

Monetary Policy Expectations at the Zero Lower Bound*

Michael D. Bauer[†] and Glenn D. Rudebusch[‡]

February 13, 2014

Abstract

To obtain monetary policy expectations from the yield curve near the zero lower bound (ZLB), it is crucial to account for the distributional asymmetry of future short rates. Because conventional dynamic term structure models (DTSMs) ignore it, they severely violated the ZLB in recent years and performed significantly worse than alternative models that account for it. Shadow-rate models incorporate the ZLB, and therefore achieve superior fit and forecasting performance. In addition, they provide estimates of the modal short-rate path, which can serve to construct accurate and plausible forecasts of the time of future short-rate liftoff.

Keywords: dynamic term structure models, shadow rates, liftoff, macro-finance

JEL Classifications: E43, E44, E52

*The authors thank Todd Clark, Greg Duffee, Jim Hamilton, Leo Krippner, Anh Le, Seth Pruitt, Jean-Paul Renne, Francisco Ruge-Murcia, Eric Swanson, John Williams, and Cynthia Wu, as well as seminar and conference participants at the Banque de France, the Federal Reserve Bank of San Francisco, Research Affiliates, UC Santa Cruz, the Bank of Canada conference “Advances on Fixed Income Modeling,” and the SED 2013 meetings in Seoul for their helpful comments. Kevin Cook and Alison Flint provided excellent research assistance. All remaining errors are ours. The views expressed in this paper are those of the authors and do not necessarily reflect those of others in the Federal Reserve System.

[†]Corresponding author: Federal Reserve Bank of San Francisco, 101 Market St. MS 1130, San Francisco, CA 94109, (415) 974-3299, michael.bauer@sf.frb.org.

[‡]Federal Reserve Bank of San Francisco

1 Introduction

Expectations of future monetary policy actions are commonly obtained from the term structure of interest rates, which captures financial market participants' views regarding the prospective path of the short-term interest rate—the policy instrument of central banks. Gaussian affine dynamic term structure models (DTSMs) are the standard representations in finance that have been used to extract such expectations (e.g., Piazzesi, 2010). However, while these models have provided good empirical representations of yield curves in the past, they appear ill-suited to represent the near-zero interest rates that have prevailed in recent years in several countries. In particular, standard Gaussian DTSMs do not recognize that in the real world, with currency available as an alternative asset, interest rates are bounded below by zero because negative nominal interest rates would lead to riskless arbitrage opportunities. The fact that Gaussian affine DTSMs ignore the zero lower bound (ZLB) previously had little practical effect since interest rates were well above zero.¹ However, as nominal interest rates have fallen to near zero, the lack of an appropriate nonnegativity restriction in conventional models has become a conspicuous deficiency, and the usefulness of these models may have declined substantially.

In this paper, we examine the inconsistency of standard Gaussian affine DTSMs with the ZLB in practice, and we explore the performance of an alternative model based on the shadow-rate concept proposed by Black (1995).² This representation replaces the affine short-rate specification of standard DTSMs with an identical affine process for an unobserved shadow short rate. The short rate is determined by the maximum of zero and this shadow short rate. Specifically, given the likely importance of the ZLB for understanding the dynamics of the U.S. term structure of interest rates during the past several years, we compare the implications of Gaussian affine and shadow-rate DTSMs in recent U.S. interest rate data. We consider both yields-only models as well as macro-finance models, where in the latter case risk factors include measures of slack and inflation.

Our analysis documents the empirical relevance of the ZLB constraint and the importance of accounting for it when carrying out inference about interest rates and monetary policy during the recent period in the United States. Shadow-rate models fit the cross section of yields substantially better than affine models during the ZLB period. Affine models frequently violate the ZLB and produce substantial estimated probabilities of negative future short rates. Shadow-rate models avoid such violations by construction. Affine models also cannot capture

¹We use the ZLB terminology as a convenience, but as discussed below, the precise lower bound on interest rates depends on institutional factors.

²Alternatively, one could consider stochastic-volatility models with square-root processes or Gaussian quadratic models. However, the shadow-rate model has the advantage of matching the canonical Gaussian DTSM when interest rates are away from the ZLB.

the phenomenon of the short rate remaining near zero for many years at a time, which has been the case in the United States and Japan. Consequently, such models produce inaccurate short-rate forecasts at the ZLB. In contrast, shadow-rate models can accurately represent and forecast prolonged near-zero policy rates, as we document in an out-of-sample forecast exercise.

Importantly, shadow-rate models can capture the substantial distributional asymmetry of future short rates that is due to the floor under nominal interest rates. The expected future shadow rates corresponds to the most-likely value of the future short rate—the *modal path*—which is an object of fundamental interest when yields are near zero. Regarding monetary policy expectations at the ZLB, the key question is how to estimate the time of liftoff of the short rate from the zero bound. A common approach among financial market researchers and participants to assess liftoff expectations is to use the horizon at which *forward rates*—the risk-neutral expected future short rates—cross a given threshold, say 25 basis points, as an estimate of the future liftoff date. However, this practice is problematic because it completely ignores the distributional asymmetry of short rates. To appropriately forecast the time of short-rate liftoff, one needs to consider instead the modal path of future short rates, taking the horizon where it crosses the threshold as the liftoff forecast. This simple approach delivers optimal forecasts of the liftoff horizon in the plausible case of an absolute-error loss function. We also compute the full distribution of the liftoff horizon in order to better understand the forecasting problem, to show the optimality of the modal-path-based liftoff estimate, and to obtain interval forecasts for liftoff. Our preferred liftoff estimate closely accords with private-sector and survey forecasts of policy liftoff. Furthermore, the expected duration until liftoff from the ZLB provides a useful partial summary of the stance of monetary policy. In contrast, model-implied shadow short rates, which have been advocated as measures of the policy stance near the ZLB in some academic and policy circles (Bullard, 2012; Krippner, 2013; Wu and Xia, 2013), are largely uninformative, highly sensitive to model specification, and depend on the exact data at the short end of the yield curve. Their lack of robustness is striking and raises a warning flag about using shadow short rates as a measure of monetary policy. Finally, we compare the modal path to forward rates, which reveals how tightly the ZLB constraint is binding. The wedge between the two forecast paths reflects the asymmetry induced by the ZLB on the distribution of future short rates. Furthermore, it also measures the option cost of the ZLB, i.e., the value of the option of holding physical currency. Using this measure, we document that the tightness of the ZLB has increased substantially over the period from 2009 to 2012.

Only a few studies have used shadow-rate DTSMs that respect the ZLB, in large part

because the associated nonlinearity makes it difficult to solve for bond prices. [Bomfim \(2003\)](#) employs a two-factor shadow-rate model to estimate the probability of the future policy rate hitting the ZLB during the 2002–2003 period. [Ueno et al. \(2006\)](#) analyze Japanese interest rates over the period 2001–2006 using a one-factor model, for which [Gorovoi and Linetsky \(2004\)](#) have derived an analytical solution, and [Ichiue and Ueno \(2007\)](#) apply a two-factor model to the same data. [Kim and Singleton \(2012\)](#) estimate two-factor models using Japanese yield data and demonstrate the good performance of shadow-rate models in comparison to alternative DTSM specifications that incorporate the ZLB. [Christensen and Rudebusch \(2013\)](#) estimate one-, two- and three-factor models on the Japanese data and document the sensitivity of shadow-rate estimates to the model specification. Two studies have considered the recent U.S. experience: [Krippner \(2012\)](#) calibrates a restricted two-factor model to obtain estimates of the shadow short rate, and [Ichiue and Ueno \(2013\)](#) compare term premium estimates obtained from approximate shadow-rate and affine Gaussian models. Our study goes beyond these papers in several ways, most notably by documenting the empirical relevance of the ZLB and the problems of affine models in this context, showing the fragility of shadow short rate estimates, and by estimating the tightness of the ZLB over time. Most importantly, our paper shows how to capture monetary policy expectations at the zero lower bound using the modal path, and carefully investigates how to estimate the key object of interest, namely the time of future policy liftoff. Our paper is also the first to introduce macroeconomic factors into a shadow-rate model, following a productive avenue for term structure modeling (e.g., [Rudebusch, 2010](#)). Intuitively, the ZLB limits the information content of the yield curve as shorter-term interest rates are pinned at zero. The associated risk factors of a standard yields-only model are also constrained and not as informative. A shadow-rate model partially lifts this veil by allowing shadow interest rates to take on a broader range of movement—even into negative territory. Those negative shadow rates are of course unobservable; however, the correlation of macro variables with observed interest rates away from the ZLB can help identify the shadow yield curve when observed rates are pinned near zero. In our analysis, we demonstrate empirically the benefits of including macro factors in the information set.

The paper is structured as follows. [Section 2](#) lays out our econometric framework. In [Section 3](#) we document the empirical relevance of taking into account the ZLB constraint, demonstrate the strong model-dependence of estimated shadow short rates, and estimate the modal path and how tightly the ZLB is restraining yields. [Section 4](#) discusses estimation of monetary policy expectations at the ZLB, with a focus on forecasting the future liftoff date, and argues in favor of using the modal short-rate path for this purpose. [Section 5](#) concludes.

2 Dynamic term structure models

In this section, we describe our model specifications, the role of the ZLB constraint in these models, and our empirical implementation, which uses monthly U.S. data.

2.1 Gaussian affine models

The canonical affine Gaussian DTSM is based on three assumptions. First, the short-term interest rate—the one-month rate in our context—is affine in the N risk factors X_t , i.e.,

$$r_t = \delta_0 + \delta_1' X_t. \quad (1)$$

Second, it is assumed that there exists a risk-neutral probability measure \mathbb{Q} which prices all financial assets—hence, there are no arbitrage opportunities—and that under \mathbb{Q} the risk factors follow a Gaussian vector autoregression (VAR),

$$X_t = \mu^{\mathbb{Q}} + \phi^{\mathbb{Q}} X_{t-1} + \Sigma \varepsilon_t^{\mathbb{Q}}, \quad (2)$$

where Σ is lower triangular and $\varepsilon_t^{\mathbb{Q}}$ is an *i.i.d.* standard normal random vector under \mathbb{Q} . Third, under the real-world probability measure \mathbb{P} , X_t also follows a Gaussian VAR,

$$X_t = \mu + \phi X_{t-1} + \Sigma \varepsilon_t, \quad (3)$$

where ε_t is an *i.i.d.* standard normal random vector under \mathbb{P} . Note that these assumptions imply the existence of a stochastic discount factor which is essentially-affine as in [Duffee \(2002\)](#). The price of a bond with a maturity of m periods is determined by

$$P_t^m = E_t^{\mathbb{Q}} \left[\exp \left(- \sum_{i=0}^{m-1} r_{t+i} \right) \right]. \quad (4)$$

In an affine model, this expectation can be found analytically, and it is exponentially affine in the risk factors. Model-implied yields therefore are affine functions of the factors. The details are well-known, but for completeness, we summarize them in [Appendix A](#). Importantly, a Gaussian model implies that interest rates can turn negative with non-zero probability. During times of near-zero interest rates, violations of the ZLB can be quite prevalent, and we document this empirically in [Section 3.2](#).

2.2 Shadow-rate models

Following [Black \(1995\)](#), our shadow-rate DTSMs are specified exactly like standard affine models except the affine short-rate equation (1) is replaced by a shadow-rate specification:

$$r_t = \max(s_t, \underline{r}), \quad s_t = \delta_0 + \delta_1' X_t. \quad (5)$$

The shadow short rate, s_t , is modeled as affine Gaussian. Equation (5) ensures that the short rate and all other model-implied interest rates cannot go below \underline{r} . [Black \(1995\)](#) set $\underline{r} = 0$, and this is our choice as well. This ZLB on nominal interest rates is typically motivated by the presence of physical currency.³ Since the storage and use of large amounts of physical currency can incur significant transaction costs, the ZLB has been violated at times in the past when interest rates have dipped into negative territory but remained close to zero. We could account for this fact by specifying a slightly negative value for \underline{r} . On the other hand, the federal funds rate, the key short-term interest rate managed by the Federal Reserve, in practice typically remains above zero (in part because it pertains to an unsecured loan), which would be an argument in favor of a slightly positive value for \underline{r} . Different authors have made alternative choices, e.g., [Wu and Xia \(2013\)](#) set $\underline{r} = 25$ basis points, and [Kim and Priebsch \(2013\)](#) treat \underline{r} as a parameter and estimate it. Below, we investigate the sensitivity of our results to different choices of \underline{r} .

A key advantages of shadow-rate models is that, except for the alternative short-rate equation, they are the same as affine Gaussian models, a mainstay of term structure analysis. Therefore, a shadow-rate model retains many of the features and advantages of an affine Gaussian model, and away from the ZLB, it behaves exactly as the standard Gaussian DTSM. Another major advantage of shadow-rate models is that in contrast to other tractable non-Gaussian models that respect the ZLB constraint, such as square-root diffusion (Cox-Ingersoll-Ross) models and quadratic models, the probability of a zero future short rate is non-zero. This becomes crucial when addressing the issue of the duration of near-zero policy rates and the time of future liftoff, as we do in this paper. Appendix B provides the moments of future values of s_t and r_t , that are key for that type of analysis.

A shadow-rate model does not lead to closed-form solutions for yields and bond prices. Hence the need arises for numerical solution methods, and [Priebsch \(2013\)](#) discusses and evaluates alternative approaches.⁴ We use Monte Carlo simulations to evaluate the expectation

³[Fisher \(1896\)](#) pointed out that investors would rather store their wealth in currency than to lend it at a loss.

⁴[Krippner \(2012\)](#) has proposed expressing forward rates in a shadow-rate context as the sum of shadow forward rates and an option effect. This can lead to quasi-analytical solutions if the value of the option effect

in equation (4), which is a flexible and very reliable method. Importantly, the computational cost of Monte Carlo simulation does not substantially increase with a higher number of risk factors, and one can attain arbitrarily high accuracy. Details about the implementation and evidence of the accuracy of our method are in Appendix C.

2.3 Risk factors

A key modeling choice is which risk factors to include in the DTSM. We estimate both “yields-only” models, where X_t reflects only information in the yield curve, and “macro-finance” models, where X_t also includes macroeconomic variables.

We estimate yields-only affine and shadow-rate models using both two-factor and three-factor specifications. We denote the affine models by YA(2) and YA(3) and the shadow-rate models by YZ(2) and YZ(3). We use the canonical form of Joslin et al. (2011). The risk factors are linear combinations of yields, with the weights corresponding to the first N principal components of observed yields. We assume that all yields are measured with error, hence X_t is latent. In the YA models, the risk factors are linear combinations of model-implied yields—they correspond to level, slope, and curvature of the yield curve.⁵ In the YZ models, the yield factors are linear combinations of *shadow yields*—the yields that obtain when the shadow short rate is used for discounting payoffs—so that they can be interpreted as shadow level, shadow slope, and shadow curvature.⁶

Macroeconomic variables are likely to be particularly informative when the yield curve is constrained by the ZLB. Accordingly, we estimate macro-finance DTSMs that include measures of inflation and economic activity in addition to the yield factors. Here, we use the canonical form of Joslin et al. (2013). We estimate affine and shadow-rate models with one ($L = 1$) or two ($L = 2$) yield factors in addition to the two macro factors, and denote these models by MA(1), MA(2), MZ(1), and MZ(2). Again, yield factors are linear combinations of (model-implied/shadow) yields, with weights corresponding to principal components. Yields are assumed to be measured with error, while macro factors are taken as observed.⁷

In our macro-finance models, the macroeconomic variables are spanned by the yield curve.⁸

can be derived using an option pricing framework. Christensen and Rudebusch (2013) perform the necessary derivations for their affine Nelson-Siegel model and are able to apply this approach empirically.

⁵Our affine yields-only models correspond to the RKF model specification in Joslin et al. (2011).

⁶Since shadow yields are calculated from shadow bond prices, where discounting is carried out using the affine shadow short rate, they are calculated with the usual affine loadings.

⁷We do not allow for measurement errors on the macro factors, because in that case “the likelihood function largely gives up on fitting the observed macro factors in favor of more accurate pricing of bonds” (Joslin et al., 2013). The affine macro-finance models correspond to the TS^f specification in Joslin et al. (2013).

⁸In the shadow-rate models, the macro factors are spanned by the (unobservable) shadow yields.

An alternative would have unspanned macro risks as in [Joslin et al. \(2012\)](#). In such models, the current short rate and yield curve depends only on the yield factors. Here we maintain the assumption that macroeconomic conditions affect the current short-term interest rate and yield curve, so that they are informative for inferring policy expectations under the risk-neutral measure. This specification is consistent with the expressed view of the FOMC that the short rate will be based on the unemployment and inflation rates.⁹

2.4 Data, filtering, and estimation

Our data consist of monthly observations of interest rates and macroeconomic variables from January 1985 to December 2013. For the short end of the yield curve, we use three-month and six-month T-bill rates.¹⁰ The remaining rates are smoothed zero-coupon Treasury yields with maturities of one, two, three, five, seven, and ten years from [Gürkaynak et al. \(2007\)](#).¹¹ We measure economic activity by the unemployment gap, using the estimate of the natural rate of unemployment from the Congressional Budget Office. Inflation is measured by the year-over-year percent change in the consumer price index (CPI) for all items excluding food and energy, i.e., by core CPI inflation. We include the inflation and gap measures because these are closely linked to the target federal funds rate, the policy instrument of the Federal Reserve ([Rudebusch, 2006, 2009](#)). [Figure 1](#) gives a view of our data. The top panel shows the yields with maturities of three months, two years, and ten years. The bottom panel shows the macroeconomic variables.

Denote the vector of $J = 8$ model-implied yields by Y_t . Throughout, we assume homoskedastic Gaussian measurement error for yields. For the affine models, we have $Y_t = A + BX_t$, with J -vector A and $J \times N$ -matrix B containing the usual affine loadings. In this case, the state-space system is linear, and the Kalman filter can be used for inferring the latent factors and calculating the likelihood. We carry out fast and reliable maximum likelihood estimation using the following approach: First, we assume that the yield factors are observed, so that μ and ϕ can be obtained using least squares ([Joslin et al., 2011](#)), and the remaining parameters are found by maximizing the likelihood function for given VAR parameters. Second, we use these estimates as starting values to estimate the models with latent yield factors using the Kalman filter. The optimization quickly converges to the global maximum because of the good starting values.

⁹One consequence of this assumption is that the model implies a simple policy rule for the short rate, which during normal times can be used to assess the stance of monetary policy.

¹⁰T-bill rates are obtained from the Federal Reserve's H.15 release, see <http://www.federalreserve.gov/releases/h15/data.htm>.

¹¹These yields are available at <http://www.federalreserve.gov/pubs/feds/2006/200628/200628abs.html>.

For the shadow-rate models, we have $Y_t = g(X_t)$, where the function $g(\cdot)$ is nonlinear and not known in closed form. We approximate it using Monte Carlo simulation—details are in Appendix C. Since the observation equation becomes nonlinear, we use the Extended Kalman Filter (EKF). This requires only the calculation of the Jacobian of $g(\cdot)$, which we approximate numerically.

Estimation of the shadow-rate models using maximum likelihood and the EKF is computationally costly. We use a simple work-around which dramatically reduces computational cost and is satisfactory in practice: We obtain the parameters by estimating the affine models over a subsample during which the ZLB was largely irrelevant, specifically, the sample period ending in December 2007. We use these parameters for both the affine model and for the corresponding shadow-rate model. There are two assumptions underlying the validity of this approach. The first assumption is that the affine model and the shadow-rate model have close to identical implications on the estimation subsample, so that estimates of the model parameters are interchangeable between the two models. Appendix D provides evidence that this is a reasonable assumption by comparing affine and shadow-rate models over this subsample. The second assumption is that ending the sample before 2008 does not materially change the parameter estimates. We have performed the full estimation for the yields-only models $YZ(2)$ and $YZ(3)$ —these have fewer parameters than the macro-finance models and hence are time-consuming but manageable to estimate. We found that the parameter estimates barely change, relative to the sub-sample estimation, and that the economic implications remain qualitatively and quantitatively the same. An advantage of using the same parameters in both affine and shadow-rate models is that we can clearly see the effect of introducing the ZLB constraint while keeping the parameters the same.

3 Effects of the ZLB constraint

From a theoretical perspective, shadow-rate models have a fundamental advantage over affine models in that they impose the nonnegativity of nominal interest rates. But how relevant is this in practice? In this section, we first evaluate affine and shadow-rate models during a period of near-zero interest rates. Then, we discuss and measure how the ZLB constraint affects current short rates and the distribution of future short rates.

3.1 Cross-sectional fit

We first assess the cross-sectional fit of model-implied yields to observed yields for affine and shadow-rate models. Table 1 shows the root mean-squared fitting errors (RMSEs) across models for the whole cross section of yields and for each yield maturity separately. The top panel reports RMSEs for the whole sample, while the bottom panel reports the fit for the ZLB subsample, here and in the following taken as the period from December 2008 to December 2013.¹² In general, shadow-rate models fit yields more accurately than their affine counterparts. The improvements are larger for those models with worse fit, such as MA(1), while for affine models that fit the cross section well, such as YA(3), the improvements from introducing a shadow-rate specification are more modest. The improved fit of shadow-rate models is entirely due to the ZLB period. The bottom panel of Table 1 shows that improvements in RMSEs are very substantial for this subsample. For the pre-2008 period the affine and shadow-rate models have essentially identical cross-sectional fit (see also Appendix D), because away from the ZLB, the implications of these models are the same. Near the ZLB, however, shadow-rate models have additional flexibility in fitting the cross section of yields, which behaves in an unusual way due to the pronounced nonlinearity at zero.

3.2 Violations of the ZLB by affine models

To understand the relevance of the ZLB for term structure modeling in recent U.S. data, it is important to measure the extent to which affine models violate this constraint. One form of violation of the ZLB occurs when model-implied paths of future short rates drop below zero at some horizons. This can happen for either forward rates, which are equal to expected future short rates under \mathbb{Q} , or for real-world (\mathbb{P} -measure) short-rate expectations.¹³ Table 2 shows the number of months that forward rates or expected future short rates drop below zero in each affine model. Also shown is the average length of horizon that the paths stay in negative territory. All affine models imply short-rate paths that frequently and severely violate the ZLB constraint, and this holds for both forward curves and short-rate expectations.

Even when the expectation for the future short rate is positive, the model-implied probability distribution for the future short rate, which is Gaussian, may put nonnegligible mass on negative outcomes. Figure 2 plots the time series of conditional probabilities of negative

¹²On December 16, 2008, the FOMC lowered the target for the federal funds rate to a range from 0 to 25 basis points, hence we choose December 2008 as the first month of the ZLB subsample.

¹³Throughout this paper, we refer to \mathbb{Q} -measure expectations of future short rates as forward rates. These differ from the actual forward rates, which can be contracted by simultaneously buying and selling bonds of different maturities, by a convexity term.

future short rates at horizons of 6, 12, and 24 months in the future, for the period from 2000 to 2013. The four panels in the figure show these probabilities for each affine model. Note that even during the extended period of monetary easing after the 2001 recession, the probability of negative future short rates was nonnegligible. For the more recent period of near-zero short rates from 2008 to 2013, all four affine models imply that these probabilities are very high. The macro-finance affine models lead to larger probabilities over this period than the yields-only models. The reason is that the high unemployment and subdued inflation toward the end of the sample imply paths of expected future short-term rates which are very low, reflecting expectations of easier future monetary policy. This leads to substantially higher probabilities of negative future short rates.

3.3 Forecasting at the ZLB

Affine models produce frequent and severe ZLB violations in the recent U.S. data. Does this matter for forecasting interest rates? While affine models may imply negative forecasts of future interest rates, a pragmatic solution is to simply set these forecasts to zero, and “fixing” them in this way may lead to sufficiently accurate forecasts. To address this question, we investigate the out-of-sample forecast accuracy of affine and shadow-rate models during the ZLB period, focusing on the three-month T-bill rate as the forecast target. For each month from December 2008 to June 2012, we calculate model-based forecasts of this rate for horizons up to 24 months. We use a fixed-window forecast scheme, i.e., we do not re-estimate the models. For the affine models, we replace negative forecasts by zero. Table 3 shows the RMSEs in percentage points for selected forecast horizons across models.

The shadow-rate models predict the short rate more accurately than the affine models, and the differences in forecast accuracy are large at longer horizons. For the two-year forecast horizon, yields-only shadow-rate models forecast 29-36% more accurately than the corresponding affine models, and macro-finance models forecast 21-43% more accurately at this horizon. The only case where there are no or only small improvements in forecast performance is when affine models forecast negative rates that are truncated to zero, as is typically the case for the MZ(1) model horizons up to 18 horizons. This, of course, should not be viewed as a success of affine models but rather as a quick fix that is done in practice to make their forecasts at least somewhat sensible. In general, shadow-rate models are at least as accurate and typically much more accurate than affine models when forecasting interest rates near the ZLB.

This demonstrates the importance of accounting for the ZLB constraint when performing inference about the yield curve during a period of near-zero short-term interest rates. While a sufficiently flexible affine model might be able to satisfactorily fit the yield curve, any type of

economic inference is prone to be misleading. The ZLB has the effect that implied short-rate paths, forecasts, and term premia (which are implied by short-rate forecasts), produced by conventional DTSMs are likely to be seriously distorted and cannot be trusted. Hence, it is advisable for researchers to instead use models that incorporate the ZLB constraint such as shadow-rate models.

Another interesting result that emerges from Table 3 is the benefit for forecasting performance at the ZLB of incorporating macroeconomic information into DTSMs. Comparing macro-finance models to yields-only models, we observe that for horizons longer than six months the former substantially outperform the latter. The inclusion of macroeconomic factors leads to improvements in forecast accuracy that are quite dramatic. When the yield curve is constrained by the zero bound, it necessarily carries less information about the future path of monetary policy than during normal times, which explains the substantial benefits to including macroeconomic variables.

3.4 The short-term interest rate

The distinguishing feature of the ZLB period is that the current short-term interest rate is stuck near zero—at the time of this writing, the target for the fed funds rate has been near zero for over five years. During this period, the Federal Reserve has implemented unconventional policies to ease financial conditions, but due to the ZLB constraint the usual indicator of monetary policy, the short rate, has not changed. Because Black (1995) described the shadow short rate as the short-term interest rate that would prevail in the absence of the option of holding physical currency, i.e., in the absence of the ZLB, some have interpreted the shadow short rate as an alternative indicator of the stance of monetary policy—see, in particular, Krippner (2012), Ichiue and Ueno (2013), and Wu and Xia (2013).¹⁴

Figure 3 shows estimated shadow short rates, together with the three-month T-bill rate, from 2005 to 2013. The top panel shows those estimates implied by our four models. When the short rate is well above zero, the various estimated shadow short rates generally match the observed short rate. During the recent ZLB period, however, the models disagree substantially about the value of the shadow short rate. The MZ(1) model implies by far the most negative shadow short rate, the MZ(2) and YZ(2) models lead to shadow short rates that are slightly negative, and the YZ(3) model produces a shadow short rate that is mostly positive and very

¹⁴Bullard (2012) has taken Krippner’s estimates of a very negative shadow rate in the United States as evidence of a very easy stance of monetary policy. Researchers at the Federal Reserve Bank of Atlanta also have taken an estimated shadow short rate as the “effective” fed funds rate during the ZLB period. See <http://macroblog.typepad.com/macroblog/2013/11/the-shadow-knows-the-fed-funds-rate.html>.

close to zero. Evidently, the estimates of the shadow rate near the ZLB are highly model-dependent.¹⁵ Shadow short rate estimates are also highly sensitive to the numerical value of the lower bound. The bottom panel of Figure 3 shows the estimated shadow short rate series for model YZ(3) when we vary the lower bound \underline{r} from 0 to 25 basis points in 5 basis-point increments. The lack of robustness is quite striking. Importantly, given the frequent observation of actual short-term interest rates within a few basis points of zero during the recent period—see, for example, Figure 3 in Gagnon and Sack (2013)—the choice of $\underline{r} = 0$ is surely as defensible as the choices of other authors (Kim and Priebsch, 2013; Wu and Xia, 2013). Consequently, it appears difficult to draw robust inference about the contemporaneous shadow short rates.

An issue closely related to the observation that shadow-rate estimates are not robust is the fact that these reflect only information at the short end of the term structure. Shadow short rate estimates are, in our experience, highly sensitive to the short-term yields included in the estimation. Essentially, shadow short rates are determined by the behavior of short-term yields, and their relation to the model-imposed lower bound on yields, which is somewhat arbitrary. However, in a ZLB situation, the short end of the term structure does not convey much information about the stance of monetary policy. The Fed’s unconventional policies, for example, have explicitly targeted longer-term interest rates.

Overall, the lack of robust estimates of shadow short rates raises a warning flag and questions their usefulness. While Krippner (2012), Ichiue and Ueno (2013), and Wu and Xia (2013) each focus on one specific model and take the estimated shadow short rate as the literal truth, our findings strongly suggest that these estimates are not robust and should be considered with a large grain of salt. For these reasons, we caution against putting too much confidence into estimated shadow short rates as indicators of monetary policy at the ZLB.

3.5 The tightness of the ZLB

To understand the effects of the ZLB on the term structure of interest rates, we focus on the asymmetry of the distribution of future short rates introduced by the ZLB. The extent of this asymmetry reveals how strongly the ZLB is binding. Here and in the remainder of the paper, we will focus exclusively on risk-neutral (Q-measure) distributions and its moments. Importantly, our methodology and all the points we make below are equally applicable to the

¹⁵Christensen and Rudebusch (2013) show that shadow-rate DTSMs estimated on Japanese bond yields differ substantially in their implications about the level of the current shadow short rate depending on the number of factors used. Kim and Singleton (2012) similarly find strong model-dependence of their shadow rate estimates.

case when the analysis is carried over to distributions under the real-world (P-) measure.¹⁶

Figure 4 shows risk-neutral probability densities for the future shadow short rate and future short rate implied by model MZ(2) on December 31, 2012, at a horizon of $h = 48$ months. For the future shadow rate, this density is Gaussian, centered around $E^Q(s_{t+h}|X_t)$. The density of the future short rate has a point mass at zero, indicated in the graph with a vertical line, and for positive values equals the shadow-rate density. For what follows, it will be useful to define the mode of this distribution uniquely as $\max(0, E^Q(s_{t+h}|X_t))$ (as in Kim and Singleton, 2012). The probability mass of the short rate density at zero corresponds to the probability of a negative future shadow short rate. During normal times, this probability is negligibly small, so that the mean and the mode of the short rate distribution will approximately coincide. The more relevant the ZLB becomes, the larger the asymmetry of the distribution of future short rates, and the larger the difference between mean and mode, the “ZLB wedge,” will become. Notably, it depends not only on the distance of yields to zero, but also on second moments. The ZLB wedge captures the option cost of the ZLB in the sense that it equals the cost introduced by the optionality in equation (5). Differently put, this measure captures the value of the option of holding physical currency, which restrains nominal interest rates when these approach zero.

The *modal path* corresponds to the mode of the future short rate distribution across horizons, i.e., the most-likely path of future short rates. It is identical to risk-neutral expectations of future shadow short rates—shadow forward rates—when these are positive, and equal to zero when these are negative. In contrast, forward rates are risk-neutral expectations of future short rates. The difference between the forward curve and the modal path corresponds to the ZLB wedge. Figure 5 displays these paths for two dates, June 30, 2011 and December 31, 2012, for models YZ(3) and MZ(2)—here and in the following sections, we will focus on these two flexible models. For the earlier date, both models imply that there is a slight difference between forward rates and modal path at short horizons, but that this difference becomes small for horizons beyond two years. In contrast, 18 months later, the ZLB wedge has become very pronounced, and even for more distant horizons there is a substantial difference between expectations of future short rate and future shadow short rates. Evidently, the ZLB constraint

¹⁶The main advantage is that the parameters μ^Q and ϕ^Q are estimated very accurately, due to the large amount of cross-sectional information in the yield curve (Cochrane and Piazzesi, 2008; Kim and Orphanides, 2012). In contrast, inference about the VAR parameters μ and ϕ and about the real-world distribution of future short rates is fraught with severe statistical problems due to the high persistence of interest rates (Bauer et al., 2012; Duffee and Stanton, 2012). The reader should keep in mind that from now on, the term “expectations” will be understood to include a term premium component that adjusts for risk. In practice, forward rates over short and medium horizons are typically taken to closely correspond to real-world expectations, based on the reasonable assumption that risk premia are likely small at such maturities.

played a more important role in December 2012 than in June 2011.¹⁷

To visualize the impact of the ZLB on the yield curve, Figure 6 shows actual yields together with the fitted yield curves and shadow yield curves implied by the shadow-rate models YZ(3) and MZ(2) on June 30, 2011, and on December 31, 2012. The left panels of Figure 6 show that in June 2011, shadow yields were, for longer maturities, about 20 to 30 basis points below actual yields. Thus, there was a noticeable effect of the ZLB on the entire yield curve. Although forward rates on this date were only noticeably affected at short maturities, yields at all maturities are constrained to some extent by the ZLB, simply because long-term yields reflect the behavior of average forward rates up to the specific maturity. The right panels show that for the later date, the differences between fitted and shadow yields at long maturities is much larger, around 60 to 70 basis points. The ZLB clearly was constraining yields more tightly at the end of 2012, in line with what we saw for forward rates.

The ZLB wedge in a long-term interest rate measures how tightly the ZLB constrains the entire term structure of interest rates, because it equals the cumulated difference between forward rates and shadow forward rates. We now focus on the ten-year maturity and show, in Figure 7, the evolution of the yield, shadow yield, and the corresponding ZLB wedge over time, for models YZ(3) and MZ(2). Evidently, the ZLB has increasingly constrained interest rates over the period from 2009 to 2012. This finding is consistent with Swanson and Williams (2012), who measure the tightness of the ZLB using the sensitivity of different interest rates to macroeconomic news, and document that this sensitivity has decreased for most yields over this period.¹⁸

Above, we have raised concerns about estimated shadow short rates due to their sensitivity to the particular model specification used. In contrast, we have found that inference about shadow forward rates, the modal path, and shadow yields is generally much more robust to the specific choice of model. Pairwise correlations for estimates from two different models, and in particular for different choices of \underline{r} , are close to one (results not shown). The reason that the model dependence is much less pronounced for these estimates than for the shadow short rate is that the yields that pin these down have longer maturities and carry more information than the short end of the yield curve, which is pinned at zero.

¹⁷Note that this figure also demonstrate the very limited amount of information of the shadow short rate at a given point in time. Its value is close to zero on both dates, which is clearly not representative of the whole curve or of the tightness of the ZLB constraint.

¹⁸Increases in the tightness of the ZLB often coincided with key Fed announcements of easier monetary policy, which pushed long-term interest rates closer to their lower bound, as evident also in the top panel of the figure. A notable example is the switch to more explicit forward guidance by the FOMC in fall 2011, which pushed out the expected duration of near-zero policy rates.

4 Monetary policy liftoff

Evidently, the ZLB has had a substantial impact on the term structure of interest rates in the United States in recent years. It is important to account for this fact, and shadow-rate DTSMs do so. In this section, we turn to the issue of how these models can be used to estimate monetary policy expectations at the zero lower bound. During a period of near-zero policy rates, the timing of the future liftoff of the short rate from zero naturally is a key object of interest. Here, we discuss how to use yield curve information to appropriately estimate it. We compare alternative approaches and argue in favor of using the modal path for this purpose, which leads to optimal liftoff forecasts.

4.1 Optimal liftoff forecasts based on the target distribution

The natural starting point for estimating future policy liftoff is the distribution of this random variable. Hence, one needs to first obtain this target distribution. In a shadow-rate model, liftoff corresponds to the first hitting time of a given threshold by the shadow rate. In the case of a univariate model, its probability density can in some cases be derived analytically.¹⁹ In multi-factor shadow-rate models, the liftoff distribution has to be obtained using simulation. Our approach to do so is the following: Starting from the current term structure at t , we simulate 10,000 sample paths for the shadow rate using the risk-neutral dynamics of the risk factors in equation (2). For each simulation, the liftoff horizon is determined by the time that the shadow rate hits a threshold of 25 basis points from below—this choice is motivated by the fact that the Federal Reserve has kept the policy rate in a range from 0 to 25 basis points over the ZLB period, and the time when it leaves this range is an obvious choice for defining policy liftoff.²⁰ Figure 8 shows smoothed Kernel densities of the empirical liftoff distribution on December 31, 2012, based on simulations from models YZ(3) and MZ(2), and reports the mean, median, mode, and interquartile range in each case. Notably, the distribution is very strongly skewed to the right—even very distant horizons for policy liftoff are not uncommon. The figure also reports alternative liftoff estimates based on the forward curve and the modal path that we will discuss below.

How to forecast policy liftoff based on these distributions depends on the forecasters loss

¹⁹See [Linetsky \(2004\)](#), who uses analytical results for hitting times of an Ornstein-Uhlenbeck process to calculate the distribution of the liftoff horizon in a shadow-rate Vasicek model. These results are also used in [Ueno et al. \(2006\)](#) and [Ichiue and Ueno \(2012\)](#), who consider one-factor shadow-rate DTSMs.

²⁰Due to the fact that at some times the shadow rate at t starts out above zero, and due to the erratic nature of the sample paths, we require that the shadow rate stays above the threshold for 12 months before we determine a crossing of the threshold as policy liftoff. Our definition of liftoff leads to better-behaved liftoff distribution than simply taking the first horizon of crossing the threshold as the liftoff.

function. For squared-error loss, the mean is the optimal forecast. Given the large weight this gives to outliers, the mean is not a very useful summary of the central tendency, and it leads to unreasonably large liftoff forecasts. The median is the optimal forecast under absolute-error loss. This loss function appears reasonable in this context, given the fat right tail of the target distribution, since outliers do not have an unduly large effect. The mode is the optimal forecast for the case of step loss (also called 0–1 loss). This loss function is unappealing for the problem at hand, since the magnitude of the error is not taken into account, and the skewness is effectively ignored.²¹ Given the nature of the problem and the pronounced skewness of the liftoff distribution, we view absolute-error loss and the median as the optimal forecast as a very reasonable choice to estimate future policy liftoff. It will turn out that this choice leads to empirical results that are closely in line with FOMC guidance and private-sector forecasts.

4.2 Forecasting liftoff using the term structure

A common practice is to base estimates of liftoff on current forward rates or money market futures rates, using the horizon when these rates rise above 25 basis points.²² At first sight, this approach seems to be a natural choice, since forward rates reflect (risk-neutral) expectations of future short-term interest rates. Furthermore, it does not require a model and is therefore not dependent on specific modeling assumptions. However, this approach is flawed, because it ignores the asymmetry of the distribution of the future short-term interest rate induced by the ZLB. Forward rates do not reflect the most likely value of the policy rate at a future point in time, as noted in Section 3.5. To see that estimates based on forward rates can be misleading, consider a mean-preserving spread in the future *shadow* rate distribution. This would raise forward rates, because the ZLB wedge—the difference between mean and mode of future short-rate expectations—would increase, and therefore push earlier those liftoff estimates that are based on forward rates. However, the modal path for the policy rate was unaffected, and the liftoff estimate should not have changed.

Instead of focusing on forward or futures rates, the modal path should form the basis for liftoff estimates. The reason is that it appropriately takes into account the distributional asymmetry of future short rates near the ZLB. Some professional forecasters appear to estimate

²¹Ichiue and Ueno (2012) have taken the mode of the liftoff distribution as their forecast of policy liftoff, without any reference to forecast optimality. They argue in favor of the mode over the mean based on what they view as better comparability with survey forecasts. They do not discuss the median of the hitting-time distribution.

²²Ueno et al. (2006) take the horizon where Euroyen futures rates exceed a given threshold as an estimate of future policy liftoff by the Bank of Japan. There are many examples in financial market commentary that base liftoff estimates on this approach, including “Fed Likely to Push Back on Market Expectations of Rate Increase,” Wall Street Journal from 6/13/2013.

liftoff in this way: They first construct their most likely path for the future policy rate and treat the first increase in this path above a 25 basis point threshold as an estimate of policy liftoff.²³ We will show below that this approach coincides with the optimal forecast under absolute-error loss, i.e., the median of the liftoff distribution. Furthermore, it leads to liftoff forecasts that are well in line with to outside estimates.

Figure 5 illustrates the alternative liftoff estimates based on the term structure. Horizontal lines at 25 basis points indicate the threshold for liftoff. As shadow forward rates and the modal path are always below forward rates, the liftoff implied by the modal short rate path is always later than the liftoff estimated from forward rates. Focusing on December 2012, the differences are quite pronounced, due to the stronger asymmetry of the future short rate distribution. For model YZ(3), the modal path lifts off at a horizon of 33 months, corresponding to September 2015. The forward curve, in contrast, crosses the threshold much earlier, at a horizon of 14 months, which would put liftoff at February 2014. For model MZ(2) the modal-path liftoff forecast is also 33 months, while the estimate based on the forward curve is 21 months, i.e., September 2014.

It is useful to put these estimates in perspective based on FOMC communication and survey forecasts. The FOMC has provided explicit forward guidance about the likely duration of the period of near-zero policy rates, at first based on calendar dates and later based on economic outcomes. In December 2012, it stated that it expected policy liftoff to occur only after the unemployment rate had fallen to at least 6.5% in a context of price stability. Out of 19 FOMC participants, 13 expected policy liftoff to occur in 2015, according to the Summary of Economic Projections (SEP), which also indicated that the central tendency of the participants' unemployment rate forecast for the fourth quarter of 2015 was 6.0 to 6.6%. By comparison, in the January 2013 Blue Chip Financial Forecasts survey, which was based on responses collected near the end of December, 80% of the respondents expected the unemployment rate to first dip below 6.5% in 2015 or later. And of the respondents to the Federal Reserve Bank of New York's Primary Dealer survey in January 2013, 84% expected liftoff to occur in 2015 or later. In sum, in December 2012 the FOMC and market participants appeared to expect liftoff around the middle of 2015 or later. While the liftoff forecasts based on the model-implied modal paths appear reasonable from this perspective, straight readings from the fitted forward curves lead to estimates of liftoff distribution that are substantially too early. In the following, we show that this observation for December 2012 is a general pattern.

²³The responses in the Primary Dealer Survey are consistent with the view that respondents base their liftoff estimate on the modal path. Other examples of analysis in line with this approach include "Reading the Tea Leaves of Rate Expectations," Goldman Sachs US Economic Analyst from 7/3/2013.

4.3 Comparing liftoff estimates over the recent ZLB period

We now consider the time series of alternative estimates for policy liftoff. Figure 9 compares liftoff estimates based on the forward curve and the modal short-rate path, as well as the median and interquartile range of the simulated liftoff distribution. The top panel shows the estimates for the yields-only model YZ(3), while the bottom panel displays them for the macro-finance model MZ(2). The period shown is from January 2008 to December 2013.

To evaluate the model-based estimates, we compare them to two alternative calculations of future liftoff dates by the private sector. The first is from the Primary Dealer Survey of the Federal Reserve Bank of New York, which is publicly available going back to January 2011.²⁴ In particular, we report the median of the Primary Dealers' modal forecasts for the time of policy liftoff.²⁵ The second alternative source of liftoff estimates is from the projections of the future path of the federal funds rate by Macroeconomic Advisers, which are also modal forecasts, i.e., the most likely scenario for Fed policy in their view.

Over time, our liftoff estimates based on the modal short-rate path match closely the median of the liftoff distribution. For model MZ(2), the two are essentially identical from 2010 onward. The close correspondence of the modal-path forecast and the median of the liftoff distribution is intuitive: Denote by h^* the horizon where the modal path crosses the threshold. Due to the symmetry of the shadow-rate distribution, it is equally likely for the shadow-rate path to be above or below the threshold at $t + h^*$. Since all paths that are above the threshold at this horizon have lifted off already, the probability mass for the event of liftoff between t and $t + h^*$ will be *at least* 0.5. Since most paths that are below the threshold have not lifted off yet, the probability of liftoff after $t + h^*$ will be below but close to 0.5. A small discrepancy between these probabilities and 0.5 arises because in some cases the shadow rate path might rise above the threshold and then fall again below it before $t + h^*$, but the chance of this happening will generally be small. Hence, the median of the liftoff distribution will always be very close to h^* , so the modal-path forecast of liftoff are approximately optimal under absolute-error loss. This optimality property is comforting, and for practical purposes, it is satisfactory to see that forecasting liftoff in a straightforward fashion using the modal path is justified based on a more thorough analysis of the forecasting problem.

Our preferred liftoff estimates are also generally close to those from the Primary Dealer survey and Macroeconomic Advisers, which provides further support for the relevance of our estimates. In particular, the models capture the substantial increase in the most likely ZLB

²⁴See http://www.newyorkfed.org/markets/primarydealer_survey_questions.html (accessed February 8, 2013) for the questions and answers for each survey.

²⁵In the survey, the respondents are asked to provide the “estimate for [the] most likely quarter and year of [the] first target rate increase.”

duration between mid 2011 and late 2011. The FOMC announced on August 9, 2011, that it expected a near-zero policy rate until at least mid-2013, which caused the expected duration to jump up (see also [Swanson and Williams, 2012](#)). This information is reflected in the outside estimates as well as in our model-based estimates.

The forward-curve based estimates are quite unsatisfactory. They imply liftoff horizons that are much too early, by as much as one year or more, both relative to our preferred estimates and relative to outside estimates. The difference is particularly pronounced late in the ZLB period. The increasing downward bias reflects the greater constraining effect of the ZLB and the consequent greater asymmetry in the future short-rate distribution.

Figure 9 also displays the horizons corresponding to the FOMC’s calendar-based forward guidance—the “mid-2013,” “late-2014,” and “mid-2015” language first used in September 2011, January 2012, and September 2012, respectively. The FOMC had indicated that it expected the period of near-zero policy rates to last at least as long as these horizons, so that reasonable liftoff estimates would have to be equal to or larger than these horizons. The figure shows that the modal-path liftoff estimates satisfy this criterion, while the liftoff estimates based on the forward curve do not.

The uncertainty underlying liftoff forecasts is denoted in the graphs by shaded confidence bands based on the interquartile range of the liftoff distribution.²⁶ Uncertainty increased with the lengthening of the expected duration of the ZLB period in 2011 and 2012, and became very substantial. Since then, uncertainty has decreased very significantly toward the end of our sample. Another important point to note is that the liftoff estimates based on forward curves are near and mostly below the lower quartile of the liftoff distribution, which reinforces the finding that these estimates are significantly too early and therefore unsatisfactory.

Our liftoff estimates also demonstrate again the benefit of incorporating macroeconomic information. Figure 9 indicates that model MZ(2) is more accurate than model YZ(3), in the sense that the modal-path forecasts are closer, on average, to the outside estimates. To quantify this we calculate the root-mean-squared difference of model-based vs. alternative estimates. For the Primary Dealer estimates, this difference is 5.2 months for model MZ(2), compared to 5.8 months for YZ(3). For the estimates from Macro Advisers, the distance is 4.4 months for MZ(2) and 4.9 months for YZ(3). From this perspective, the macro-finance model MZ(2) estimates the duration of the ZLB period more accurately than the yields-only model YZ(3). Notably, liftoff estimates that use macroeconomic information are also more stable than those that only use yield-curve information, due to the more volatile nature

²⁶These confidence bands only reflect uncertainty based on shocks to the term structure and not parameter uncertainty. To incorporate such estimation uncertainty, one would need an extensive bootstrap analysis or, ideally, a fully Bayesian estimation framework.

of yield movements relative to macroeconomic variables. Furthermore, the incorporation of macroeconomic information also has the effect that it lowers the uncertainty around the estimated liftoff horizon. Figure 9 shows the narrower confidence bands for model MZ(2) compared with model YZ(3), so adding macroeconomic information improves the precision of the liftoff estimates. We have argued above that macro information can augment the information in the yield curve, which is constrained due to the ZLB, and these estimates confirm this intuition.

In sum, the modal short-rate path provides estimates of policy liftoff that are simple to calculate, correspond to optimal forecasts under a plausible loss function, and perform well empirically. It should be noted that while we use a shadow-rate model to obtain the modal path, one can also use option prices and derive the modal path from the estimated risk-neutral probability densities, which provides largely model-free estimates. However, an advantage of having a parametric model for the risk-neutral distribution of future interest rates is that we are able to simulate the distribution of the liftoff horizon. In this way, we are able to provide an assessment of the uncertainty around point forecasts of liftoff, and to demonstrate the near-optimality of the modal-path liftoff estimate.

4.4 Interpretation of the liftoff horizon

For interpreting the estimated liftoff horizon, we first note that it is highly correlated with our measure of the tightness of the ZLB constraint. The correlation between the ZLB wedge in the ten-year yield and the liftoff estimate based on the modal path over the ZLB period is 0.97 for model YZ(3) and 0.96 for model MZ(2). This makes much intuitive sense: The fact that the period of near-zero policy rates is expected to last is the main reason that yields beyond very short maturities are constrained by the ZLB. Variation in the expected duration should therefore closely correlate with variation in the measured tightness of the ZLB constraint.

The estimated liftoff horizon provides a partial summary of the stance of monetary policy. It captures one of the most important dimensions of monetary policy at the ZLB, namely, how long the policy rate can be expected to stay near zero. For example, based on the liftoff estimates in Figure 9, the stance of policy became increasingly accommodative from 2009 to 2012. In contrast, at the ZLB, the observed short-term interest rate does not convey any information about the policy stance. Figure 9 displays, as vertical lines, the times of key forward guidance announcements by the Federal Reserve, including the changes in the statement language about the appropriate path of the funds rate in 2008 and 2009, the explicit calendar-based forward guidance in 2011 and 2012, and the outcome-based forward guidance announcement in December 2012. Most of these announcements had a noticeable impact and

substantially lengthened the estimated liftoff horizon. This underscores the usefulness of the liftoff horizon as a univariate summary of monetary policy at the ZLB.

Still, the expected duration of the ZLB period can provide only an incomplete measure of monetary policy because it ignores the anticipated pace of tightening after liftoff, which is another important dimension of monetary policy expectations at the ZLB. However, as a practical matter, the anticipated pace of tightening following liftoff appears to have held fairly steady during the recent episode.²⁷ And of course the expected pace of tightening can also be calculated based on the modal path, and serve as a complement to the liftoff measure of policy expectations. Another point is that our liftoff estimates based on the risk-neutral modal path do not capture the effects of bond purchases conducted by the Federal Reserve, which appear to have significantly lowered longer-term interest rates, as documented by [Gagnon et al. \(2011\)](#) and others. To the extent that these purchases resulted in lower interest rates at intermediate maturities mainly through changes in the term premium—i.e., without corresponding changes in the expected policy path and the likely duration of the period of near-zero policy rates—our liftoff estimates would lengthen without any change in the expected policy path. However, in practice, these purchases often went hand in hand with corresponding changes in policy expectations and a longer expected liftoff horizon ([Bauer and Rudebusch, 2013](#)). In any case, our approach could be extended to take explicitly into account changes in risk premia due to Federal Reserve balance sheet policies. In sum, we view our time series of expected policy liftoff as a useful and accurate summary of the stance of monetary policy over recent years.

5 Conclusion

Using U.S. data, we estimate Gaussian affine and shadow-rate DTSMs with a variety of risk factors and elucidate some important issues about U.S. monetary policy at the zero bound. We estimate mean and modal paths for future short rates, taking into account the asymmetric probability distribution of future short rates at a range of projection horizons, and assess the associated dates for monetary policy liftoff from the ZLB. We argue that forecasts of policy liftoff using the term structure should be based on the modal path of future short rates, which is a near-optimal forecast and performs well empirically. We find that the increasing model-implied expectations of liftoff from 2009 to 2012 are very closely matched by private-sector and survey forecasts. Furthermore, the expected duration of the ZLB period can provide a useful measure of the stance of monetary policy and the tightness of the ZLB. Finally, we

²⁷For example, the median of the Primary Dealer survey’s modal forecasts has fairly consistently indicated about a 2-2.5 percentage point increase in the fed funds rate over the two years following liftoff.

document the benefits of including macroeconomic information in shadow-rate models, which improves inference at the ZLB about future monetary policy.

Our work can be extended in several promising directions. A natural application of our framework would be the recent period of near-zero interest rates in euro countries and the United Kingdom, as well as the long period of low interest rates in Japan over the past 20 years. Regarding our modeling framework, one could evaluate and impose restrictions on the risk pricing, which would tighten the link between the cross-sectional and time series dynamics of the risk factors (Joslin et al., 2012; Bauer, 2011). Such restrictions would improve model parsimony, alleviate some of the statistical issues related to the highly persistent nature of interest rates, and lead to more precise inference about short-rate expectations and policy liftoff under the real-world \mathbb{P} -measure. Alternative techniques for pinning down \mathbb{P} -measure expectations more accurately that could be incorporated in our framework include bias correction (Bauer et al., 2012) and the inclusion of survey-based interest rate forecasts (Kim and Orphanides, 2012). Finally, the use of a Bayesian inference framework holds some promise for estimation of shadow-rate models: On a practical level, the use of modern Markov chain Monte Carlo methods could improve the computational efficiency of the estimation. More fundamentally, such a framework would allow for an accurate description of the uncertainty around current and future shadow rates and liftoff estimates by taking into account not only shocks to the dynamic system but also parameter uncertainty and possibly model uncertainty.

References

- Bauer, Michael D.**, “Bayesian Estimation of Dynamic Term Structure Models under Restrictions on Risk Pricing,” Working Paper 2011-03, Federal Reserve Bank of San Francisco November 2011.
- and **Glenn D. Rudebusch**, “The Signaling Channel for Federal Reserve Bond Purchases,” Working paper 2011-21, Federal Reserve Bank of San Francisco April 2013.
- , – , and **Jing Cynthia Wu**, “Correcting Estimation Bias in Dynamic Term Structure Models,” *Journal of Business and Economic Statistics*, July 2012, 30 (3), 454–467.
- Black, Fischer**, “Interest Rates as Options,” *Journal of Finance*, December 1995, 50 (5), 1371–1376.
- Bomfim, Antulio N.**, “Interest Rates as Options: Assessing the markets’ view of the liquidity trap,” Finance and Economics Discussion Series 2003-45, Federal Reserve Board of Governors 2003.
- Bullard, James**, “Shadow Interest Rates and the Stance of U.S. Monetary Policy,” speech at the Annual Corporate Finance Conference, Olin Business School, Washington University in St. Louis November 8 2012.
- Christensen, Jens H. E. and Glenn D. Rudebusch**, “Estimating Shadow-Rate Term Structure Models with Near-Zero Yields,” Working Paper 2013-07, Federal Reserve Bank of San Francisco 2013.
- Cochrane, John H. and Monika Piazzesi**, “Decomposing the Yield Curve,” unpublished manuscript 2008.
- Duffee, Gregory R.**, “Term Premia and Interest Rate Forecasts in Affine Models,” *Journal of Finance*, 02 2002, 57 (1), 405–443.
- and **Richard H. Stanton**, “Estimation of Dynamic Term Structure Models,” *Quarterly Journal of Finance*, 2012, 2 (2).
- Fisher, Irving**, *Appreciation and Interest*, August M. Kelley Bookseller, New York, 1896.
- Gagnon, Joseph E. and Brian Sack**, “Monetary Policy with Abundant Liquidity: A New Operating Framework for the Federal Reserve,” Policy Brief, Peterson Institute for International Economics 2013.

- Gagnon, Joseph, Matthew Raskin, Julie Remache, and Brian Sack**, “The Financial Market Effects of the Federal Reserve’s Large-Scale Asset Purchases,” *International Journal of Central Banking*, March 2011, 7 (1), 3–43.
- Gorovoi, Viatcheslav and Vadim Linetsky**, “Black’s Model of Interest Rates as Options, Eigenfunction Expansions, and Japanese Interest Rates,” *Mathematical Finance*, January 2004, 14 (1), 49–78.
- Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright**, “The U.S. Treasury yield curve: 1961 to the present,” *Journal of Monetary Economics*, 2007, 54 (8), 2291–2304.
- Ichiue, Hibiki and Yoichi Ueno**, “Equilibrium Interest Rates and the Yield Curve in a Low Interest Rate Environment,” Working Paper 2007-E-18, Bank of Japan 2007.
- and —, “Monetary Policy and the Yield Curve at Zero Interest,” unpublished manuscript 2012.
- and —, “Estimating Term Premia at the Zero Bound: An Analysis of Japanese, US, and UK Yields,” Working Paper 2013-E-8, Bank of Japan 2013.
- Joslin, Scott, Anh Le, and Kenneth J. Singleton**, “Why Gaussian Macro-Finance Term Structure Models Are (Nearly) Unconstrained Factor-VARs,” *Journal of Financial Economics*, September 2013, 109, 604–622.
- , **Kenneth J. Singleton, and Haoxiang Zhu**, “A New Perspective on Gaussian Dynamic Term Structure Models,” *Review of Financial Studies*, March 2011, 24 (3), 926–970.
- , **Marcel Pribsch, and Kenneth J. Singleton**, “Risk Premiums in Dynamic Term Structure Models with Unspanned Macro Risks,” working paper October 2012.
- Kim, Don H. and Athanasios Orphanides**, “Term Structure Estimation with Survey Data on Interest Rate Forecasts,” *Journal of Financial and Quantitative Analysis*, February 2012, 47 (1), 241–272.
- and **Kenneth J. Singleton**, “Term Structure Models and the Zero Bound: An Empirical Investigation of Japanese Yields,” *Journal of Econometrics*, September 2012, 170 (1), 32–49.
- and **Marcel A. Pribsch**, “Estimation of Multi-Factor Shadow-Rate Term Structure Models,” Unpublished manuscript, Federal Reserve Board 2013.

- Krippner, Leo**, “Modifying Gaussian term structure models when interest rates are near the zero lower bound,” Discussion Paper 2012/02, Reserve Bank of New Zealand 2012.
- , “Measuring the stance of monetary policy in zero lower bound environments,” *Economics Letters*, January 2013, *118* (1), 135–138.
- Linetsky, Vadim**, “Computing hitting time densities for CIR and OU diffusions: Applications to mean-reverting models,” *Journal of Computational Finance*, 2004, *7*, 1–22.
- Piazzesi, Monika**, “Affine Term Structure Models,” in Yacine Ait-Sahalia and Lars Hansen, eds., *Handbook of Financial Econometrics*, Vol. 1, Elsevier, 2010, chapter 12.
- Pribsch, Marcel A.**, “Computing Arbitrage-Free Yields in Multi-Factor Gaussian Shadow-Rate Term Structure Models,” Finance and Economics Discussion Series, Federal Reserve Board 2013.
- Rudebusch, Glenn D.**, “Monetary Policy Inertia: Fact or Fiction?,” *International Journal of Central Banking*, December 2006, *2* (4), 85–135.
- , “The Fed’s Monetary Policy Response to the Current Crisis,” *FRBSF Economic Letter*, May 2009.
- , “Macro-Finance Models of Interest Rates and the Economy,” *The Manchester School*, September 2010, *78* (s1), 25–52.
- Swanson, Eric T. and John C. Williams**, “Measuring the Effect of the Zero Lower Bound On Medium- and Longer-Term Interest Rates,” Working Paper 2012-02, Federal Reserve Bank of San Francisco 2012.
- Ueno, Yoichi, Naohiko Baba, and Yuji Sakurai**, “The Use of the Black Model of Interest Rates as Options for Monitoring the JGB Market Expectations,” Working Paper 2006-E-15, Bank of Japan 2006.
- Wu, Jing Cynthia and Fan Dora Xia**, “Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound,” Research Paper, Chicago Booth October 2013.

A Affine bond pricing

Under the assumptions of Section 2.1, bond prices are exponentially affine functions of the pricing factors:

$$P_t^m = e^{\mathcal{A}_m + \mathcal{B}_m X_t},$$

and the loadings $\mathcal{A}_m = \mathcal{A}_m(\mu^Q, \phi^Q, \delta_0, \delta_1, \Sigma)$ and $\mathcal{B}_m = \mathcal{B}_m(\phi^Q, \delta_1)$ follow the recursions

$$\begin{aligned}\mathcal{A}_{m+1} &= \mathcal{A}_m + (\mu^Q)' \mathcal{B}_m + \frac{1}{2} \mathcal{B}_m' \Sigma \Sigma' \mathcal{B}_m - \delta_0 \\ \mathcal{B}_{m+1} &= (\phi^Q)' \mathcal{B}_m - \delta_1\end{aligned}$$

with starting values $\mathcal{A}_0 = 0$ and $\mathcal{B}_0 = 0$. Model-implied yields are determined by $y_t^m = -m^{-1} \log P_t^m = A_m + B_m X_t$, with $A_m = -m^{-1} \mathcal{A}_m$ and $B_m = -m^{-1} \mathcal{B}_m$. Risk-neutral yields, the yields that would prevail if investors were risk-neutral, can be calculated using

$$\tilde{y}_t^m = \tilde{A}_m + \tilde{B}_m X_t, \quad \tilde{A}_m = -m^{-1} \mathcal{A}_m(\mu, \phi, \delta_0, \delta_1, \Sigma), \quad \tilde{B}_m = -m^{-1} \mathcal{B}_m(\phi, \delta_1).$$

Risk-neutral yields reflect policy expectations over the life of the bond, $m^{-1} \sum_{h=0}^{m-1} E_t r_{t+h}$, plus a convexity term. The yield term premium is defined as the difference between actual and risk-neutral yields, $ytp_t^m = y_t^m - \tilde{y}_t^m$.

B Moments for the shadow-rate model

Of particular interest then are the conditional moments of the future shadow rate at horizon h :

$$\begin{aligned}\tilde{\mu}_t^h &= E(s_{t+h} | X_t) = \delta_0 + \delta_1' E(X_{t+h} | X_t) = \delta_0 + \delta_1' [(1 - \phi^h) E(X_t) + \phi^h X_t] \\ (\tilde{\sigma}^h)^2 &= Var(s_{t+h} | X_t) = \delta_1' Var(X_{t+h} | X_t) \delta_1 = \delta_1' \left[\sum_{i=0}^{h-1} \Psi_i \Sigma \Sigma' \Psi_i' \right] \delta_1,\end{aligned}$$

where Ψ_i is the i -th coefficient matrix in the Wold representation for X_t (for a VAR(1), $\Psi_i = \phi^i$), and the conditional variance $(\tilde{\sigma}^h)^2$ is time-invariant because X_t is homoskedastic. Importantly, note that these moments are identical to the conditional moments of the future short rate in the affine model,.

The conditional mean of the future short rate is:

$$\begin{aligned}E(r_{t+h} | X_t) &= P(s_{t+h} > 0) E(s_{t+h} | X_t, s_{t+h} > 0) \\ &= N(\tilde{\mu}_t^h / \tilde{\sigma}^h) \left[\tilde{\mu}_t^h + \tilde{\sigma}^h \frac{n(-\tilde{\mu}_t^h / \tilde{\sigma}^h)}{1 - N(-\tilde{\mu}_t^h / \tilde{\sigma}^h)} \right] \\ &= N(\tilde{\mu}_t^h / \tilde{\sigma}^h) \tilde{\mu}_t^h + \tilde{\sigma}^h n(-\tilde{\mu}_t^h / \tilde{\sigma}^h),\end{aligned}$$

where $N(\cdot)$ is the cumulative distribution function of the standard normal distribution, and

$n(\cdot)$ is its density function.²⁸ The second line follows from well-known results about truncated normal distributions. Using the formulas given above, one can easily calculate the path of expected future short rates for any given value of X_t . By using the \mathbb{Q} -measure parameters instead of μ and ϕ to calculate the conditional moments, the forward curve can also be obtained analytically. Note that the above results are for the special case $\underline{r} = 0$, but are easy to generalize (see, for example [Wu and Xia, 2013](#)).

C Monte Carlo bond pricing

To obtain bond prices and yields for the shadow-rate model, we resort to Monte Carlo simulation. For given values of the risk factors, the price of a bond with maturity m is

$$P_t^m = E_t^{\mathbb{Q}} \left[\exp \left(- \sum_{i=0}^{m-1} r_{t+i} \right) \right].$$

Since this expectation cannot be found analytically, we approximate it by simulating $M = 500$ paths of the risk factors of length m , where each sample path is obtained using the \mathbb{Q} -measure VAR in equation (2), starting from the given initial value X_t . Using the ZLB short-rate equation (5) we obtain the sampled paths for the short rate. Denote the value of the short rate in simulation j at time $t + i$ by $r_{t+i}^{(j)}$. The approximate bond price is given by

$$\tilde{P}_t^m = M^{-1} \sum_{j=1}^M \exp \left(- \sum_{i=0}^{m-1} r_{t+i}^{(j)} \right).$$

We use antithetic sampling to obtain the shock sequences, by taking the shock sequence for replication j as the negative of the shock sequence for replication $j - 1$. This improves the accuracy of the approximation for any given M , because it introduces negative dependence between pairs of replications.

To assess the accuracy of our Monte Carlo simulation for bond prices, we compare Monte Carlo bond prices and analytical bond prices for the affine model. The Monte Carlo bond prices are obtained as described above, with the only difference being that we use the affine short-rate equation (1) instead of the ZLB short-rate equation (5). We use our estimated parameters, together with the risk factors filtered using the Kalman filter in the affine model, and consider the accuracy of Monte Carlo yields over the full sample from 1985 to 2012.

Table C.1 shows the approximation error, measured as the root mean-squared difference between exact affine yields and approximate Monte Carlo yields, in basis points. The approximation error is miniscule at short and medium maturities, and rarely exceeds one basis point. Evidently, the approximation of model-implied yields using our Monte Carlo method, even for a moderate amount of replications ($M = 500$), is very accurate.

²⁸This formula is also given in the appendix of [Kim and Singleton \(2012\)](#).

Table C.1: Accuracy of Monte Carlo yields

Model	Total	3m	6m	1y	2y	3y	5y	7y	10y
YA(2)	0.14	0.00	0.00	0.01	0.04	0.08	0.06	0.11	0.35
YA(3)	0.85	0.00	0.00	0.01	0.04	0.20	0.72	1.17	1.96
MA(1)	0.47	0.00	0.00	0.01	0.04	0.11	0.42	0.75	1.01
MA(2)	0.38	0.00	0.00	0.01	0.03	0.07	0.48	0.78	0.56

Notes: Root mean-squared difference, in basis points, between analytical and Monte Carlo yields for affine models. Sample period: January 1985 to December 2012.

D Affine vs. shadow-rate models over the 1985–2007 subsample

Here we provide evidence that the implications of the affine and shadow-rate models on the estimation subsample 1985 to 2007 are very similar. This is part of the justification for using in our shadow-rate models those parameters that are estimated for the affine models.

We use the parameters estimated for the affine model and calculate model-implied yields using both the affine and shadow-rate models over the pre-2008 sample period. For each model we first filter the latent risk factors appropriately from the observed yields, using the Kalman filter for the affine models, and the Extended Kalman filter for the shadow-rate models.

Table D.1 compares the model fit of the affine and shadow-rate models over the estimation subsample. In the first panel, it shows the RMSEs in basis points, to assess whether affine and shadow-rate models differ in terms of cross-sectional fit over this subsample. The RMSEs are very similar for each pair of affine and shadow-rate models, which demonstrates that the cross-sectional fit is essentially identical. The second panel of the table shows the discrepancy between yields implied by the affine and shadow-rate models, again measured as a root mean-squared difference. We see that the discrepancy is very small—model-implied yields from affine and shadow-rate models typically differ by less than one basis point.

These results show that the affine model and the shadow-rate model have close to identical implications on the estimation subsample. Therefore, it appears that the assumption is reasonable that maximum likelihood estimates of the model parameters are interchangeable between each affine model and its ZLB counterpart, if estimation is constrained to the pre-2008 subsample.

Table D.1: Comparison of affine and shadow-rate model from 1985 to 2007

Model	Total	3m	6m	1y	2y	3y	5y	7y	10y
<i>Cross-sectional fit</i>									
YA(2)	13.98	17.73	11.22	19.60	15.98	11.25	9.06	8.91	13.97
YZ(2)	13.96	17.77	11.27	19.59	16.03	11.23	8.85	8.88	13.86
YA(3)	6.07	6.36	9.25	9.29	2.89	4.44	4.17	2.76	5.46
YZ(3)	6.10	6.33	9.22	9.39	2.86	4.56	4.25	2.72	5.45
MA(1)	28.77	40.85	33.80	22.86	12.08	13.55	22.93	30.78	38.51
MZ(1)	28.75	40.85	33.71	22.90	12.06	13.56	22.95	30.76	38.52
MA(2)	9.56	13.74	10.26	10.31	9.98	8.45	4.15	5.15	10.80
MZ(2)	9.58	13.64	10.39	10.36	9.98	8.44	4.15	5.19	10.88
<i>Discrepancy between affine and ZLB fitted yields</i>									
YA(2)/YZ(2)	0.41	0.26	0.19	0.08	0.20	0.37	0.45	0.13	0.93
YA(3)/YZ(3)	0.28	0.14	0.13	0.25	0.24	0.37	0.42	0.24	0.31
MA(1)/MZ(1)	0.47	0.34	0.36	0.36	0.32	0.23	0.21	0.61	0.89
MA(2)/MZ(2)	0.58	0.58	0.32	0.12	0.57	0.70	0.34	0.34	1.10

Notes: Cross-sectional fit, measured by root mean-squared errors of model-implied yields, and discrepancy between affine and ZLB fitted yields, measured by root mean-squared differences, all measured in basis points. Sample period: January 1985 to December 2007.

Table 1: Cross-sectional fit

Model	Total	3m	6m	1y	2y	3y	5y	7y	10y
<i>Full sample</i>									
	Total	3	6	12	24	36	60	84	120
YA(2)	13.9	16.9	10.7	17.8	16.1	13.2	8.9	8.8	15.4
YZ(2)	13.4	16.5	10.9	18.0	15.2	12.0	8.6	8.5	14.1
YA(3)	6.8	7.2	8.6	10.1	3.0	6.0	5.8	2.8	7.2
YZ(3)	6.7	7.0	8.7	10.0	3.4	6.0	6.0	2.9	6.7
MA(1)	60.7	78.7	72.3	55.7	19.4	19.9	52.8	69.1	81.1
MZ(1)	28.7	37.0	31.2	24.7	22.0	20.4	21.9	29.0	37.7
MA(2)	12.1	17.3	9.8	15.1	13.0	9.8	8.1	7.7	12.2
MZ(2)	10.2	13.5	10.4	12.4	10.6	9.1	6.3	5.6	11.0
<i>ZLB subsample</i>									
	Total	3	6	12	24	36	60	84	120
YA(2)	12.7	13.0	7.9	5.7	16.2	17.6	7.4	8.2	18.3
YZ(2)	9.3	8.5	9.2	8.7	9.5	11.9	6.7	5.9	11.9
YA(3)	7.8	6.9	4.0	11.2	2.7	8.6	9.3	2.8	10.8
YZ(3)	7.4	6.2	5.8	10.3	4.5	7.9	9.7	3.4	8.6
MA(1)	130.9	166.2	156.8	123.3	37.7	36.7	115.8	151.3	175.1
MZ(1)	29.2	11.4	18.5	31.6	45.1	38.0	16.0	20.5	34.8
MA(2)	19.2	27.5	7.0	27.2	22.1	13.2	16.6	14.3	16.2
MZ(2)	11.5	9.8	10.9	17.8	12.8	10.0	11.1	6.9	10.1

Notes: Root mean-squared errors of model-implied yields in basis points. Full sample: January 1985 to December 2012. ZLB subsample: January 2008 to December 2013.

Table 2: Violations of the ZLB

Model	Forward rates		Short-rate expectations	
	frequency	avg. length	frequency	avg. length
YA(2)	31	1.5	37	3.5
YA(3)	17	4.5	22	8.6
MA(1)	62	10.5	61	27
MA(2)	46	10.1	61	21.2


Notes: Number of months, between January 2008 and December 2013, in which some forward rates (column two) or short-rate expectations (column four) drop below zero, and the average length (in months) of horizon over which the forward curve/short-rate path stays negative.

Table 3: Out-of-sample forecast accuracy at the ZLB

Model	6m	12m	18m	24m
YA(2)	12.4	30.5	49.7	69.1
YZ(2)	11.3	16.7	28.2	44.1
YA(3)	10.2	21.6	48.6	82.4
YZ(3)	9.3	13.6	32.4	58.7
MA(1)	11.5	10.7	9.7	9.5
MZ(1)	11.5	10.7	9.6	7.5
MA(2)	11.5	10.7	8.5	19.9
MZ(2)	11.0	10.3	8.1	11.3


Notes: RMSEs in basis points for out-of-sample forecasts (under \mathbb{P} -measure) of the future three-month T-bill rate at various forecast horizons. Forecast period: December 2008 to December 2011.

Figure 1: Yields and macroeconomic data




Notes: Top panel shows three selected yields, for the three-month, two-year, and ten-year maturities. Bottom panel shows the macroeconomic data. Sample period: January 1985 to December 2012

Figure 2: Affine model probabilities of negative future short rates




Notes: Model-implied real-world (\mathbb{P}) probabilities of negative future short-term interest rates at horizons of six months, one year, and two years. Shaded areas correspond to NBER recessions. Sample period: January 2000 to December 2012

Figure 3: Shadow short rates




Notes: Shadow short rates implied by shadow-rate models and the three-month T-bill rate. Shaded areas correspond to NBER recessions. Top panel shows shadow short rates for each model, with $\underline{r} = 0$. Bottom panel shows shadow short rates for model YZ(3) with \underline{r} equal to 0, 5, 10, 15, 20, and 25 basis points. Sample period: January 1985 to December 2012.

Figure 4: Distribution of future short rate




Notes: Densities of future shadow rate and of future short rate, based on model MZ(2), at horizon of 48 months, on December 31, 2012. Vertical lines show the mode and mean of the distribution of the future short rate.

Figure 5: Forward rates, shadow forward rates, and modal paths




Notes: Forward rates, shadow forward rates, and modal path on June 30, 2011, and on December 31, 2012.

Figure 6: Fitted and shadow yield curves




Notes: Actual, fitted, and shadow yield curves on June 30, 2011, and on December 31, 2012.

Figure 7: Ten-year yield




Notes: The top panel shows the observed ten-year yield and the corresponding shadow yield, estimated from models YZ(3) and MZ(2). The bottom panel shows the difference between these two yields, for both models.

Figure 8: Distribution of liftoff horizon



Notes: Approximate densities for the distribution of the future liftoff horizon, for December 31, 2012

Figure 9: Alternative liftoff estimates



Notes: Estimated horizon (in months) until policy liftoff from the ZLB, based on the forward curve and the modal path, as well as the median and the interquartile range (gray-shaded) of the liftoff distribution. Also shown is the predicted duration of the ZLB period from the Primary Dealer survey (median response) and from Macroeconomic Advisers. Vertical lines indicate FOMC statements that contained forward guidance (FG) announcements, and crosses indicate the