{t+h,21}} & \underset{(m-r)\times(m-r)}{\hat{\Sigma}_{t+h,22}} \end{pmatrix}.$$

Thus, the solution to ((5)) 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 $\hat{\Sigma}_{t+h}$. Forecasts of yields at different maturities are generally positively correlated. This implies that when the model forecast is larger than that of the survey, that is, when $\mu_{t+h,1:r} - \tilde{\mu}_{t+h,1:r} > 0$, $\mu_{t+h,r+1:N}$ is adjusted downwards; on the other hand, when the model forecast is smaller than the survey, $\mu_{t+h,r+1:N}$ is increased.

3.2 Where to anchor the yield curve?

A natural question to ask is which survey data one should use to anchor the yield curve, given that the survey expectations are in principle available for a number of yields. In this section, we provide guidance on where to anchor the yield curve by showing that, if the survey expectation of a given yield is informationally efficient relative to the corresponding model-based forecast, using this expectation to anchor the yield curve delivers an improvement in forecast accuracy for the yield curve as a whole. The notion of informational efficiency is equated here to forecast encompassing, which conveniently lends itself to developing a testable condition that can be used to decide which survey data to use for anchoring.

Formally, we have the following result:

Proposition 1. *Let $\hat{e}_{t+h}(\tau)$ and $\tilde{e}_{t+h}(\tau)$ denote the model- and the survey-based h-step-ahead forecast errors for a yield with maturity τ, respectively. If the survey expectation for the yield of maturity τ encompasses the model-based forecast of the same yield, that is, if*

$$E\left[(\hat{e}_{t+h}(\tau) - \tilde{e}_{t+h}(\tau))\tilde{e}_{t+h}(\tau)\right] \leq 0, \tag{6}$$

then the anchored density forecast $h_t^*(y_{t+h})$ is more accurate than the base forecast $f_t(y_{t+h})$, according to the logarithmic scoring rule of Amisano and Giacomini (2007), i.e.,

$$E\left[\log\left(\frac{h_t^*(y_{t+h})}{f_t(y_{t+h})}\right)\right] > 0. \tag{7}$$

Proof. Since

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

it is sufficient to show that the expectations of both terms are positive. In particular, we show that $\xi_t = \hat{\Sigma}^{-1}_{t+h,11}(\tilde{\mu}_{t+h}(\tau) - \hat{\mu}_{t+h}(\tau))$ and $\zeta_t = \frac{1}{2}\hat{\Sigma}^{-1}_{t+h,11}(\tilde{\mu}_{t+h}(\tau) - \hat{\mu}_{t+h}(\tau))^2$, from which it follows that $E\left[\zeta_t\right] \geq 0$, and that $E\left[\xi_t \tilde{e}_{t+h}(\tau)\right] = E\left[\hat{\Sigma}^{-1}_{t+h,11}(\tilde{e}_{t+h}(\tau) - \hat{e}_{t+h}(\tau))\tilde{e}_{t+h}(\tau)\right] \geq 0$ if condition ((6)) is satisfied.

Analytical expressions for ξ_t and ζ_t can be obtained by completing the square, as follows. First write $h_t^*(y_{t+h}) = \exp\left\{\zeta_t + \xi_t'\left[y_{t+h}(\tau) - \tilde{\mu}_{t+h}(\tau)\right]\right\} f_t(y_{t+h})$ as

$$\begin{aligned} h_t^*(y_{t+h}) = (2\pi)^{-\frac{m}{2}}&\left|\hat{\Sigma}_{t+h}\right|^{-\frac{1}{2}} \\ &\times \exp\left\{-\frac{1}{2}(y_{t+h} - \hat{\mu}_{t+h})'\hat{\Sigma}^{-1}_{t+h}(y_{t+h} - \hat{\mu}_{t+h}) + \zeta_t + \xi_t'\left[Jy_{t+h} - \tilde{\mu}_{t+h}(\tau)\right]\right\}, \end{aligned}$$

where J is a selection vector selecting the element of y_{t+h} corresponding to maturity τ. We have

$$-\frac{1}{2}(y_{t+h} - \hat{\mu}_{t+h})'\hat{\Sigma}^{-1}_{t+h}(y_{t+h} - \hat{\mu}_{t+h}) + \zeta_t + \xi_t'\left[Jy_{t+h} - \tilde{\mu}_{t+h}(\tau)\right] = y_{t+h}'Ay_{t+h} + y_{t+h}'b + c$$

where $A = -\frac{1}{2}\hat{\Sigma}^{-1}_{t+h}$ and $b = \hat{\Sigma}^{-1}_{t+h}\hat{\mu}_{t+h} + J'\xi_t$ and $c = -\frac{1}{2}\hat{\mu}_{t+h}'\hat{\Sigma}^{-1}_{t+h}\hat{\mu}_{t+h} - \xi_t'\tilde{\mu}_{t+h}(\tau) + \zeta_t$. We can thus write $y_{t+h}'Ay_{t+h} + y_{t+h}'b + c = \left(y_{t+h} + \frac{1}{2}A^{-1}b\right)' A \left(y_{t+h} + \frac{1}{2}A^{-1}b\right) + k$ with $k = c - \frac{1}{4}b'A^{-1}b$, which gives

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

Imposing the constraint $E\left[Jy_{t+h}\right] = \tilde{\mu}_{t+h}(\tau)$ implies that $-\frac{1}{2}JA^{-1}b = \tilde{\mu}_{t+h}(\tau)$, which in turn gives $\xi_t = \hat{\Sigma}^{-1}_{t+h,11}(\tilde{\mu}_{t+h}(\tau) - \hat{\mu}_{t+h}(\tau))$. To obtain the expression for ζ_t, note that we must have that $k = 0$ and thus set $c = \frac{1}{4}b'A^{-1}b$ and solve for ζ_t to obtain, after a few straightforward manipulations, $\zeta_t = \frac{1}{2}\hat{\Sigma}^{-1}_{t+h,11}(\tilde{\mu}_{t+h}(\tau) - \hat{\mu}_{t+h}(\tau))^2$. This completes the proof. □

Condition ((6)) can be empirically tested using a modification of the Giacomini and Rossi (2010) fluctuation test, which accounts for the possibility that the expectation might be changing over time.[5]

For ease of exposition, in the following we omit the reference to the forecast horizon h, with

[5] Note that, even though Giacomini and Rossi (2010) restrict attention to a rolling window forecasting scheme to avoid the complications that arise when conducting pairwise comparisons of forecast accuracy in the context of estimated nested models, the fact that one of our forecasts here is model-free prevents the need to limit attention to the rolling scheme.

Table 1. Critical values for the encompassing test ($k_{\delta,\alpha}$)

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

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, $\tilde{e}_t(\tau)$ and $\hat{e}_t(\tau)$ for $t = 1, ..., T$. A test of the null hypothesis (6) of encompassing, $H_0 : E\left[(\hat{e}_{t+h}(\tau) - \tilde{e}_{t+h}(\tau))\tilde{e}_{t+h}(\tau)\right] \leq 0$ against the one-sided alternative that the survey forecast does not encompass the model forecast can be obtained by letting $\Delta L_t = (\hat{e}_{t+h}(\tau) - \tilde{e}_{t+h}(\tau))\tilde{e}_{t+h}(\tau)$ 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:

$$F_{t,\delta} = \hat{\sigma}^{-1}(\delta T)^{-1/2} \sum_{j=t-\delta T+1}^{t} \left[(\hat{e}_j(\tau) - \tilde{e}_j(\tau))\tilde{e}_j(\tau)\right], \quad t = \delta T, ..., T,$$

where $\hat{\sigma}$ is an HAC estimator of the standard deviation of $(\hat{e}_{t+h}(\tau) - \tilde{e}_{t+h}(\tau))\tilde{e}_{t+h}(\tau)$ computed over the rolling window, typically with truncation lag $h - 1$, where h is the forecast horizon. The null hypothesis is rejected when

$$\max_{t \leq T} F_{t,\delta} > k_{\delta,\alpha},$$

where the critical value $k_{\delta,\alpha}$ is given in Table 1.

4 Empirical results

In this section, we apply the anchoring method described in Section 3 to the DNS model and the 3-month yield BC forecast, which was the only forecast to satisfy the informational efficiency condition discussed in Section 3.2. Our goal is to assess the out-of-sample performance of the individual yield forecasts, relative to the DNS forecasts and to the random walk benchmark. Recall that the anchored forecast for the whole vector of yields for forecast horizon h is given by

$$\mu^*_{t+h} = \begin{pmatrix} \tilde{\mu}_{t+h,1} \\ \hat{\mu}_{t+h,2:m} - \hat{\Sigma}_{t+h,21}\left(\hat{\Sigma}_{t+h,11}\right)^{-1}(\hat{\mu}_{t+h,1} - \tilde{\mu}_{t+h,1}) \end{pmatrix},$$

where $\tilde{\mu}_{t+h,1}$ is the 3-month yield BC forecast, $\hat{\mu}_{t+h,2:m}$ the vector of DNS forecasts for yields with maturity 6,...,120 months, $\hat{\Sigma}_{t+h,11}$ and $\hat{\Sigma}_{t+h,21}$ are blocks of the partitioned variance matrix

$\hat{\Sigma}_{t+h}$:

$$\hat{\Sigma}_{t+h} = \begin{pmatrix} \underset{1\times 1}{\hat{\Sigma}_{t+h,11}} & \underset{1\times 16}{\hat{\Sigma}_{t+h,12}} \\ \underset{16\times 1}{\hat{\Sigma}_{t+h,21}} & \underset{16\times 16}{\hat{\Sigma}_{t+h,22}} \end{pmatrix}.$$

4.1 Anchoring works

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. 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 clearly shows that the anchored forecasts significantly and strongly outperform the DNS forecasts for almost all maturities and forecast horizons, with a typical forecast accuracy gain of about 30% and up to 52%. The only exception is for a few long maturities and short forecast horizons, at which the anchored forecasts and the DNS forecasts perform equally well. The table also gives evidence that the anchored method was able to 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. For space-saving reasons, we do not report the results of the comparison between the anchored forecasts and the variants of the DNS forecasts which incorporate macroeconomic information or survey expectations, but they paint a very similar picture and are available upon request.

Figure 4 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 3. The figure shows that, whereas before the announcement the DNS and anchored yield curve forecasts similarly overpredicted the actual yield curve, after the announcement the anchored forecast quickly incorporates the information contained in the FOMC statement and shifts the entire yield curve downwards towards the actual realization with a sizable adjustment relative to the previous month. The DNS forecasts, instead, continue to largely overpredict 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.

4.2 Why does anchoring work?

In this section, we provide some insight into the possible reasons for the superior performance of the anchored forecast relative to the baseline DNS model. Our goal is to understand the nature of the possible additional information contained in the anchored forecast. In order to do so, we proceed as follows. We consider the factors extracted by principal components from the 3-month-ahead anchored yield curve forecasts and the corresponding factors associated with the 3-month-ahead forecasts from the base DNS model. Note that, by construction, only three factors can be extracted from the DNS forecasts because of the structure of the model. For the anchored forecasts, we find that the first three factors are essentially identical to those for the

Table 2. Relative MSFEs of anchored forecasts

	Anchored vs DNS				Anchored vs RW			
Maturity	h=3	h=6	h=9	h=12	h=3	h=6	h=9	h=12
3	0.48*	0.53*	0.58*	0.67*	0.78*	0.83	0.90	1.00
6	0.49*	0.52*	0.57*	0.66*	0.62*	0.72*	0.81	0.93
9	0.54*	0.54*	0.59*	0.67*	0.71*	0.79*	0.87	0.98
12	0.58*	0.57*	0.60*	0.68*	0.78*	0.84	0.92	1.04
15	0.62*	0.59*	0.62*	0.69*	0.83*	0.89	0.97	1.09
18	0.64*	0.61*	0.63*	0.70*	0.87	0.94	1.02	1.15
21	0.66*	0.63*	0.64*	0.70*	0.91	0.98	1.07	1.20
24	0.68*	0.64*	0.65*	0.71*	0.94	1.02	1.12	1.26
30	0.70*	0.67*	0.67*	0.71*	1.01	1.10	1.22	1.37
36	0.72*	0.69*	0.68*	0.72*	1.07	1.18	1.32	1.48
48	0.75*	0.73*	0.70*	0.73*	1.16	1.28	1.47	1.66
60	0.77*	0.75*	0.73*	0.74*	1.19	1.32	1.56	1.78
72	0.80*	0.78*	0.74*	0.75*	1.17	1.31	1.57	1.81
84	0.84*	0.81*	0.76*	0.76*	1.12	1.26	1.53	1.77
96	0.89	0.84*	0.78*	0.76*	1.06	1.20	1.45	1.69
108	0.95	0.87*	0.80*	0.77*	1.03	1.13	1.36	1.59
120	1.00	0.90	0.82*	0.78*	1.02	1.08	1.27	1.48

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

Figure 4. DNS and Anchored forecasts before and after the monetary policy announcement

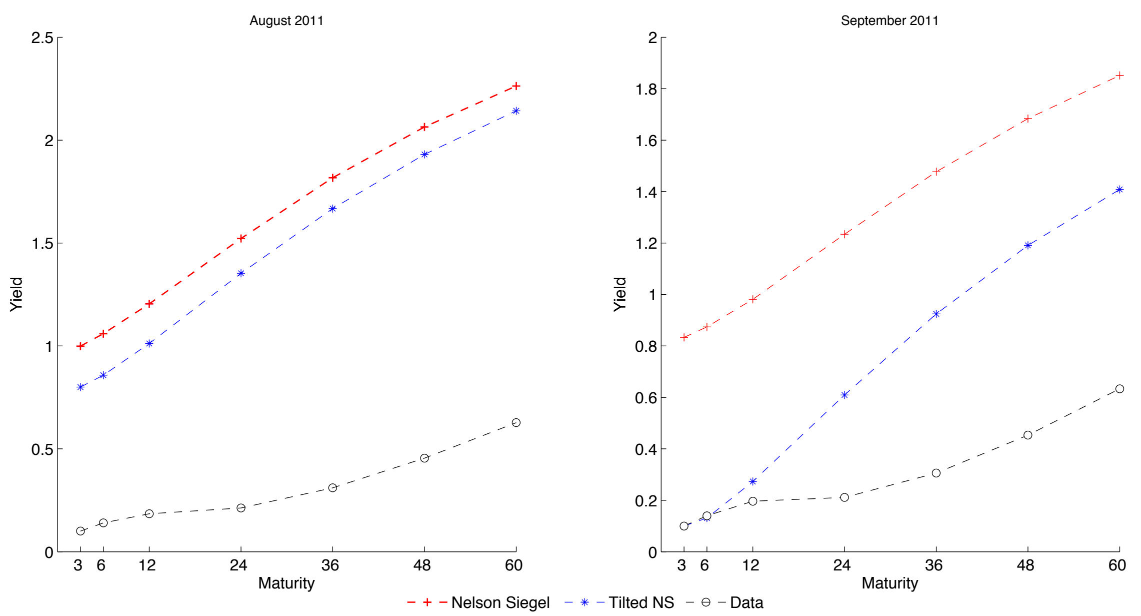

Notes: The figure shows the 3-month-ahead yield curve forecast implied by the DNS model and the corresponding anchored forecast made before and after the FOMC Statement of August 9, 2011, together with the actual yield curve realization.

Figure 5. Factors extracted from DNS and anchored forecasts

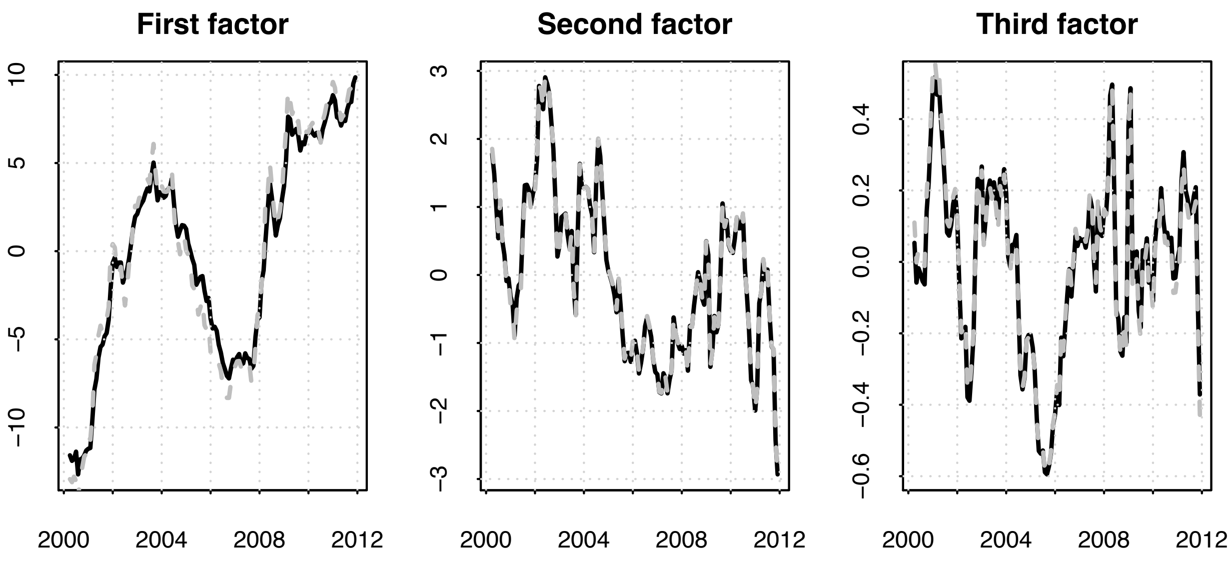

Note: The figure shows the first 3 factors extracted from the DNS (solid line) and anchored forecasts (dashed line).

DNS model, as can be seen from Figure 5.

We then investigate whether the fourth factor extracted from the anchored forecasts, ${f_t}^{(4)}$, captures the additional information that is embedded into the forecasts by the anchoring procedure and that is not already contained in the cross-section of yields. We do so by relating the factor to two of the possible sources of informational advantage of survey forecasts over model-based forecasts that we discussed in Section 2: the access of survey participants to information about the state of the economy and to forward-looking information such as that contained in monetary policy announcements. To measure the information about the state of the economy, we consider the same two factors that we utilized in Section 2.2 to augment the DNS model, $f_t^{(real)}$and $f_t^{(nominal)}$. We then construct an index of forward-looking information in monetary policy announcements, $I_t^{(forward)}$, which equals one if in the month before the release of the survey forecast there was an FOMC statement that contained forward-looking information, which we assess by putting together Table 4 of Gürkaynak et al. (2007) for 2000:1 to 2004:12 and Table 1 of Campbell et al. (2012) for 2007:2011. For the years 2005 and 2006, we follow the conclusion by Kool and Thornton (2012) that there was no forward guidance during this period and let the index equal zero. We then estimate the following regression (with t-statistics within parentheses):

$$ {f_t}^{(4)} = \underset{(-0.5)}{-0.05} - \underset{(-5.7)}{0.37} f_t^{(real)} - \underset{(-0.4)}{0.03} f_t^{(nominal)} + \underset{(2.4)}{0.41} I_t^{(forward)} + error. $$

These estimates confirm that the superior accuracy of the anchored forecasts is related to their ability to incorporate information about real economic activity and forward-looking information contained in monetary policy announcements.

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. We applied the method to the Dynamic Nelson and Siegel model of Diebold and Li (2006) because of its popularity in financial and policy institutions, but we stress that the method could be applied to the forecasts from any other base model. The method also offers a way to establish which information to incorporate, and we found grounds for using only the expectations about the three-month yield from the Blue Chip Financial Forecasts survey.

The results are stark. We find large and significant improvements in out-of-sample accuracy across maturities and forecast horizons, with typical accuracy gains of about 30% and up to 52% relative to the base model. To the best of our knowledge, our forecast is the only one that was able to outperform a random walk benchmark over the period 2000-2012, 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 now adopting forward guidance as a nonstandard monetary policy measure.

Finally, our method offers a way to formally incorporate into yield curve forecasts "hidden" 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.

Appendix

A Appendix. Yield and macroeconomic data

In this Appendix, we describe the data on the yields and the macroeconomic data that we used in the paper.

The data on the 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 Gurkaynak et al. (2007). [6] 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 2012. 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 3 and a plot of the data is in Figure 6.

The macroeconomic factors that appear in (3) are the first two principal components extracted from a dataset of 23 variables. The dataset consists of monthly observations on 23 U.S. macroeconomic time series from 1985:1 through 2011:12. Table 4 lists the variables and the data transformations that we applied to them.

Figure 6. Bond yields data in three dimensions.

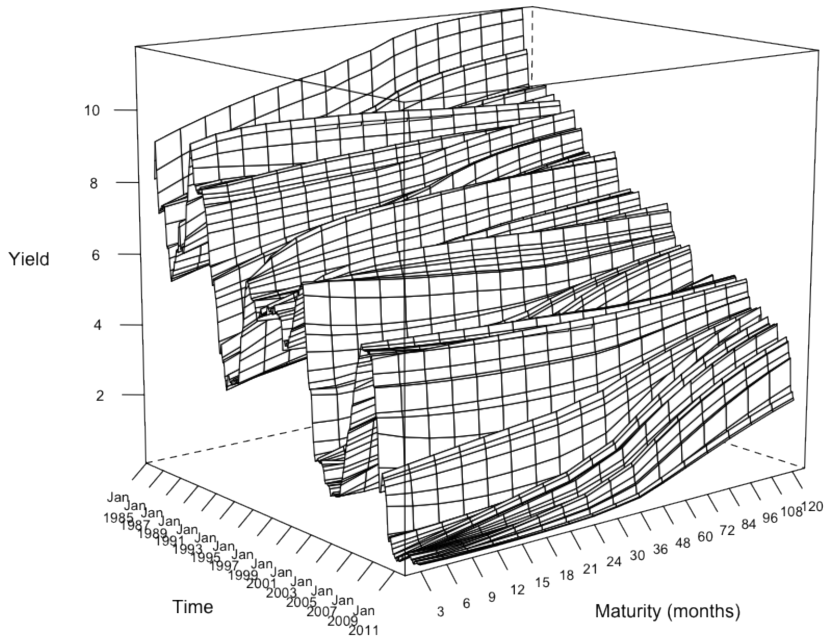

Notes: The figure plots average-of-the-month U.S. Treasury bill and bond yields at maturities ranging from 6 months to 10 years. The three-month yield is taken from the Fed's H-15 release. For longer maturities, we use zero-coupon yields constructed in Gürkaynak et al. (2007). The sample period is January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

[6]This 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 3. Yield data, 1985:1-2011:12. Descriptive statistics

Maturity (months)	*Mean*	*SD*	*Min*	*Max*	*MAE*	*RMSE*	$\hat{\rho}(1)$	$\hat{\rho}(12)$	$\hat{\rho}(30)$
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 4. 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 nonfarm	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

Figure 7. Estimated DNS factors and empirical counterparts

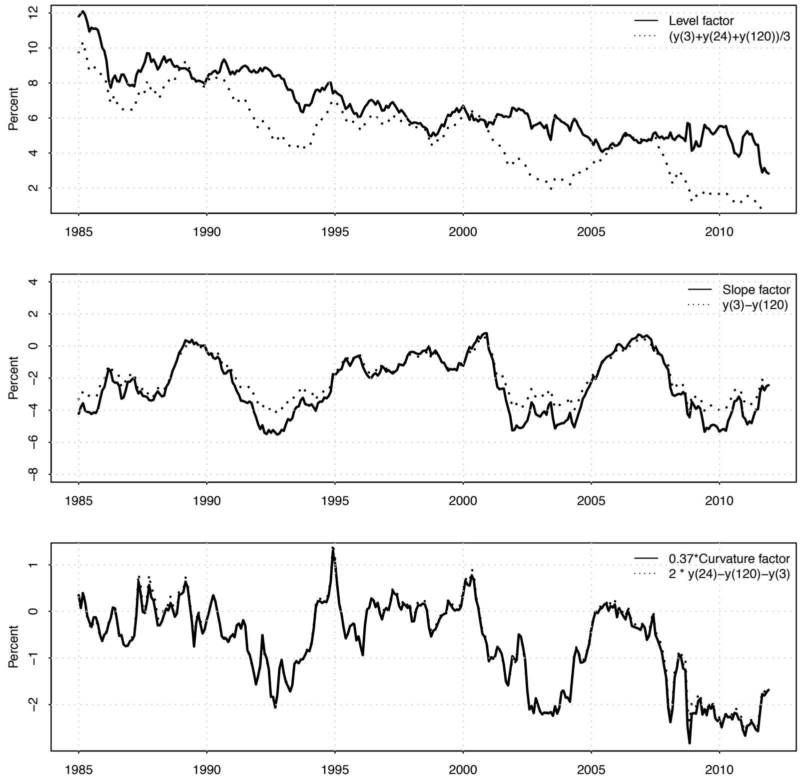

Note. The first factor β_{1t} controls the yield curve level, as it can be verified that $\lim_{\tau\to\infty} y_t(\tau) = \beta_{1t}$. The second factor β_{2t} is related to the yield curve slope, defined as the difference between the 10-year and three-month yields. The third factor β_{3t} governs the curvature of the yield curve, defined as twice the two-year yield minus the sum of the t10-year and three-month yields.

B Appendix. Estimation and in-sample fit of the DNS model

Here we report results regarding the estimation and the in-sample fit of the DNS model.

Figure 7 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 2 basis points and a standard deviation of 4 basis points. Also, the slope factor matches very closely the empirical proxy for the slope with a correlation of .99. The level factor shows 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 0.77, from January 2001 to December 2011 the correlation drops to 0.24. The mean and standard deviations of the difference increase from 122 basis points to 216 basis points and from 99 basis points to 132 basis points, respectively.

Table 5. In-sample fit of the DNS model

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

Table 6. In-sample fit statistic of the DNS model. The model is estimated using monthly yield data from January 1985 to December 2011 with λ_t fixed at 0.0609. The descriptive statistics refer to the corresponding residuals at various maturities. The last three columns present residual sample autocorrelations at lag 1, 12, and 30, respectively.

Table 6 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 11 and 9 basis points. For the other maturities, the residual means oscillate between 1.4 and 7 basis points. Similarly, the residuals' standard deviation is fairly stable, with largest values at the beginning and at the end of the curve.

C Appendix. Survey data

The empirical analysis in the paper considers survey expectations about interest rates collected in the Blue Chip Financial Forecasts (BC). Each month, the BC publishes the forecasts made by approximately 50 professional forecasters at leading consulting firms, investment banks, and academic institutions. The interest rates that the 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 7 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 quarter, i.e. the nowcast, contained in the survey (BCNow, 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 (BCJan,1Q, BCJan,2Q BCJan,3Q in the table).

Table 7. 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^{Now}_{Jan}	BC^{1Q}_{Jan}	BC^{2Q}_{Jan}	BC^{3Q}_{Jan}
February 1^{st}	BC^{1Q}_{Feb}	BC^{2Q}_{Feb}	BC^{3Q}_{Feb}	BC^{4Q}_{Feb}
March 1^{st}	BC^{1Q}_{Mar}	BC^{2Q}_{Mar}	BC^{3Q}_{Mar}	BC^{4Q}_{Mar}
April 1^{st}	BC^{Now}_{Apr}	BC^{1Q}_{Apr}	BC^{2Q}_{Apr}	BC^{3Q}_{Apr}

References

Amisano, G. and R. Giacomini (2007): "Comparing density forecasts via weighted likelihood ratio tests," *Journal of Business & Economic Statistics*, 25, 177–190. 1

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. 1

BIS (2005): "Zero-Coupon Yield Curves: Technical Documentation," *BIS papers*, 25. 1

Campbell, J. R., C. L. Evans, J. D. Fisher, and A. Justiniano (2012): "Macroeconomic Effects of FOMC Forward Guidance," . 4.2

Chernov, M. and P. Mueller (2012): "The term structure of inflation expectations," *Journal of financial economics*. 1

Chun, A. L. (2011): "Expectations, bond yields, and monetary policy," *Review of Financial Studies*, 24, 208–247. 1

——— (2012): "Forecasting Interest Rates and Inflation: Blue Chip Clairvoyants or Econometrics?" in *EFA 2009 Bergen Meetings Paper*. 1

Coroneo, L., D. Giannone, and M. Modugno (2013): "Unspanned Macroeconomic Factors in the Yields Curve," Tech. rep., ULB–Universite Libre de Bruxelles. 1

Csiszár, I. (1975): "I-divergence geometry of probability distributions and minimization problems," *The Annals of Probability*, 146–158. 3.1

Dai, Q. and K. J. Singleton (2000): "Specification analysis of affine term structure models," *The Journal of Finance*, 55, 1943–1978. 1

Diebold, F. X. and C. Li (2006): "Forecasting the term structure of government bond yields," *Journal of econometrics*, 130, 337–364. 1, 2.1, 2.1, 2.2, 5

Diebold, F. X. and R. S. Mariano (1995): "Comparing Predictive Accuracy," *Journal of Business & Economic Statistics*, 13, 253–265. 4.1

Diebold, F. X. and G. D. Rudebusch (2012): *Yield Curve Modeling and Forecasting: The Dynamic Nelson-Siegel Approach*, Princeton University Press. 1, 2.2

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. 1

Duffee, G. R. (2011): "Information in (and not in) the term structure," *Review of Financial Studies*, 24, 2895–2934. 1

Duffie, D. and R. Kan (1996): "A yield-factor model of interest rates," *Mathematical finance*, 6, 379–406. 1

Giacomini, R. and G. Ragusa (2013): "Forecasting with judgment," Mimeo. 1, 3.1

Giacomini, R. and B. Rossi (2010): "Forecast comparisons in unstable environments," *Journal of Applied Econometrics*, 25, 595–620. 3.2

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. 2.2, 2, 4.2, **??**, 3

Hördahl, P., O. Tristani, and D. Vestin (2006): "A joint econometric model of macroeconomic and term-structure dynamics," *Journal of Econometrics*, 131, 405–444. 1

Joslin, S., M. Priebsch, and K. Singleton (2010): "Risk premiums in dynamic term structure models with unspanned macro risks," *Manuscript, Stanford University.* 1

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. 1

Kitamura, Y. and M. Stutzer (1997): "An information-theoretic alternative to generalized method of moments estimation," *Econometrica: Journal of the Econometric Society*, 861–874. 3.1

Kool, C. and D. L. Thornton (2012): "How Effective Is Central Bank Forward Guidance?" FRB of St. Louis Working Paper No 2012. 4.2

Litterman, R. B., J. Scheinkman, and L. Weiss (1991): "Volatility and the yield curve," *The Journal of Fixed Income*, 1, 49–53. 1

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. 1, 2.2

Newey, W. K. and R. J. Smith (2004): "Higher order properties of GMM and generalized empirical likelihood estimators," *Econometrica*, 72, 219–255. 3.1

Ragusa, G. (2011): "Minimum divergence, generalized empirical likelihoods, and higher order expansions," *Econometric Reviews*, 30, 406–456. 3.1

Robertson, J. C., E. W. Tallman, and C. H. Whiteman (2005): "Forecasting Using Relative Entropy," *Journal of Money, Credit and Banking*, 37, 383–401. 1, 3.1

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. 1

Stark, T. (2010): "Realistic evaluation of real-time forecasts in the Survey of Professional Forecasters," *Federal Reserve Bank of Philadelphia Research Rap, Special Report*. 1

Van Dijk, D., S. J. Koopman, M. Van der Wel, and J. Wright (2012): "Forecasting Interest Rates with Shifting Endpoints," Tinbergen Institute Discussion Paper 12-076/4. 1