Monetary Policy Expectations at the Zero Lower Bound

Michael D. Bauer,
Federal Reserve Bank of San Francisco

Glenn D. Rudebusch,
Federal Reserve Bank of San Francisco

May 2015

Suggested citation:

Bauer, Michael D. and Glenn D. Rudebusch. 2015. "Monetary Policy Expectations at the Zero Lower Bound." Federal Reserve Bank of San Francisco Working Paper 2013-18. http://www.frbsf.org/economic-research/publications/working-papers/wp2013-18.pdf

Monetary Policy Expectations at the Zero Lower Bound[*]

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

May 21, 2015

Abstract

We show that conventional dynamic term structure models (DTSMs) estimated on recent U.S. data severely violate the zero lower bound (ZLB) on nominal interest rates and deliver poor forecasts of future short rates. In contrast, shadow-rate DTSMs account for the ZLB by construction, capture the resulting distributional asymmetry of future short rates, and achieve good forecast performance. These models provide more accurate estimates of the most likely path for future monetary policy—including the timing of policy liftoff from the ZLB and the pace of subsequent policy tightening. We also demonstrate the benefits of including macroeconomic factors in a shadow-rate DTSM when yields are constrained near the ZLB.

Keywords: dynamic term structure models, shadow rates, policy liftoff, macro-finance
JEL Classifications: E43, E44, E52

1   Introduction

Divining the path of future monetary policy has been of special interest during the Great Recession and its aftermath. 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 representation in finance used to extract such short-rate expectations (e.g., Piazzesi, 2010). However, while these models have provided good empirical representations of yield curves in the past, they may be ill-suited to represent the dynamics of recent near-zero interest rates that have prevailed in many 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.[1]

The fact that Gaussian affine DTSMs ignore the zero lower bound (ZLB) was of little consequence when interest rates were well above zero. However, as nominal interest rates have fallen to near zero, the lack of an appropriate nonnegativity restriction in conventional models has become a conspicuous theoretical deficiency. This paper presents evidence showing that the theoretical failure of standard Gaussian affine DTSMs to account for the ZLB has been an important practical deficiency in recent years in terms of fit, point and distributional forecasting ability, and accuracy of estimated monetary policy expectations. Our benchmark for comparison is an alternative model based on the shadow-rate concept proposed by Black (1995).[2] The shadow-rate representation replaces the affine short-rate specification of standard DTSMs with an identical affine process for an unobserved shadow short rate. The observed short rate is set equal to this shadow short rate when it is positive; otherwise, it is set to zero (or some other near-zero minimum value). 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. In particular, a key advantage of affine DTSMs—namely, analytical affine bond pricing—is lost. Instead, numerical solution methods are required, so the calculation of model-implied interest rates is computationally quite intensive.

---

[1]The value of the lower bound on nominal interest rates is not precisely zero. As discussed below, the effective lower bound depends on a variety of institutional factors including the size of costs associated with storing, transferring, and spending large amounts of currency. Indeed, recent nominal European yields have been slightly negative. Still, for convenience, we will describe this constraint as a *zero* lower bound even though our model and analysis will allow for a non-zero lower bound.

[2]At the ZLB, one could also consider other alternatives, such as stochastic-volatility models with square-root processes or Gaussian quadratic models, but the shadow-rate model has the advantage of matching the canonical Gaussian DTSM when interest rates are away from the ZLB.

The focus of our paper is the estimation of monetary policy expectations at the ZLB. To this end, we consider both yields-only and macro-finance shadow-rate models, where the latter includes measures of economic activity and inflation as risk factors. There is now a sizable literature arguing that a joint macro-finance approach is a very productive research avenue for term structure modeling (e.g., Rudebusch, 2010), but this paper is the first to include macroeconomic factors into a shadow-rate model. We show that when the nominal term structure is constrained by the ZLB, the addition of macroeconomic variables to the DTSM information set is useful for inference about the future evolution of the yield curve. Intuitively, the ZLB limits the information content of the yield curve because its short end is pinned at zero. In such a situation, macro variables provide important additional information for forecasting future yields, particularly for predicting how long the policy rate will remain near zero.[3]

We begin our analysis with an evaluation of affine and shadow-rate models during the past near-decade of very low interest rates in the United States. Given the close proximity of interest rates to the ZLB during this period, we find that shadow-rate DTSMs provide a statistically significant and economically relevant improvement in fit and forecasting performance compared with standard Gaussian affine DTSMs. Affine models frequently violate the ZLB and produce substantial estimated probabilities of negative future short rates. Shadow-rate models avoid such violations by construction and fit the cross section of yields substantially better than affine models. Most notably, affine models cannot capture 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 quite inaccurate short-rate forecasts at the ZLB. In contrast, we document that shadow-rate models can accurately forecast prolonged near-zero policy rates in an out-of-sample forecast exercise.

Shadow-rate models account for the substantial asymmetry in the distribution of future short rates during periods of near-zero policy rates. This feature is especially valuable for assessing monetary policy expectations embedded in the yield curve at the ZLB. For example, one key question is how to estimate the anticipated timing of the liftoff of the policy rate from the ZLB. A common approach among financial market researchers and investors is to use the horizon at which forward rates cross a given threshold, say 25 basis points, as an estimate of the expected date of liftoff. But forward rates correspond to (risk-neutral) expectations of future short rates, and this practice of using this *mean path* to estimate liftoff is problematic because it ignores the asymmetry of the distribution of future short rates near the ZLB.

---

[3]The value of a macro-finance approach is also consistent with the many central bank statements that have stressed that the timing of liftoff from the ZLB is dependent on the flow of incoming macroeconomic data.

Instead, to assess the expected time of short-rate liftoff, one needs to consider the *modal path*—the most-likely path for future short rate rates. In a shadow-rate model, this modal path corresponds to the expected path of future shadow short rates (or to zero when the expected shadow short rate is negative). A comparison of the mean path to the modal path reveals how tightly the ZLB constraint is binding. The difference between the two paths, which we term a "ZLB wedge," 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. We use the ZLB wedge between the ten-year yield and the corresponding shadow yield as a measure of the tightness of the ZLB constraint, and document that it increased substantially over the period from 2009 to 2012, and then gradually decreased over 2013 and 2014, a period when macroeconomic conditions improved notably.

To measure monetary policy expectations at the ZLB, we focus on two key metrics: the time until liftoff and the subsequent pace of tightening. In the first, the expected date at which the modal short-rate path escapes from near zero provides a forecast of the time until liftoff that is optimal under an absolute-error loss function. We also compute the full forecast distribution of the liftoff horizon in order to verify the modal-path-based liftoff estimate and to obtain interval forecasts for liftoff. We find that model-based liftoff estimates based on a macro-finance yield curve model closely accord with private-sector forecasts of the timing of monetary policy liftoff, and are consistent with the Federal Open Market Committee's (FOMC) calendar-based forward guidance. Overall, the liftoff horizon can therefore serve as a useful univariate summary of monetary policy at the ZLB.[4] Our second metric, the initial pace of policy tightening is calculated as the expected cumulative increase in the modal short rate path during the first two years after liftoff. Our macro-finance term structure model forecasts a much more gradual increase in the policy rate than in previous policy tightening cycles, which is consistent with statements by Federal Reserve policymakers.

Overall, 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. 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

---

[4]In contrast to our measures of monetary policy expectations, 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, 2014), are highly sensitive to model specification the exact data at the short end of the yield curve. Their lack of robustness raises a warning flag about using shadow short rates as a measure of monetary policy.

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. Using Japanese yield curve data, Kim and Singleton (2012) estimate two-factor models and demonstrate the good performance of shadow-rate models compared to alternatives, and Christensen and Rudebusch (2013) document the sensitivity of shadow-rate estimates to model specification in estimated one-, two- and three-factor models. Several other studies have considered the recent U.S. experience, including Krippner (2015), Ichiue and Ueno (2013), and Christensen and Rudebusch (2015). 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 various aspects of monetary policy expectations at the ZLB using the modal path.

The paper is structured as follows. Section 2 lays out our modeling 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 5 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 6 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 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.,

$$r_t = \delta_0 + \delta_1' X_t. \tag{1}$$

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

4

factors follow a Gaussian vector autoregression (VAR),

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

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,

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

where ε_t is an *i.i.d.* standard normal random vector under $\mathbb{P}$.[5] 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

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

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 closely similar to our affine models except that the affine short-rate equation (1) is replaced by a shadow-rate specification:

$$r_t = \max(s_t, r_{min}), \quad s_t = \delta_0 + \delta_1' X_t. \tag{5}$$

The shadow short rate, s_t, is modeled as affine Gaussian, exactly as the short rate in affine models. Equation (5) ensures that the short rate and all other model-implied interest rates cannot go below r_{min}. Black (1995) set $r_{min} = 0$, and this is our choice as well. This ZLB on

---

[5]That is, as is standard, forecasts for the state variables can be calculated under two different probability measures: the real-world $\mathbb{P}$ measure (also know as the physical or historical or objective measure) and the risk-neutral $\mathbb{Q}$ measure that investors use to value assets because of their risk aversion. Specifically, investors value assets just as a risk-neutral agent would if that agent believed that the dynamics of state variables were characterized by the $\mathbb{Q}$ measure.

nominal interest rates is typically motivated by the presence of physical currency. 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 r_{min}. 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 r_{min}. Different authors have made alternative choices, e.g., Wu and Xia (2014) set $r_{min} = 25$ basis points, and Kim and Priebsch (2013) treat r_{min} as a parameter and estimate it. Below, we investigate the sensitivity of our results to different choices of r_{min}.

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 corresponding affine 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.

A shadow-rate model does not lead to closed-form solutions for yields and bond prices so that the need arises for approximative solution methods. One appealing approach was proposed by Krippner (2014), who approximates forward rates in a shadow-rate context as the sum of shadow forward rates and an option effect, both of which are available analytically. This method is fast and convenient, and has successfully been used in other studies.[6] However, Priebsch (2013) proposes a solution method that appears significantly more accurate than Krippner's approach, at the expense of being somewhat more computationally intensive. In this paper, we use the Priebsch method, which we adapt to the discrete-time context, to approximate bond prices in our shadow-rate models. The first and second moments of the future short rate in the shadow-rate model are in Appendix C, and details on how to use the moments to approximate bond prices are in Appendix D.

---

[6]Christensen and Rudebusch (2015) perform the necessary derivations for their affine Nelson-Siegel model and in this way apply this approach empirically. Wu and Xia (2014) independently derive an approximation for bond prices in a discrete-time shadow-rate model that is equivalent to the method proposed by Krippner, as shown in Krippner (2015). We show how the Krippner approach relates to moments of the shadow-rate model 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 use yields-only affine and shadow-rate models with three risk factors, denoting the affine model by YA(3) and the shadow-rate model by 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 loadings of the first N principal components of observed yields. In the affine model, the risk factors are linear combinations of model-implied yields—they correspond to level, slope, and curvature of the yield curve.[7] In the shadow-rate model, 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.[8]

Macroeconomic variables are likely to be particularly informative when the yield curve is constrained by the ZLB. To investigate this, 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. (2013b). We use affine and shadow-rate models with two ($L = 2$) yield factors in addition to the two macro factors, and denote our models by $MA(2)$ and $MZ(2)$.[9] As in the case of yields-only models, the yield factors are linear combinations of (model-implied/shadow) yields, with weights corresponding to principal components of observed yields.

In our macro-finance models, the macroeconomic variables are spanned by the yield curve.[10] An alternative is to use models with unspanned macro risks as in Joslin et al. (2014), where the current short rate and yield curve depends only on the yield factors. Here we maintain the assumption that macroeconomic conditions directly 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 Federal Open Market Committee (FOMC) that the short rate will be based on the unemployment and inflation rates. For further discussion of this issue and a defense of spanned macro-finance DTSMs see Bauer and Rudebusch (2015).

---

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

[8]Shadow yields can be calculated by using the risk factors of a shadow-rate model in combination with affine loadings.

[9]We have also considered models with two macro factors and only one yields factor, as in Joslin et al. (2013b). We found that these models are not able to accurately fit observed yields, and hence focus on models with two yields factors (which were also used in Joslin et al., 2013a).

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

2.4 Data, measurement error, and estimation

Our data consist of monthly observations of interest rates and macroeconomic variables from January 1985 to December 2014. For the short end of the yield curve, we use three-month and six-month T-bill rates.[11] 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).[12] 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 displays these data on yields (for three maturities) and macroeconomic variables.

Denote the vector of $J = 8$ model-implied yields by Y_t. For the affine models, we have $Y_t = A + B X_t$, with J-vector A and $J \times N$-matrix B containing the usual affine loadings. The observed bond yields used for estimation and inference are $\hat{Y}_t = Y_t + e_t$, where e_t is a vector of *iid* normal measurement error. We include measurement error on yields because an N-dimensional factor model cannot perfectly price $J > N$ yields. In line with the large literature on macro-finance DTSMs, we do not include measurement errors on macro variables.[13]

Estimation of the affine models is standard, both for yields-only and macro-finance models. In the estimation we assume that the yield factors are observed, as in Joslin et al. (2011) and Joslin et al. (2013b), so that μ and ϕ can be obtained using least squares and the remaining parameters are found by maximizing the likelihood function for given VAR parameters. This is particularly advantageous for macro-finance models, which have many parameters. Our estimation method delivers fast and reliable maximum likelihood estimates.[14] Instead of estimating the shadow-rate models, we take a different approach in this paper. We estimate parameters *only for the affine models* $YA(3)$ and $MA(2)$, using the pre-ZLB sample ending in December 2007. Over this period, affine and shadow-rate models are essentially indistinguishable, because yields are sufficiently far from the ZLB. Then, we use the same pre-ZLB

---

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

[12]These yields are available at http://www.federalreserve.gov/pubs/feds/2006/200628/200628abs.html.

[13]We 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., 2013b). Note that our affine macro-finance model corresponds to the TS^f specification in Joslin et al. (2013b), with the difference that we use $\mathcal{L} = 2$ yield factors instead of just one.

[14]Denote by W the $\mathcal{L} \times J$ matrix with the principal component loadings. The assumption that X_t is observable, i.e., that the $\mathcal{L}$ linear combination of yields in W are priced exactly by the model, implies $X_t = W\hat{Y}_t = WY_t$ and $We_t = 0$ so that there are effectively only $J - \mathcal{L}$ independent measurement errors.

estimated parameters in the affine models and the corresponding shadow-rate models $YZ(3)$ and $MZ(2)$, and apply the models to the full sample period until December 2014.[15]

Hence, we use shadow-rate and affine models with the same parameters estimated from the pre-ZLB sample to answer questions regarding the full sample. One important reason for this approach is that estimation of shadow-rate models incurs very high computational costs, as it requires both numerical bond pricing and nonlinear filtering. This is particularly problematic for the macro-finance models that we are interested in due to their many parameters. In contrast, estimation of affine models is extremely fast and much more reliable. While one may be concerned about using parameters in the shadow-rate models that are not the maximum likelihood estimates, we show in Section 3 that in spite of this, shadow-rate models in fact perform very well in our data along several dimensions. They perform much better than affine models, although using the affine-model parameters in our comparison would of course tend to give the advantage to the affine models. Another advantage to holding the parameters the same for each pair of affine and shadow-rate models is that the effects of the ZLB constraint when comparing each pair can be clearly seen. More generally, we view our use of only pre-ZLB data for estimation as a defensible compromise. As a practical matter, it is an open question as to whether the pre-ZLB data or the full sample best approximates the expanding realtime information set that the bond market investors actually use for pricing. Certainly, early on in the ZLB period, using the pre-ZLB sample is likely a better approximation than using the full sample because the latter would raise issues of look-ahead bias by potentially resolving much of the uncertainty about how long the ZLB period would really last. Of course, for our macro-fiance models, a month-by-month expanding-sample procedure with repeated re-estimation of the parameters is computationally infeasible.

3 Model evaluation

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.

---

[15]In the full sample, the shadow-rate model yield factors must be latent, so for consistency we use latent factors for the affine model as well. We use the Kalman filter for the affine models and the Extended Kalman filter for the shadow-rate models. For the shadow-rate models, $Y_t = g(X_t)$, where the function $g(\cdot)$ is nonlinear and not known in closed form, so we approximate it using the Priebsch method explained in Appendix D. The Extended Kalman filter requires the calculation of the Jacobian of $g(\cdot)$, which we approximate numerically.

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.[16]

Overall, shadow-rate models fit yields better than their affine counterparts. This improvement in fit is 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-ZLB period the affine and shadow-rate models have essentially identical cross-sectional fit (results not shown), 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.

The macro-finance models exhibit slightly worse yield fit than the yields-only models. While these models have four risk factors, more than the yields-only models, only two of these are yield factors—compared to the three yield factors in our yields-only models—hence they are more constrained in fitting the cross section of yields.[17] For our purposes here, however, the cross-sectional fit of model $MZ(2)$ is sufficient.

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 (i.e., expected future short rates under $\mathbb{Q}$) or for (real-world, $\mathbb{P}$-measure) short-rate expectations.[18] 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

---

[16]On 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.

[17]One possibility would be to go to a model with three yield factors and two macro factors (as in Bauer and Rudebusch, 2015).

[18]Throughout this paper, for simplicity 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. Note that short-rate expectations under either measure are available in closed form even in the shadow-rate model—see Appendix C.

territory. Both 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 future short rates at horizons of 6, 12, and 24 months in the future, for the period from 2000 to 2013. The top panel show these probabilities for model $YA(3)$, and the bottom panel for model $MA(2)$. 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 2014, both affine models imply that these probabilities are very high. The macro-finance model implies 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 even higher probabilities of negative future short rates than for model $YA(3)$.

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. A second question is whether incorporating macroeconomic information improves interest-rate forecasts near the ZLB, because of the limited information content of yields that are constrained.

To address these questions, 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 December 2012, we calculate model-based forecasts of this short rate for horizons up to 24 months. We use a fixed-window forecast scheme, i.e., we do not re-estimate the models but instead use our baseline parameter estimates, obtained over the estimation sample from January 1985 through December 2007.

For a given forecast date and horizon, we obtain model-based forecasts by first calculating conditional expectations of the risk factors, $E_t(X_{t+h})$, and then plugging these into the relevant yield formulas, where for the affine models we replace negative forecasts by zeros. These forecasts are not the conditional expectations of future yields, since we plug in the conditional expectations of the risk factors into non-linear functions. However, while these forecasts are not optimal under mean-squared-error loss, they are optimal under absolute error loss,

because they correspond to the *median* of the forecast distribution of future yields.[19] We use the median instead of the mean because the target distribution is highly asymmetric due to the ZLB, and the median is less affected by this asymmetry. The median is optimal under an absolute-error forecast loss function. See Section 5.1 for further discussion of the use of the median in ZLB situations, where the forecast target is different but the issue is the same.

Table 3 shows in the top panel the mean absolute forecast errors in basis points for selected forecast horizons across models. The bottom panel shows relative forecast accuracy (the ratio of mean absolute errors) for four pairs of models, with asterisks indicating the significance level of the test for equal accuracy of Diebold and Mariano (2002). Our main result is that the shadow-rate models predict the short rate more accurately than the affine models. The differences in forecast accuracy are very substantial, with the shadow-rate models in several cases producing forecasts that are twice as accurate as those from the affine models. In most cases, the null for equal forecast accuracy is rejected. Overall, shadow-rate models are at least as accurate and typically much more accurate than affine models when forecasting interest rates near the ZLB.

This evidence, together with the results above, 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.

The results in Table 3 also show the benefit of incorporating macroeconomic information for forecasting at the ZLB. With only one exception, forecasts from macro-finance models outperform those from yields-only models, and the improvements in forecast accuracy are sizable. For example, at horizons longer than six months, the forecasts from the macro-finance shadow-rate model $MZ(2)$ have average errors that are almost an order of magnitude smaller than those of the yields-only shadow-rate model $YZ(3)$. These dramatic differences in forecast accuracy illustrate the importance of accounting for macroeconomic information at the ZLB. In contrast, during normal times—away from the ZLB—the yield curve itself likely contains most or all of the information necessary to predict the future course of interest rates (Duffee, 2013;

---

[19]The reason is that the median goes through nonlinear functions. Note in particular that these forecasts correspond to the target that is approximated by the following Monte Carlo simulation: First, simulate draws from the (Gaussian) distribution of the risk factors X_{t+h}, given information at time t. Second, calculate the model-implied three-month rate for each of these draws, replacing negative yields by zero for the affine model. Third, obtain the point forecast as the median of this model-implied forecast distribution of the short rate.

Bauer and Hamilton, 2015). But when the yield curve is constrained by the ZLB, yields cannot fully incorporate all relevant information and cannot reflect information in other important state variables. Hence it is particularly important to incorporate macroeconomic variables when making inference about monetary policy expectations near the ZLB. For these reasons, the macro-finance model $MZ(2)$ is our preferred model for the remainder of this paper.

4   Interpreting shadow rates

In this section we first discuss the shadow short rate, which unfortunately turns out to be not a very robust measure of the constraining influence of the ZLB on interest rates. Instead, we consider the entire distribution of future short rates and show how the asymmetry of this distribution can convey the relevance of the ZLB.

4.1   The shadow short rate

During the ZLB period, the usual indicator of monetary policy, the short rate, has not changed, but the Federal Reserve has implemented unconventional monetary policies in order to try to ease financial conditions further. To account for the effect of these atypical easing actions, some have interpreted the shadow short rate as an alternative indicator of the stance of monetary policy—see, in particular, Krippner (2014), Ichiue and Ueno (2013), and Wu and Xia (2014).[20]

Figure 3 shows estimated shadow short rates, together with the three-month T-bill rate, from 2005 to 2014. For the models $YZ(3)$ (top panel) and $MZ(2)$ (bottom panel), we show estimated shadow short rates assuming values of the numerical lower bound for interest rates, r_{min}, ranging from 0 to 25 basis points in 5 basis-point increments. When the short rate is well above zero, the various estimated shadow short rates and the observed short rate all are almost identical. However, during the recent ZLB period, the estimated shadow short rates differ markedly depending on the model and the assumed r_{min}. For our baseline $r_{min} = 0$ specification, the $MZ(2)$ model generally implies a more negative shadow short rate than the $YZ(3)$ model, and the two estimates have often moved in different directions. Shadow short rate estimates are also highly sensitive to the assumed numerical value of the lower bound r_{min}. Indeed, the resulting variation across this dimension in Figure 3 is quite striking.[21]

---

[20]Bullard (2012) considers Krippner's estimates of a very negative shadow rate in the United States as evidence of a very easy stance of monetary policy. Similarly, researchers at the Federal Reserve Bank of Atlanta 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.

[21]Kim and Singleton (2012), Christensen and Rudebusch (2015), Christensen and Rudebusch (2013), Krippner (2015), and Ichiue and Ueno (2013) provide complementary evidence that shadow short-rate es-

The lack of robustness exhibited by estimates of the shadow short rate reflects the subtle interplay of observed short-term yields with the imposed lower bound on yields. Although short-term yields are stuck near zero, they show some modest variation due to idiosyncratic factors in money markets. Apparently, these small fluctuations interact with the assumed r_{min} to determine the estimated shadow short rate. Accordingly, we have found that estimated shadow short rates are quite sensitive to both the specific short-term yields included in the model and the assumption about r_{min}. However, the choice of r_{min} in model estimations is largely arbitrary. For example, the choice of $r_{min} = 25$ basis points in Wu and Xia (2014), is hard to reconcile with 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 (2014).

The sensitivity of estimated shadow short rates raises a warning flag about their use as a measure of monetary policy, as in Ichiue and Ueno (2013) and Wu and Xia (2014). Our findings show that such estimates are not robust and strongly suggest that their use as indicators of monetary policy at the ZLB is problematic. More promising approaches have recently been suggested by Lombardi and Zhu (2014), who infer a shadow short rate that is consistent with other observed indicators of monetary policy and financial conditions, and Krippner (2015), who considers the area between shadow rates and their long-term level.

4.2 The asymmetric distribution of future short rates

Given the inadequacy of focusing on just the current shadow short rate, we now turn to a consideration of the entire path of future shadow short rates using our preferred $MZ(2)$ macro-finance model. The ZLB leads to an asymmetry in the distribution of future short rates, and the extent of this asymmetry reveals how strongly the ZLB is binding, i.e., how relevant it is for the yield curve at a certain point in time. Figure 4 illustrates this asymmetry by showing the probability densities for the distributions of the future short rate and the future shadow (short) rate, as implied by model $MZ(2)$ on December 31, 2012, for a horizon of $h = 48$ months. The densities shown are for the risk-neutral ($\mathbb{Q}$-measure) distribution, and the same arguments apply to the real-world ($\mathbb{P}$-measure) distribution. For the future shadow rate, the density is Gaussian and centered around the conditional mean $E(s_{t+h}|X_t)$. The future short rate has a mixed discrete-continuous distribution: it has a point mass at zero (indicated in the graph with a vertical line) and for positive values the density equals that of the shadow rate. Therefore, its conditional mean is higher than that of the shadow

timates can be sensitive to model specification.

14

rate, $E(r_{t+h}|X_t) > E(s_{t+h}|X_t)$. For what follows, it will be useful to define the mode of the short-rate distribution uniquely as $\max[0, E(s_{t+h}|X_t)]$ (as in Kim and Singleton, 2012). The distribution of the future short rate is right-skewed, the mean being higher than the mode.

The probability of a zero future short rate corresponds to the probability of a non-positive future shadow rate. During normal times, this probability is negligibly small, so that the mean and the mode of the short rate distribution approximately coincide. The more relevant the ZLB becomes, the larger the asymmetry of the distribution of future short rates, and the larger the difference will become between mean and mode—the "ZLB wedge." This wedge depends on the distance of yields from zero and the second moments of yield curve distribution. 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). That is, the ZLB wedge captures the value of the option of holding physical currency, which restrains nominal interest rates as they approach zero, and therefore measures how much the ZLB constrains the yield curve.

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 expectations of future shadow rates when these are positive, and equal to zero when these are non-positive. The modal path contrasts with the *mean path*, i.e., expectations of future short rates. Figure 5 displays mean and modal paths in December 2012 and in December 2013 under both the $\mathbb{Q}$- and $\mathbb{P}$-measure. For the earlier date, the ZLB wedge between the mean and modal paths is very large and it persists out to fairly long horizons. By the end of 2013, however, there is a much smaller difference between these paths, and it becomes negligible for horizons longer than about two years. Evidently, the ZLB constraint had a greater effect constraining the yield curve in December 2012 than in December 2013. This figure also demonstrates the very limited amount of information of the shadow short rate at a given point in time. Its value is similar on both dates, which is clearly not representative of the whole curve or of the tightness of the ZLB constraint.

The paths under the risk-neutral measure $\mathbb{Q}$ are estimated using information in the cross section of interest rates, and the mean path under $\mathbb{Q}$ essentially corresponds to fitted forward rates.[22] In contrast, the paths under the real-world probability measure $\mathbb{P}$ also take into account the macroeconomic information, in addition to the current shape of the yield curve. In December 2012, policy expectations under $\mathbb{P}$ and $\mathbb{Q}$ were quite similar. However, in December 2013, the $\mathbb{Q}$-measure paths were notably flatter, implying a later liftoff from the ZLB and a more gradual increase of short rates thereafter. This difference reflects a sluggish economic

---

[22]One difference is that the mean path under $\mathbb{Q}$ ignores the convexity of bond prices and the resulting Jensen-inequality terms. Consequently, it is available in closed form even in a shadow-rate model, while forward rates have to be approximated.

recovery with low underlying inflation and persistent economic slack, which in the macro-finance model results in an expectation of a very gradual easing of monetary policy.[23] To help understand the impact of the ZLB on the yield curve, Figure 6 shows actual yields together with fitted and shadow yield curves in December of 2012 and 2013. The $MZ(2)$ model provides a close fit to observed yields on both dates. However, on December 2012, shadow yields were very substantially below fitted yields, with a ZLB wedge of 3/4 to 1 percentage point, so the ZLB clearly was constraining yields. A year later, fitted and shadow yields were much closer and the effect of the ZLB was significantly attenuated.

The ZLB wedge between long-term fitted and shadow interest rate measures how tightly the ZLB constrains the entire term structure of interest rates, because it equals the cumulative difference between the mean and modal paths (under $\mathbb{Q}$). Figure 7 shows the evolution over time of the fitted and shadow ten-year yields (top panel) and the difference—the corresponding ten-year ZLB wedge (bottom panel). Over the period from 2009 to 2012, the difference between observed and shadow yields has increased substantially, indicating that the ZLB has increasingly constrained interest rates. 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.[24] Conversely, over 2013 and 2014 the ZLB constraint evidently has become less restrictive, due to improving macroeconomic conditions and a resulting higher level of the ten-year yield.

5   Monetary policy expectations

We have shown that the ZLB had a substantial impact on the term structure of interest rates in the United States in recent years and that shadow-rate DTSMs can account for this important fact. This section examines how to appropriately characterize monetary policy near the ZLB. We compare alternative approaches and argue in favor of using the modal path for this purpose, which can lead to optimal liftoff forecasts and estimates of the subsequent pace

---

[23]It is important to note, however, that estimates of real-world, $\mathbb{P}$-measure expectations are necessarily imprecise. While the risk-neutral distribution is estimated very accurately, due to the large amount of cross-sectional information in the yield curve (Cochrane and Piazzesi, 2008; Kim and Orphanides, 2012), inference about the VAR parameters μ and ϕ and about the real-world distribution of future short rates is difficult (Bauer et al., 2012; Duffee and Stanton, 2012). In other words, the policy paths under $\mathbb{Q}$ are pinned down quite precisely by the data, whereas the paths under $\mathbb{P}$ are subject to a substantial amount of uncertainty.

[24]Increases 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.

of policy tightening.

5.1 Optimal liftoff forecasts based on the target distribution

A natural starting point for estimating monetary policy expectations at the ZLB is an examination of the distribution of future policy liftoff. In a shadow-rate model, policy liftoff corresponds to the initial time the shadow short rate rises above a given threshold. We set this threshold at 25 basis points, which is consistent with 0 to 25 basis points range the Federal Reserve has kept during the ZLB period. In the case of a univariate model, the distribution of the liftoff horizon can in some cases be derived analytically.[25] In multi-factor shadow-rate models, the liftoff distribution has to be obtained using Monte Carlo simulation. Figure 8 shows the smoothed Kernel density of the liftoff distribution on December 31, 2012, based on simulations from our preferred model $MZ(2)$ under the risk-neutral distribution.[26] The figure also reports the mean, median, mode, and interquartile range and alternative liftoff estimates based on the mean and modal paths that we will discuss below. The liftoff distribution is strongly skewed to the right—even very distant horizons for policy liftoff are not uncommon.

Accordingly, the most appropriate forecast of policy liftoff based on this distribution will depend heavily on the forecaster's loss function. For squared-error loss, the mean is the optimal forecast. Given the large weight the mean gives to outliers, the long-horizon right tail of the distribution leads to unappealingly distant liftoff forecasts. Hence the mean is not a useful summary of the central tendency of the liftoff distribution. Under absolute-error loss, the median as an optimal forecast. This choice appears much more reasonable in this context, given that outliers do not have an unduly large influence. 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.[27] Hence, to estimate future policy liftoff, our focus will be on the median as the most natural choice for forecasting. We now discuss how the liftoff distribution relates to estimates of liftoff based on the mean and modal paths.

---

[25]See 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.

[26]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 date of liftoff is determined by the time that the shadow rate hits the threshold. Due to the erratic nature of the sample paths, we also require that the shadow rate stays above the threshold for 12 months before we designate a policy liftoff, which leads to better-behaved liftoff distribution.

[27]Ichiue 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 but not discuss the median.

5.2   Mean path vs. modal path

It is a common practice to base estimates of liftoff on the mean path (under $\mathbb{Q}$) obtained from forward rates or money market futures rates. The time until liftoff is given by the horizon when this path rises above a certain threshold (e.g., 25 basis points).[28] At first glance, this approach may seem to be an attractive model-free approach. However, by using the mean path, it ignores the asymmetry of the distribution of the future short-term interest rate induced by the ZLB. Because of this asymmetry, the mean path does not reflect the most likely value of the policy rate at a future point in time, as noted in Section 4.2. Instead, we advocate using the modal path for calculating liftoff estimates, and some professional forecasters and market analysts estimate liftoff in exactly 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.[29] Furthermore, liftoff estimates using the modal path essentially coincide with the median of the liftoff distribution, which again is the optimal forecast under absolute-error loss. For example, Figure 8 shows that in December 2012, these two estimates were almost identical—with the median of the distribution at 33 months and the liftoff horizon based on the modal path at 34 months. The intuition for this concordance is the following: When the modal path crosses the liftoff threshold, the (Gaussian) shadow short rate is equally likely to be above or below the modal short-rate path, so there is an equal probability for liftoff to occur earlier or later (i.e., the modal threshold crossing is also the median of the liftoff distribution.[30] Therefore, forecasting liftoff in a straightforward fashion using the modal path is justified based on a more complete analysis of the forecasting problem, which also introduces optimality under an absolute-error loss.

Figure 5 also illustrates the liftoff estimates based on mean and modal paths. Horizontal lines at 25 basis points indicate the threshold for liftoff. As the modal path is always below

[28]For example, Ueno 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, and there are many similar examples in U.S. financial market commentary, including "Fed Likely to Push Back on Market Expectations of Rate Increase," from the June 13, 2013 issue of the Wall Street Journal.

[29]The 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.

[30]To be more precise, denote by h^* the horizon where the modal path crosses the threshold. It is equally likely for the shadow rate 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 close to h^*.

the mean path, the liftoff implied by the modal path is always later than that implied by the mean path. In December 2012, the difference is particularly pronounced, due to the strong asymmetry of the future short rate distribution. For the paths under the ℚ-measure, the modal-path liftoff forecast at that point was 34 months, while the estimate based on the forward curve was 22 months.

5.3 Comparing liftoff estimates over the recent ZLB period

We now discuss the time series of alternative estimates for policy liftoff during the ZLB period from January 2008 to December 2014. The top panel of Figure 9 shows model-based estimates for liftoff based on the mean and modal paths, as well as the median and the interquartile range (IQR) of the forecast distribution for the liftoff horizon. While we focus on the estimates under the ℚ-measure, the same arguments apply to estimates under the ℙ-measure. Again, over the whole sample, the median of the liftoff distribution is extremely close to the liftoff based on the modal path. Consequently, the modal-path-based estimate will be an optimal forecast for liftoff under a mean-absolute-error loss function. In contrast, the liftoff forecasts based on the mean path—corresponding to the common practice of estimating liftoff based on forward rates—imply liftoff that is much earlier, often by more than a year. Differences between liftoff estimates from the mean and modal path are particularly pronounced during 2011 and 2012, which was when the asymmetry due to the ZLB was strongest. Our interval forecasts for future liftoff—the IQR of the target distribution—illustrate the substantial uncertainty around the point forecasts for future policy liftoff. Late in 2011 and in 2012, when the liftoff forecasts imply especially long horizons, the uncertainty is particularly pronounced as longer horizons generally come with a larger degree of uncertainty.[31]

We now put the shadow-rate model forecasts of the time until liftoff in perspective by comparing them to alternative estimates. The bottom panel of Figure 9 shows liftoff forecasts based on the model-bases modal paths under ℙ and ℚ and two alternative calculations of future liftoff dates by the private sector. The first is the median of modal forecasts for the time of policy liftoff from the Survey of Primary Dealers (SPD) of the Federal Reserve Bank of New York, which is publicly available going back to January 2011.[32] The second alternative

---

[31]The model does not have stochastic volatility; however, as noted by Christensen and Rudebusch (2013) shadow-rate models can capture some of the time variation in second moments about future monetary policy at the ZLB.

[32]See http://www.newyorkfed.org/markets/primarydealer_survey_questions.html for the questions and answers for each survey. In the survey, the respondents are asked to provide the "estimate for [the] most likely quarter and year of [the] first target rate increase." We use the middle month of the quarter to translate these responses into monthly horizons.

source of liftoff estimates is from Macroeconomic Advisers (MA), based on their published "policy call" and on their modal projections of the future path of the federal funds rate, i.e., the most likely scenario for Fed policy in their view. The bottom panel of 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 at least as large as the horizons corresponding to these calendar dates.

Liftoff forecasts based on the modal path under $\mathbb{Q}$ reflect the views about policy liftoff that were priced into the yield curve at each month in our sample. These estimates are generally close to those from the SPD and MA (with the exception of 2009, where the MA estimates imply substantially later liftoff). Both the outside estimates and the liftoff forecast based on the $\mathbb{Q}$-modal path are also generally consistent with the FOMC's forward guidance. In particular, they responded to the introduction of calendar-based forward guidance by the FOMC on August 9, 2011, when the Committee noted in its statement that it expected a near-zero policy rate until at least mid-2013. This led to a substantial increase in the expected liftoff horizon (see also Swanson and Williams, 2012).

Turning to the liftoff forecasts under the real-world measure $\mathbb{P}$, we find that over the first half of the ZLB period, from 2009 to about mid-2012, these imply later liftoff than the estimates under $\mathbb{Q}$. This is not surprising, since the former include the information in macroeconomic variables, whereas the latter are based solely on the information in the yield curve. The discrepancy between the two is due to the substantially depressed macroeconomic situation during the Great Recession and in the early years of the recovery. Our macro-finance model takes into account the elevated unemployment gap and low inflation during this period, which informs the $\mathbb{P}$-measure forecasts for short rates and liftoff. The $\mathbb{P}$-measure estimates of liftoff are also generally later than the outside estimates (the only exception being the MA forecasts in 2009), and later than the FOMC's forward guidance. It is noteworthy that the simple macro-finance shadow-rate model gave substantially longer and (with the benefit of hindsight) more reasonable estimates of liftoff during the early years of the ZLB period than most professional forecasters anticipated at the time.

Finally, we note the close connection between the tightness of the ZLB constraint and liftoff. The correlation between the ZLB wedge in the ten-year yield and the model-based liftoff horizon estimate (using the $\mathbb{Q}$-measure modal path) is 0.99. Intuitively, variation in the length of the expected period of near-zero policy rates is the main reason for variation

in yields that are constrained by the ZLB. In contrast, the connection between the ZLB wedge and the shadow short rate is much weaker, with a correlation coefficient of -0.47. The shadow short rate contains a more limited amount of information, lacks robustness and is hard to interpret, while the expected liftoff horizon more accurately summarizes the evolution of the yield curve and the stance of monetary policy at the ZLB.[33] For example, based on the liftoff estimates in Figure 9, the stance of policy became increasingly accommodative from 2009 to 2012. Federal Reserve forward guidance announcements—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—typically had a noticeable impact and lengthened the estimated liftoff horizon. Accordingly, the expected liftoff horizon appears to be a fairly comprehensive univariate summary of the stance of monetary policy at the zero lower bound.

5.4 Pace of tightening

Another key dimension of monetary policy expectations at the ZLB is the expected pace of policy tightening after liftoff. For a *given* expected liftoff horizon, a faster expected pace implies higher interest rates, and vice versa. Hence, the anticipated pace of tightening is a crucial determinant of interest rates at medium and long maturities. To estimate the pace of tightening, a first issue is what metric to use. One possibility would be the time from liftoff until the policy rate reaches a certain higher threshold such as one or two percent. Instead, we use the cumulative anticipated increase in the policy rate over the two years after liftoff—a statistic that can easily be compared to previous monetary policy normalizations. A second issue is how to use shadow-rate models to estimate the anticipated increase in the policy rate after liftoff. We use the modal path for this purpose because again the modal path correctly accounts for the asymmetry of the short-rate distribution near zero. The mean path, in contrast, is always flatter than the modal path and would imply a slower pace of tightening.

Hence, we measure the expected pace of policy tightening as the increase in the modal path over the two-year horizon after first crossing the 25 basis point liftoff threshold. Figure 10 shows this measure for the modal path under both the $\mathbb{Q}$ and $\mathbb{P}$ measures. The first thing to note is that the pace under $\mathbb{Q}$ has been very volatile, varying considerably from below 1.5 to over 3 percentage points. This is due to the fact that this measure mostly

---

[33]We also note that in contrast to estimated shadow short rates, our modal paths, forecasts for liftoff, and expected pace of tightening estimated from model $MZ(2)$ are very robust to different choices of the numerical lower bound, r_{min}—pairwise correlations are close to one and the estimates are generally very similar (results omitted). This is a very desirable feature, since it is hard to pin down r_{min} with data on short-term interest rates (see discussion in Section 4.1).

reflects information in the cross section of interest rates, and any shift in the steepness of the yield curve translates into changes in the estimated pace. In contrast, the pace under $\mathbb{P}$ is much more stable, which reflects steadier model-based forecasts of future interest rates including macroeconomic variables. This measure, which is our preferred measure of the pace of tightening, declines over 2013 and 2014, and at the end of our sample, in December 2014, indicates an anticipated increase in the policy path of slightly below 1.5 percentage points over two years.

Figure 10 also shows the anticipated pace of policy tightening that is implied by the SPD modal policy paths. Both our model-based estimates imply a slower pace than is apparently expected by the survey respondents. While our model-based estimates imply an increase of about a 1.5 to 2 percentage points in the policy rate over the two years after liftoff, the Primary Dealers have generally anticipated fairly steadily an increase in the range of 2 to 2.5 percentage points. It is noteworthy that our preferred model-based estimates have generally implied both a later liftoff and a slower pace of tightening than outside estimates.

How do these estimates compare to historical episodes of policy tightening? During the tightening cycle from February 1994 to February 1995, the Fed increased the policy rate from 3 to 6 percentage points, which corresponds to a pace of 6 percentage points over two years. From June 1999 to May 2000 the increase was from 4.75 to 6.5 percent, at a pace of about 3.5 percentage points, and from June 2004 to June 2006 the policy rate was raised from 1 percent to 5.25 percent, an increase at a pace of 4.25 percentage points. Clearly, in all three previous policy tightening cycles, the pace of tightening was substantially faster than it is expected for the period after liftoff from the ZLB. This discrepancy may be explained by the unusual situation of the U.S. economy at the ZLB. Some Fed policymakers also expected a slow pace of policy tightening after liftoff (see, for example Dudley, 2014). In a speech in March 2015, Fed Chair Yellen (see Yellen, 2015) indicated that the pace of tightening in past monetary policy cycles may be a "highly misleading guide" to the course of monetary policy in 2015 and beyond because of perceived macroeconomic uncertainties, headwinds to the domestic and global economic outlook, and a possibly slower rate of long-run growth.

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

An admitted shortcoming of our ZLB term structure model is the assumption of stationarity across pre-ZLB and ZLB periods. This assumption is made by all ZLB models that we are aware of. Given the unique character of a situation with near-zero policy rates, one could certainly imagine that the interaction and evolution of yields is fundamentally different than during normal times, in particular when macroeconomic variables are included in the model. This issue could potentially be addressed in future research by incorporating regime-switching processes or time-varying parameters into shadow-rate DTSMs.

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., 2014; Bauer, 2015). 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 for DTSM estimation, as in Bauer (2015), holds promise for shadow-rate models. In particular, Bayesian methods would allow for an accurate description of the uncertainty around current and future shadow rates and liftoff estimates that results from forecast errors, parameter uncertainty, and even model uncertainty.

References

Bauer, Michael D., "Restrictions on Risk Prices in Dynamic Term Structure Models," Working Paper 2011-03, Federal Reserve Bank of San Francisco February 2015.

__ **and Glenn D. Rudebusch**, "Resolving the Spanning Puzzle in Macro-Finance Term Structure Models," Working Paper 2015-01, Federal Reserve Bank of San Francisco January 2015.

__ **and James D. Hamilton**, "Robust Bond Risk Premia," Technical Report, unpublished manuscript 2015.

__ **, Glenn D. Rudebusch, 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, "Modeling Yields at the Zero Lower Bound: Are Shadow Rates the Solution?," working paper 2013-39, Federal Reserve Bank of San Francisco 2013.

__ **and** __ , "Estimating Shadow-Rate Term Structure Models with Near-Zero Yields," *Journal of Financial Econometrics*, Spring 2015, *13* (2), 226–259.

Cochrane, John H. and Monika Piazzesi, "Decomposing the Yield Curve," unpublished manuscript, Chicago Booth School of Business 2008.

Diebold, Francis X and Robert S Mariano, "Comparing predictive accuracy," *Journal of Business & economic statistics*, 2002, *20* (1), 134–144.

Dudley, William C., "The Economic Outlook and Implications for Monetary Policy," Speech at the New York Association for Business Economics May 20 2014.

Duffee, Gregory R., "Term Premia and Interest Rate Forecasts in Affine Models," *Journal of Finance*, February 2002, *57* (1), 405–443.

___ , "Forecasting Interest Rates," in Graham Elliott and Allan Timmermann, eds., *Handbook of Economic Forecasting*, Vol. 2, Part A, Elsevier, 2013, pp. 385–426.

___ **and Richard H. Stanton**, "Estimation of Dynamic Term Structure Models," *Quarterly Journal of Finance*, 2012, *2* (2).

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 2014.

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.

Greene, William H, *Econometric Analysis*, 7th edition ed., Prentice Hall, 2003.

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, "Gaussian Macro-Finance Term Structure Models with Lags," *Journal of Financial Econometrics*, 2013, *11* (4), 581–609.

___ **,** ___ **, and** ___ , "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 Priebsch, and Kenneth J. Singleton**, "Risk Premiums in Dynamic Term Structure Models with Unspanned Macro Risks," *Journal of Finance*, June 2014, *69* (3), 1197–1233.

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 Investifation of Japanese Yields," *Journal of Econometrics*, September 2012, *170* (1), 32–49.

— **and Marcel A. Priebsch**, "Estimation of Multi-Factor Shadow-Rate Term Structure Models," Unpublished manuscript, Federal Reserve Board 2013.

Krippner, Leo, "Measuring the stance of monetary policy in zero lower bound environments," *Economics Letters*, January 2013, *118* (1), 135–138.

—, "A Tractable Framework for Zero-Lower-Bound Gaussian Term Structure Models," CAMA Working Paper 49/2013, Reserve Bank of New Zealand 2014.

—, *Term Structure Modeling at the Zero Lower Bound: A Practitioner's Guide*, Palgrave-Macmillan, 2015.

Linetsky, Vadim, "Computing hitting time densities for CIR and OU diffusions: Applications to mean-reverting models," *Journal of Computational Finance*, 2004, *7*, 1–22.

Lombardi, Marco Jacopo and Feng Zhu, "A shadow policy rate to calibrate US monetary policy at the zero lower bound," Technical Report, Bank for International Settlements 2014.

Piazzesi, Monika, "Affine Term Structure Models," in Yacine Ait-Sahalia and Lars Hansen, eds., *Handbook of Financial Econometrics*, Vol. 1, Elsevier, 2010, chapter 12.

Priebsch, 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," Working Paper, National Bureau of Economic Research 2014.

Yellen, Janet L., "Normalizing Monetary Policy: Prospects and Perspectives," Speech at the conference "The New Normal for Monetary Policy" at the Federal Reserve Bank of San Francisco March 27 2015.

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

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

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 of an affine short rate

Here we provide expressions for the moments of an affine Gaussian short rate. These results hold for the short rate in an affine Gaussian DTSM as well as for the shadow short rate in a Gaussian shadow-rate model. To avoid confusion with the short rate in the shadow-rate model, we use s_t to denote the affine short rate, with $s_t = \delta_0 + \delta_1 X_t$, and the dynamics of X_t under $\mathbb{Q}$ and $\mathbb{P}$ given in equations (2) and (3), respectively. To save on superscripts, we give all moments in Appendix B and C under the $\mathbb{P}$-measure. Exactly analogous expressions, with μ and ϕ replaced by $\mu^{\mathbb{Q}}$ and $\phi^{\mathbb{Q}}$, give the moments under the $\mathbb{Q}$-measure, which will be needed for bond pricing in Appendix D.

The conditional mean is

$$\mu_{t \to t+h} := E_t(s_{t+h}) = \delta_0 + \delta_1' E_t(X_{t+h}) = \delta_0 + \delta_1' \left[(1 - \phi^h) E(X_t) + \phi^h X_t \right]$$

and the conditional variance is

$$\sigma_h^2 := Var_t(s_{t+h}) = \delta_1' Var_t(X_{t+h}) \delta_1 = \delta_1' \left[\sum_{i=0}^{h-1} \phi^i \Sigma \Sigma' (\phi^i)' \right] \delta_1,$$

which is time-invariant because X_t is homoskedastic. The conditional covariance is

$$\sigma_{g,h} := Cov_t(s_{t+g}, s_{t+h}) = \delta_1' \left[\sum_{i=1}^{\min(g,h)} \phi^{g-i} \Sigma\Sigma'(\phi^{h-i})' \right] \delta_1.$$

We will also need cross-moments, i.e., non-centered second moments, which are given as

$$E_t(s_{t+g}s_{t+g}) = \sigma_{g,h} + \mu_{t\to t+g}\mu_{t\to t+h}.$$

C   Moments of the short rate in a shadow-rate model

For the short rate in the shadow-rate model, r_t, the conditional mean at $t+h$ is:

$$
\begin{aligned}
E(r_{t+h}|X_t) &= P(s_{t+h} > 0|X_t)E(s_{t+h}|X_t, s_{t+h} > 0) \\
&= \Phi(\mu_{t\to t+h}/\sigma_h)\left[\mu_{t\to t+h} + \sigma_h\frac{\phi(-\mu_{t\to t+h}/\sigma_h)}{1 - N(-\mu_{t\to t+h}/\sigma_h)}\right) \\
&= \Phi(\mu_{t\to t+h}/\sigma_h)\mu_{t\to t+h} + \sigma_h\phi(-\mu_{t\to t+h}/\sigma_h),
\end{aligned}
\tag{6}
$$

where $\Phi(\cdot)$ is the cumulative distribution function of the standard normal distribution, and $\phi(\cdot)$ is its density function.[34] The second line follows from well-known results about truncated normal distributions. Based on these results the path of expected future short rates for any given value of X_t can be calculated. By using the $\mathbb{Q}$-measure parameters instead of μ and ϕ to calculate the conditional moments, we obtain $\mathbb{Q}$-expectations of future short rates. These are very close to forward rates, but not identical due to convexity/Jensen inequality effects in bond prices, which are the reason that we cannot solve for yields in the shadow-rate model analytically.

We can understand some existing approximations in the literature on shadow-rate DTSMs based on equation (6). The first-order approximation used by, for example, Ichiue and Ueno (2013) amounts to using expected short rates under $\mathbb{Q}$, based on (6), as the forward rates of the shadow-rate model. Krippner (2014) and Wu and Xia (2014) also use (6) to obtain forward rates for the shadow-rate model, but they plug in the affine forward rates[35] for $\mu_{t\to t+h}$, which significantly increases the accuracy of the approximation.

We now turn to conditional, non-centered second moments, i.e., $E_t(r_{t+g}r_{t+h})$. For $g = h$ we can use a standard result for the variance of censored normal random variables (see, for example Greene, 2003, Theorem 19.3), which implies that

$$E_t(r_{t+h}^2) = (\mu_{t\to t+h}^2 + \sigma_h^2)\Phi(\mu_{t\to t+h}/\sigma_h) + \sigma_h\mu_{t\to t+h}\phi(\mu_{t\to t+h}/\sigma_h).$$

---

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

[35]The affine forward rates are given by

$$f_t^h = \log(P_t^h) - \log(P_t^{h+1}) = \mathcal{A}_h - \mathcal{A}_{h+1} + (\mathcal{B}_h - \mathcal{B}_{h+1})X_t.$$

More generally, for $g \neq h$, we use a result provided in Priebsch (2013):

$$
\begin{aligned}
E_t(r_{t+g} r_{t+h}) &= (\mu_{t \to t+g} \mu_{t \to t+h} + \sigma_{g,h})(1 - \Phi(-\varsigma_{t,g}) - \Phi(-\varsigma_{t,h}) + \Phi_2(-\varsigma_{t,g}, -\varsigma_{t.h}, \chi_{g,h})) \\
&\quad + \sigma_h \mu_{t \to t+g} \phi(\varsigma_{t,h}) \Phi\left(\frac{\varsigma_{t,g} - \chi_{g,h} \varsigma_{t,h}}{\sqrt{1 - \chi_{g,h}^2}} \right) \\
&\quad + \sigma_g \mu_{t \to t+h} \phi(\varsigma_{t,g}) \Phi\left(\frac{\varsigma_{t,h} - \chi_{g,h} \varsigma_{t,g}}{\sqrt{1 - \chi_{g,h}^2}} \right) \\
&\quad + \sigma_g \sigma_h \sqrt{\frac{1 - \chi_{g,h}^2}{2\pi}} \phi\left(\sqrt{\frac{\varsigma_{t,g}^2 - 2\chi_{g,h} \varsigma_{t,g} \varsigma_{t,h} + \varsigma_{t,h}^2}{1 - \chi_{g,h}^2}} \right),
\end{aligned}
$$

where $\varsigma_{t,h} = \mu_{t \to t+h}/\sigma_h$, $\chi_{g,h} = \sigma_{g,h}/(\sigma_g \sigma_h)$, and $\Phi_2(\cdot, \cdot, \rho)$ is the cumulative distribution function for the bivariate standard normal distribution with correlation coefficient ρ.

Note that the above results are for the special case $r_{min} = 0$, but are easy to generalize.

D   Discrete-time bond pricing using cumulants

This appendix develops an approximative solution for bond prices, which can be used to calculate yields in shadow-rate models. It takes the ideas developed in Priebsch (2013) for continuous-time shadow-rate models, and translates them to the discrete-time case. An advantage of the discrete-time setting is that no numerical integration is necessary.

The price of a default-free, zero-coupon bond in discrete time is

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

Define the random variable $R_t^n = \sum_{i=0}^{n-1} r_{t+i}$, which is not known until time $t + n - 1$. The conditional cumulant-generating function of R_t^n under the $\mathbb{Q}$-measure is

$$
C_t(\eta) = \log\left[E_t^{\mathbb{Q}} \left(e^{\eta R_t^n} \right) \right],
$$

and the log bond price is $\log(P_t^n) = C_t(-1)$. The function $C_t(\eta)$ has the power series representation

$$
C_t(\eta) = \sum_{j=1}^{\infty} \eta^j \frac{\kappa_j}{j!},
$$

where κ_j are the cumulants of R_t^n, which are generated by the j-th derivative of $C_t(\eta)$ evaluated at $\eta = 0$. Priebsch' bond-pricing method for shadow-rate models is based on approximating

$C_t(-1)$ using the first two cumulants:

$$C_t(-1) \approx -\kappa_1 + \frac{1}{2}\kappa_2.$$

D.1   First cumulant

The first cumulant is equal to the first moment:

$$\kappa_1 = C_t'(0) = E_t^{\mathbb{Q}} R_t^n = \sum_{i=0}^{n-1} E_t^{\mathbb{Q}}(r_{t+i}).$$

D.2   Second cumulant

The second cumulant is the variance:

$$\kappa_2 = Var_t^{\mathbb{Q}}(R_t^n) = E_t^{\mathbb{Q}}\left[(R_t^n)^2\right] - \left[E_t^{\mathbb{Q}}(R_t^n)\right]^2 = E_t^{\mathbb{Q}}\left[(R_t^n)^2\right] - \kappa_1^2.$$

For the second non-centered moment of R_t^n we have

$$E_t^{\mathbb{Q}}\left[(R_t^n)^2\right] = E_t^{\mathbb{Q}}\left[\left(\sum_{i=0}^{n-1} r_{t+i}\right)^2\right] = \sum_{i=0}^{n-1}\sum_{j=0}^{n-1} E_t^{\mathbb{Q}}(r_{t+i}r_{t+j}).$$

Note that if either i or j is equal to zero, the current short rate comes outside of the expectations operator, and we just need to calculate first moments.

D.3   Yields

Approximating yields using only the first cumulant gives

$$y_t^n \approx n^{-1}\kappa_1 = n^{-1}\sum_{i=0}^{n-1} E_t^{\mathbb{Q}}(r_{t+i}).$$

This approximation, which is used for example by Ichiue and Ueno (2013), ignores Jensen inequality terms, i.e., the convexity of bond prices, and is quite inaccurate as shown by Priebsch (2013). Using a second-order approximation to the cumulant-generating function, we obtain

$$y_t^n \approx n^{-1}\left(\kappa_1 - \frac{1}{2}\kappa_2\right).$$

Priebsch shows that for the specific yields-only model that he considers, this approximation is very accurate. Note that for the affine model, this approximation is exact, since only the first two cumulants of a Gaussian random variable are non-zero. Again, these results are for $r_{min} = 0$ but are straightforward to generalize (see Priebsch, 2013).

Table 1: Cross-sectional fit

Model	Total	3m	6m	1y	2y	3y	5y	7y	10y
Full sample									
YA(3)	6.7	7.0	8.5	10.0	2.7	5.8	6.0	2.8	7.2
MA(2)	11.8	16.6	9.7	14.3	12.5	10.2	8.2	7.2	13.2
YZ(3)	6.5	6.8	8.5	9.5	3.0	5.6	6.2	3.1	6.5
MZ(2)	10.4	13.5	10.4	12.4	10.5	9.3	6.9	5.7	11.8
ZLB subsample									
YA(3)	7.6	6.7	3.7	10.8	2.4	7.8	9.5	2.9	10.8
MA(2)	17.6	23.6	7.2	23.2	19.2	14.1	15.7	12.3	19.2
YZ(3)	6.8	5.5	4.5	9.2	3.4	6.8	9.7	4.3	8.2
MZ(2)	12.1	9.8	11.2	17.1	12.1	10.6	12.2	7.5	14.2

Notes: Root-mean-squared fitting errors of model-implied yields in basis points. Full sample: January 1985 to December 2014. ZLB subsample: December 2008 to December 2014.

Table 2: Violations of the ZLB

	Forward rates		Short-rate expectations	
Model	frequency	avg. length	frequency	avg. length
YA(3)	23	4.1	7	4.3
MA(2)	46	9.8	62	18.5

Notes: Number of months, between December 2008 and December 2014, 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(s)	6m	12m	18m	24m
Mean absolute forecast errors				
YA(3)	18.4	48.2	85.6	122.1
MA(2)	9.5	8.4	7.4	22.9
YZ(3)	8.0	20.5	49.7	88.2
MZ(2)	9.2	8.1	6.2	12.2
Relative mean absolute forecast errors				
YZ(3)/YA(3)	0.44*	0.43***	0.58***	0.72***
MZ(2)/MA(2)	0.97	0.96	0.84**	0.54***
MA(2)/YA(3)	0.52	0.17***	0.09***	0.19***
MZ(2)/YZ(3)	1.15	0.39**	0.13***	0.14***

Notes: Top panel shows the mean absolute forecast errors (in basis points) for out-of-sample forecasts of the three-month T-bill rate at various forecast horizons. Bottom panel shows the relative mean absolute forecast errors for different model pairs. *, **, and *** indicate significance of the Diebold-Mariano test for equal forecast accuracy at the 10%, 5%, and 1% level, respectively. Forecast period: December 2008 to December 2012.

Figure 1: Yields and macroeconomic data
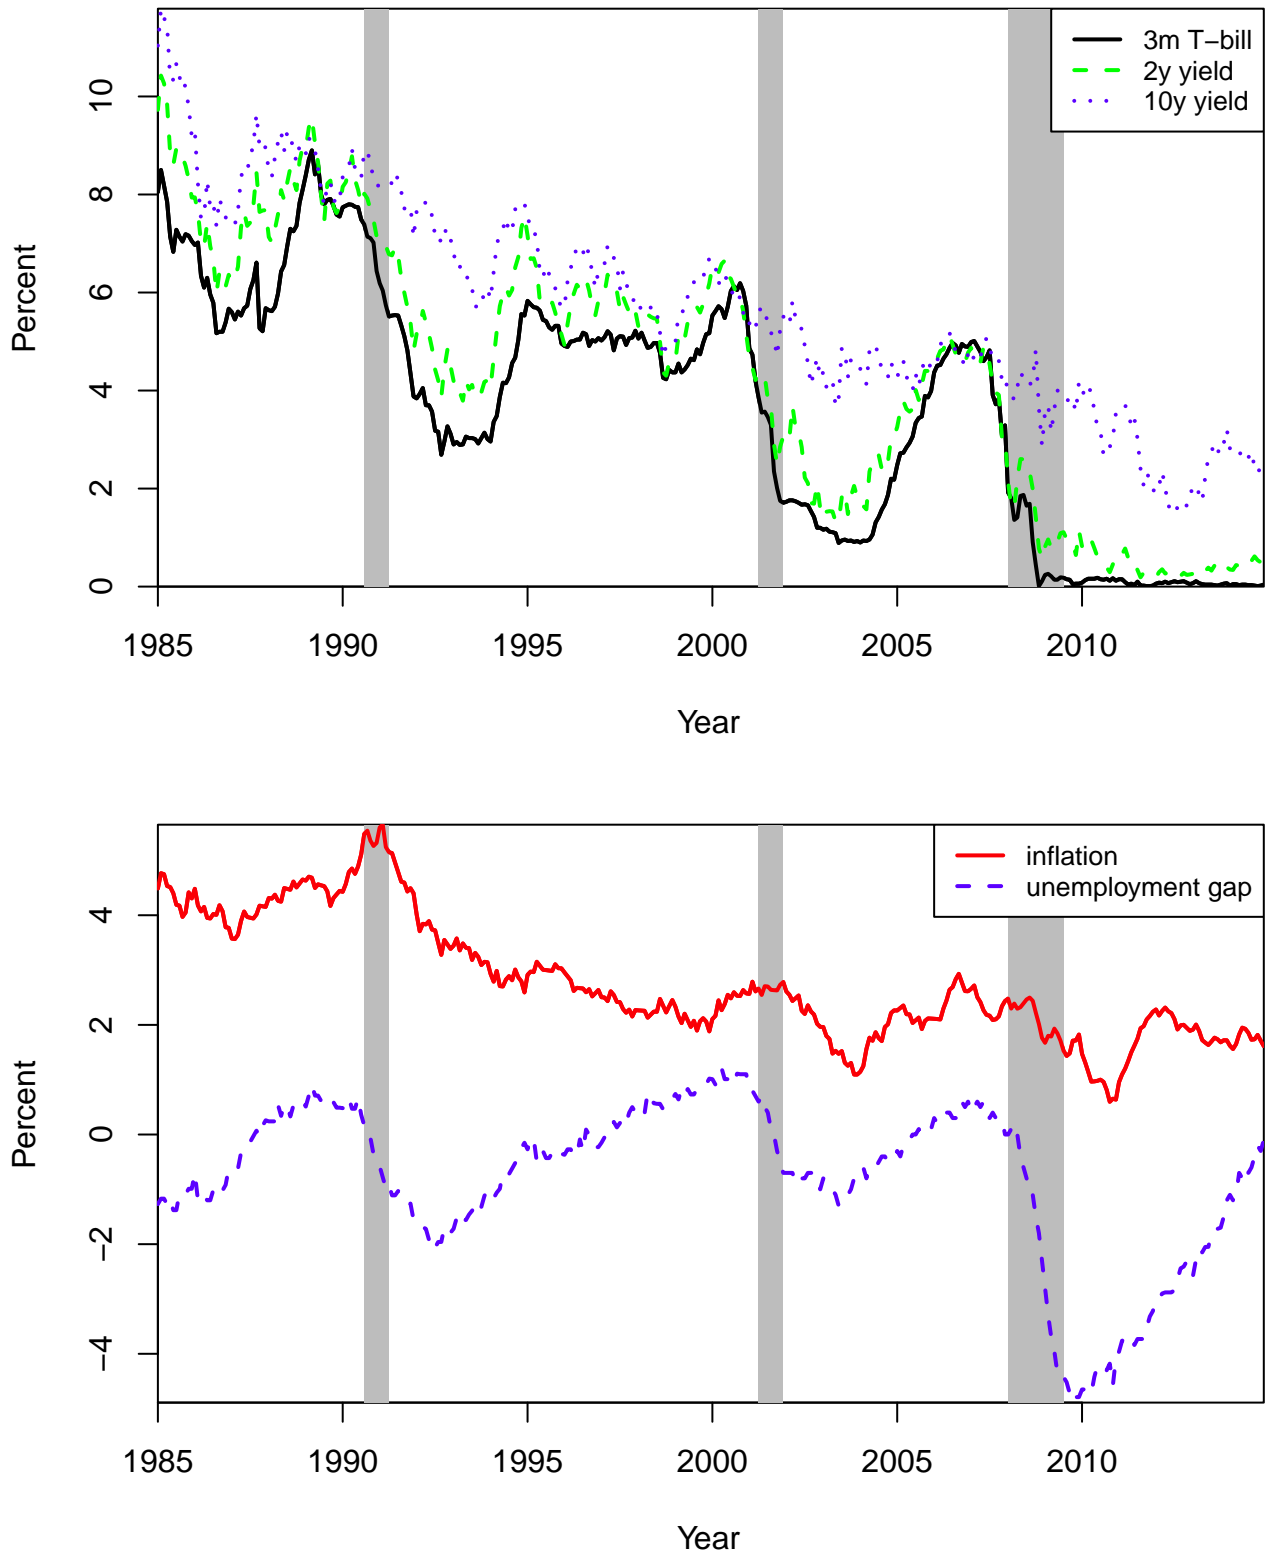
Notes: Top panel shows three selected yields, for the three-month, two-year, and ten-year maturities. Bottom panel shows the macroeconomic data. Sample period: January 1985 to December 2014

Figure 2: Affine model probabilities of negative future short rates
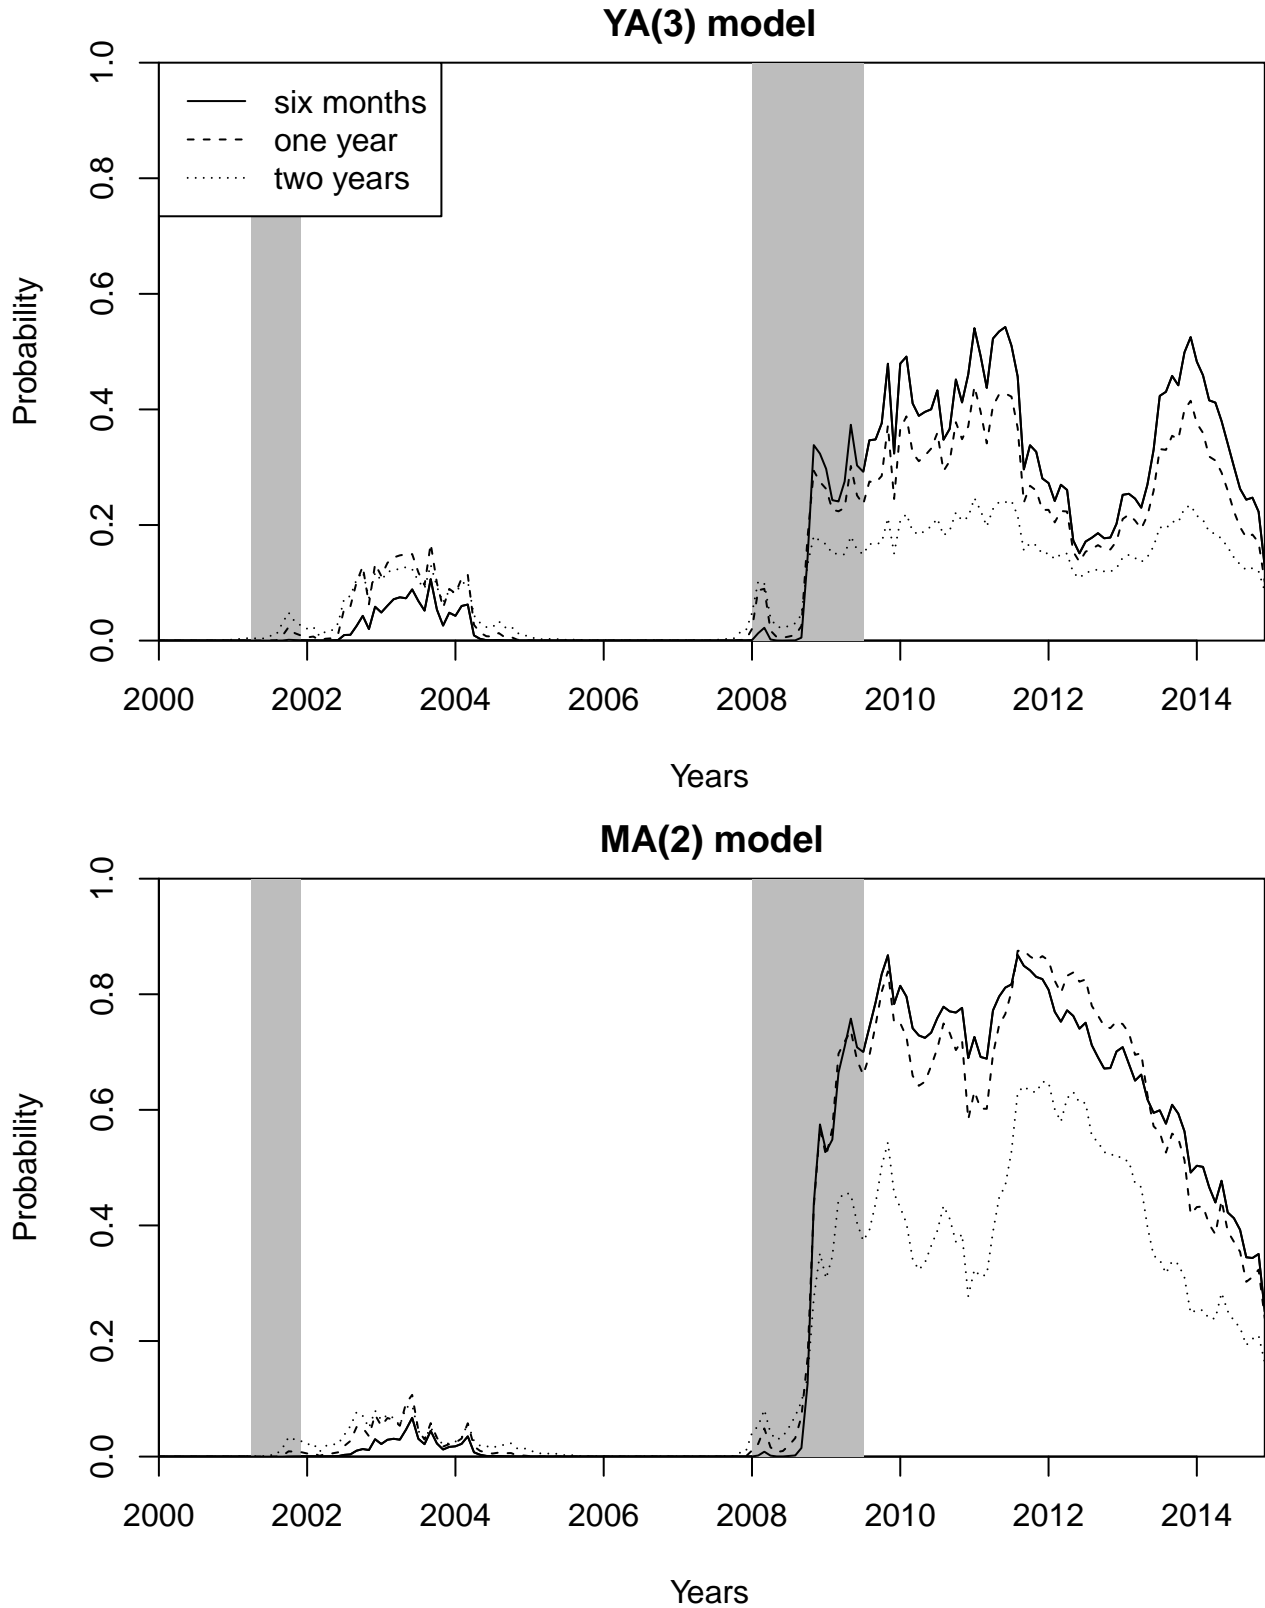
Notes: Model-implied real-world ($\mathbb{P}$) probabilities of negative future short-term interest rates at horizons of six months, one year, and two years. Shaded areas correspond to NBER recessions. Sample period: January 2000 to December 2014.

Figure 3: Shadow short rates
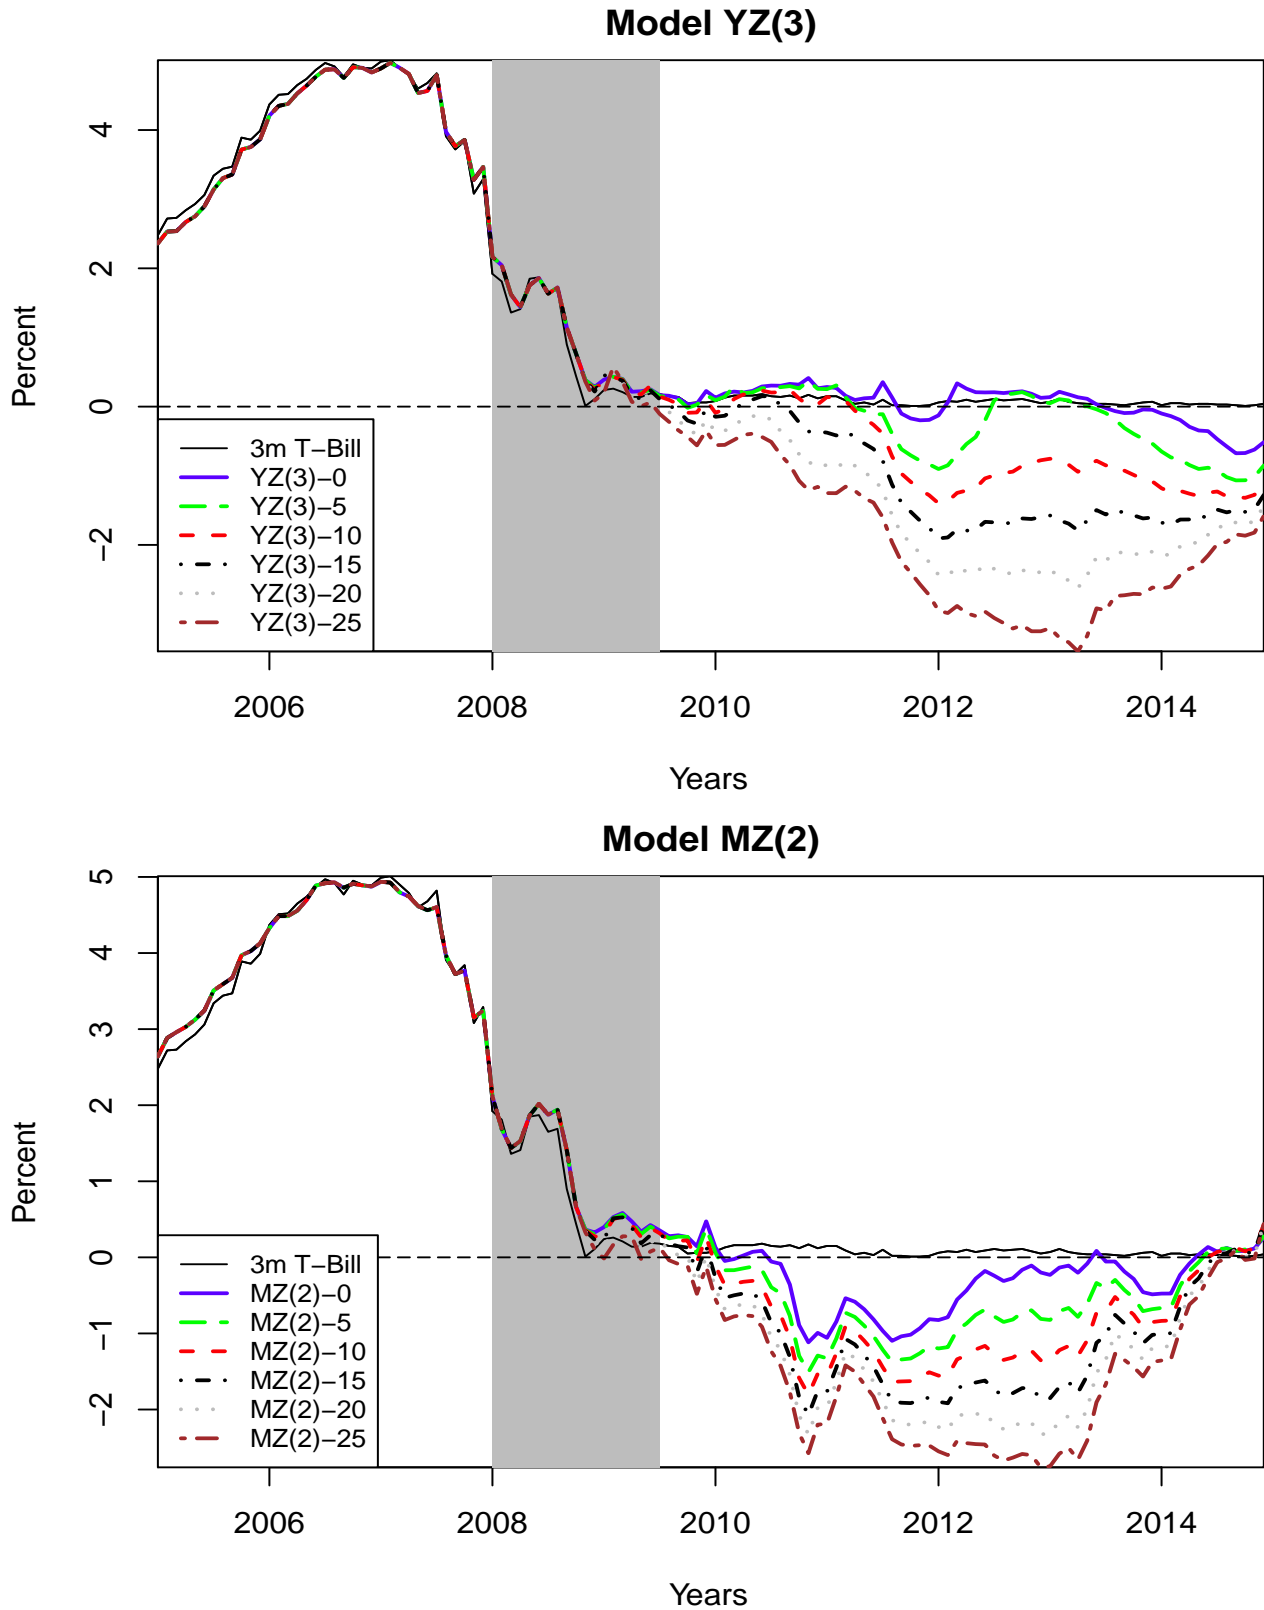
Notes: Shadow short rates implied by shadow-rate models and the three-month T-bill rate. Shaded areas correspond to NBER recessions. Each panel shows estimated shadow short rates for r_{min} varying from 0 to 25 basis points; the top panel for model $YZ(3)$, the bottom panel for model $MZ(2)$. Sample period: January 1985 to December 2014.

Figure 4: Distribution of future shadow rate and short rate
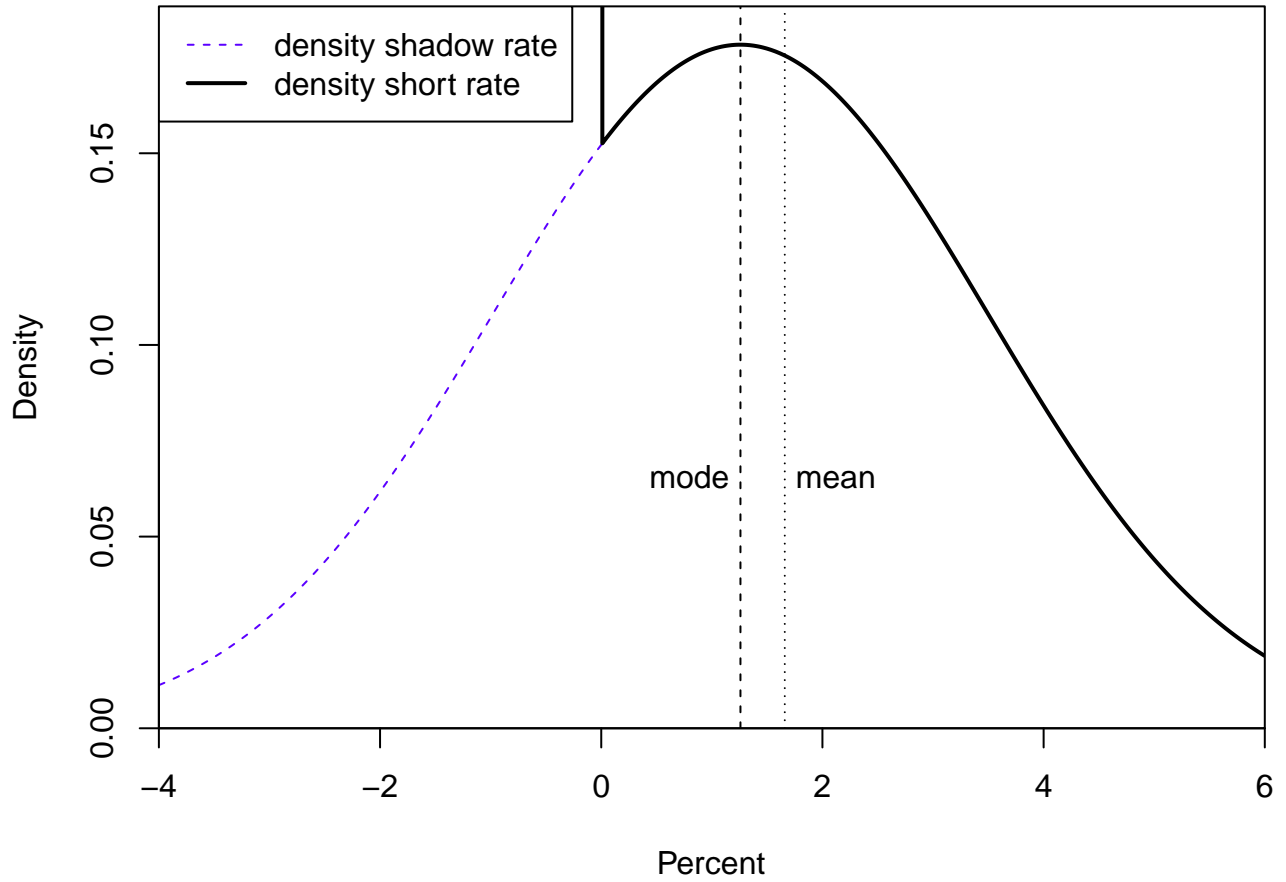
Notes: Densities of future shadow rate and of future short rate, at horizon of 48 months, on December 31, 2012. Vertical lines show the mode and mean of the distribution of the future short rate. Model: $MZ(2)$.

Figure 5: Policy paths

December 2012

Horizon

December 2013

	Q–mean short rate
	Q–mean shadow rate
	Q–mode short rate
	P–mean short rate
	P–mean shadow rate
	P–mode short rate

Horizon
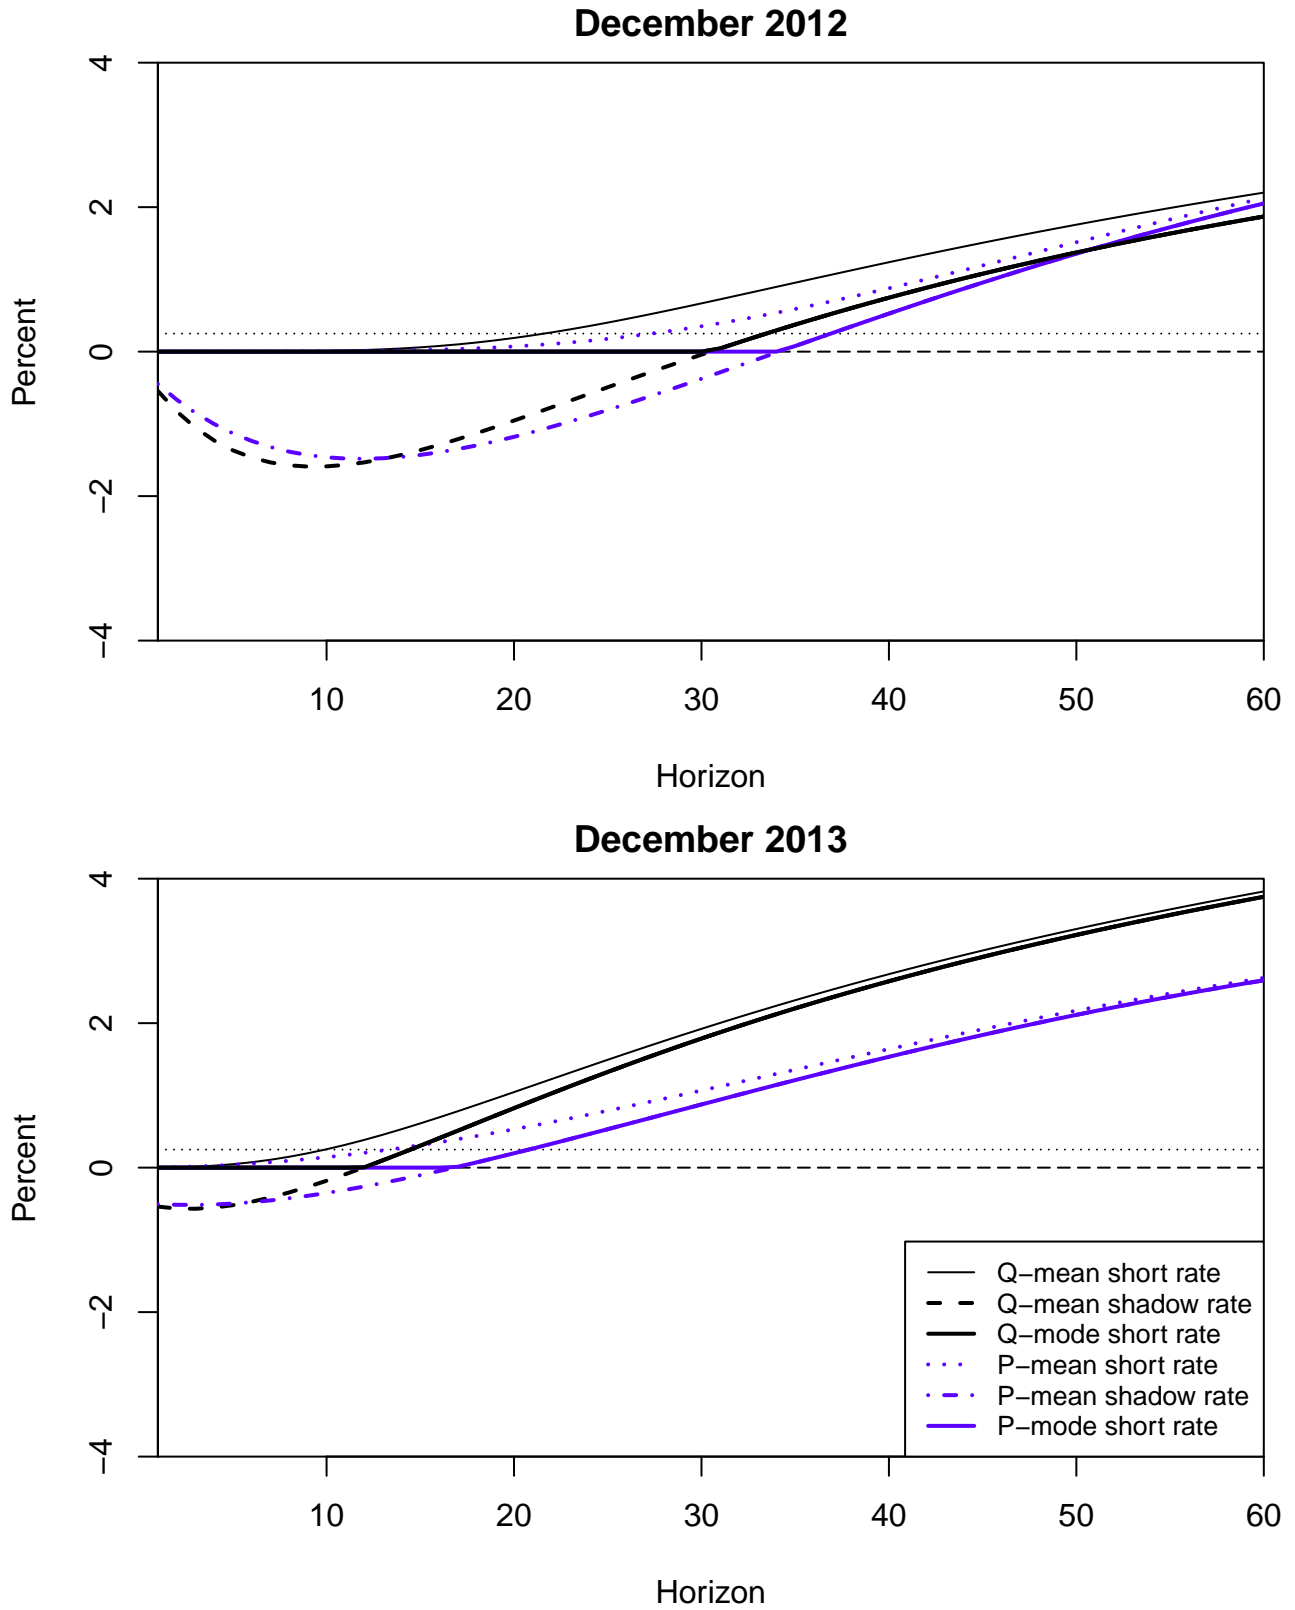
Notes: Mean and modal paths of future short rates, under real-world ($\mathbb{P}$) and risk-neutral ($\mathbb{Q}$) probability measure. The mean path under $\mathbb{Q}$ approximately corresponds to fitted forward rates. Model: $MZ(2)$.

December 2012

December 2013
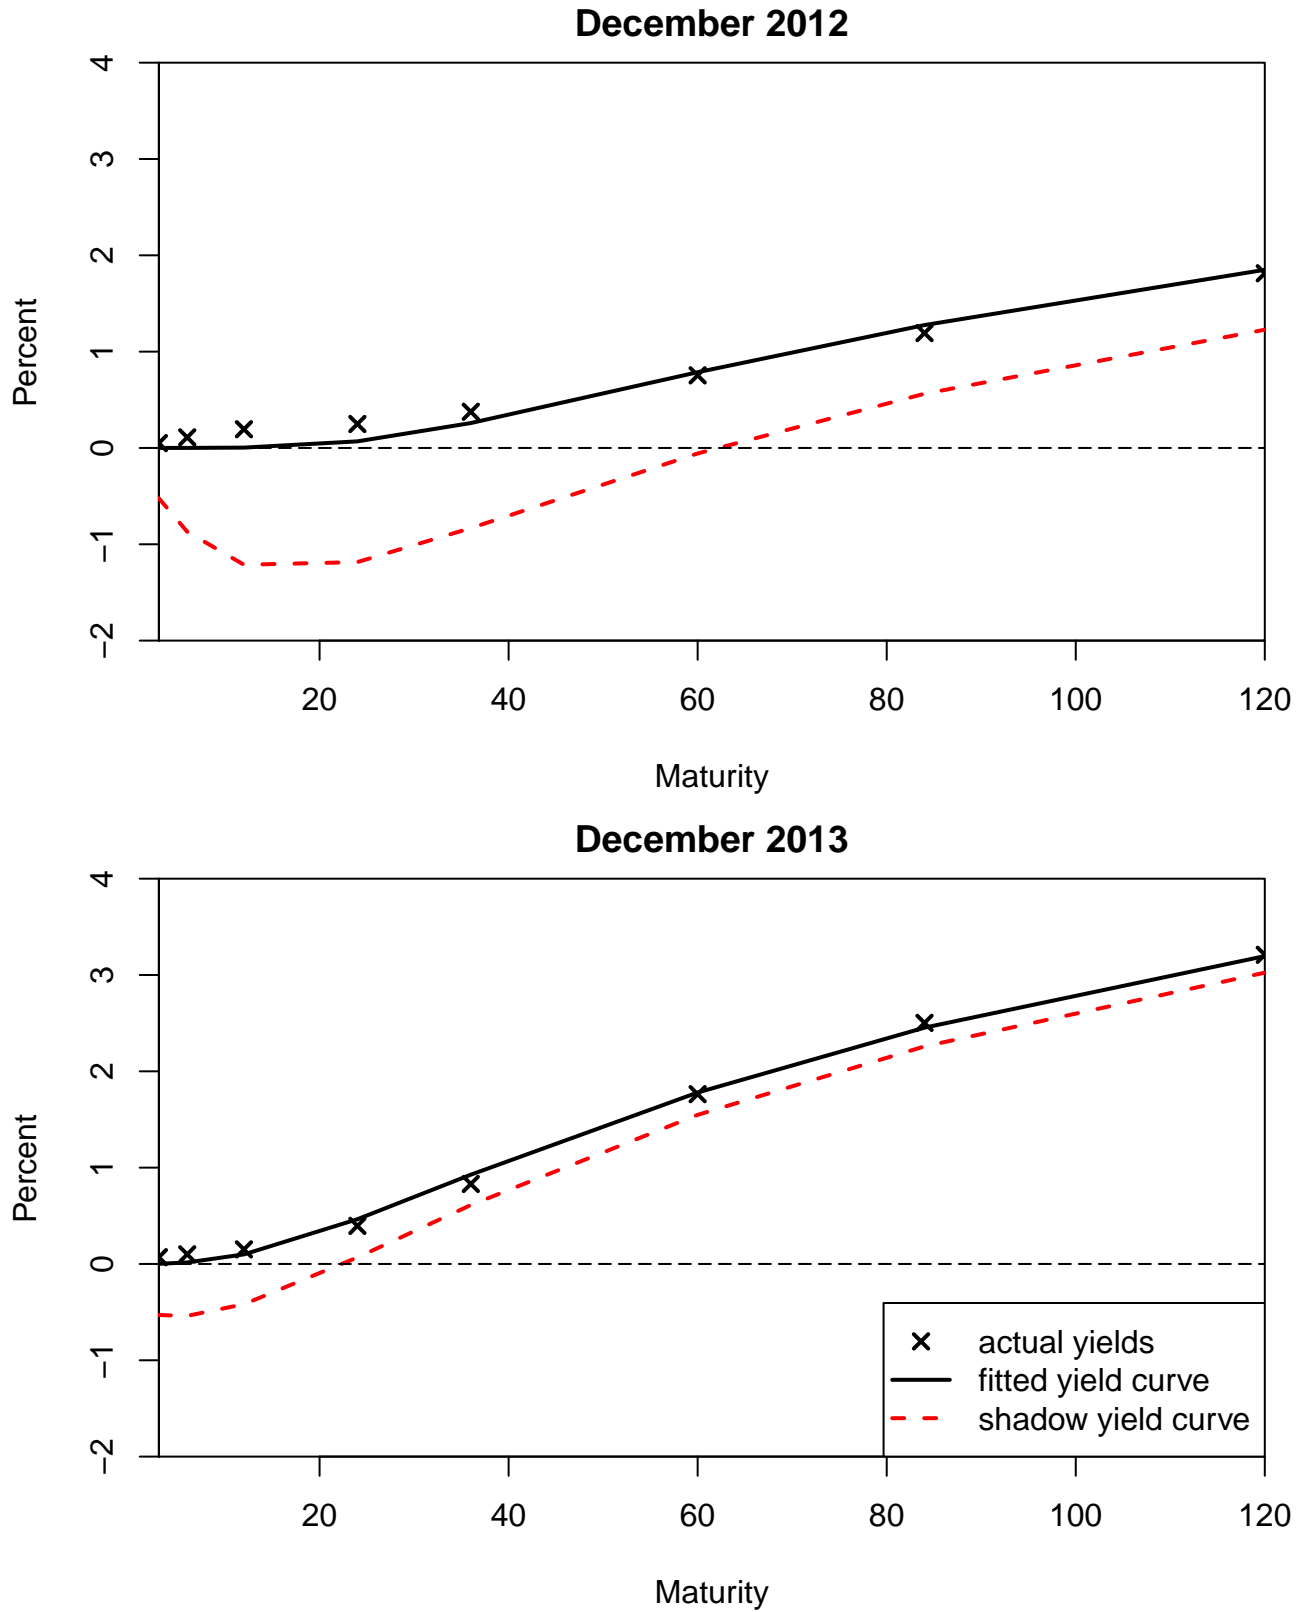
Notes: Actual, fitted, and shadow yield curves on select dates. Model: $MZ(2)$.

Figure 7: Ten-year yield and ZLB wedge

**Shadow yield vs. fitted yield**

ZLB wedge
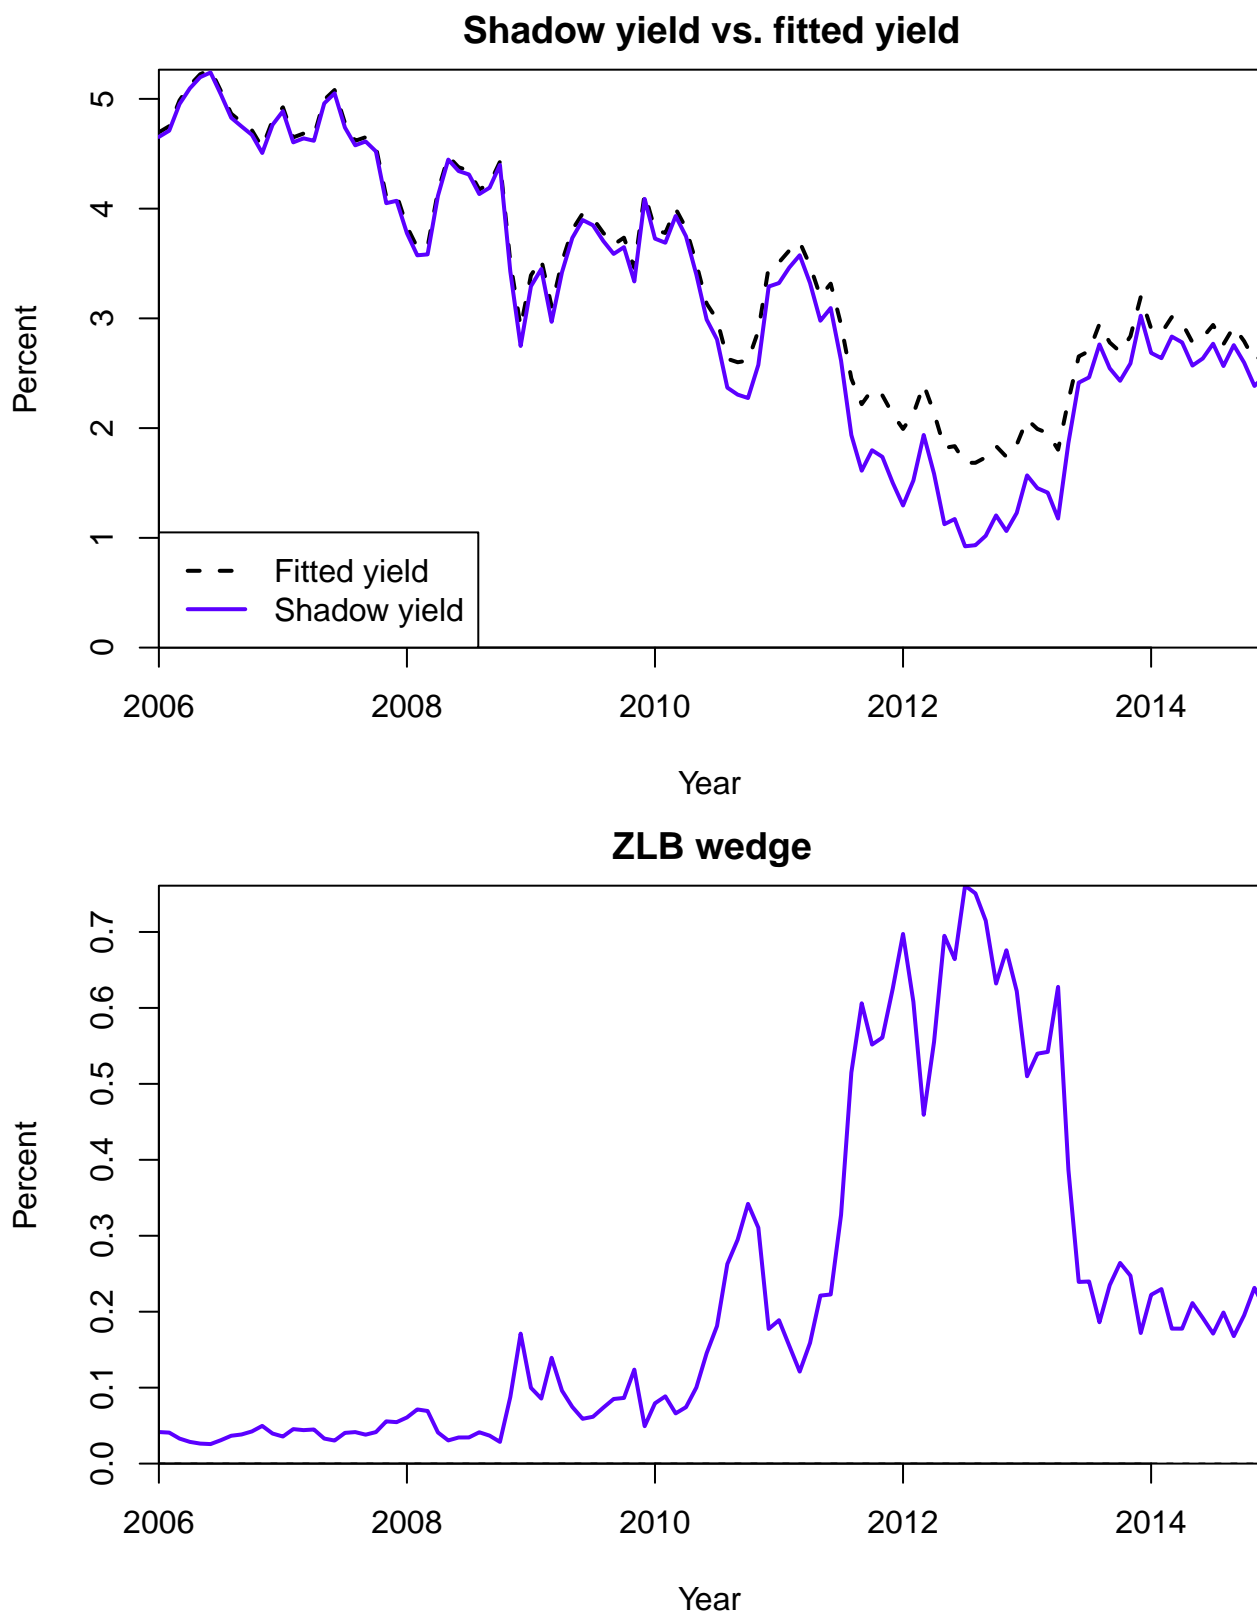
Notes: The top panel shows the fitted ten-year yield and the corresponding shadow yield. The bottom panel shows the difference between these two yields. Model: $MZ(2)$

Figure 8: Distribution of liftoff horizon
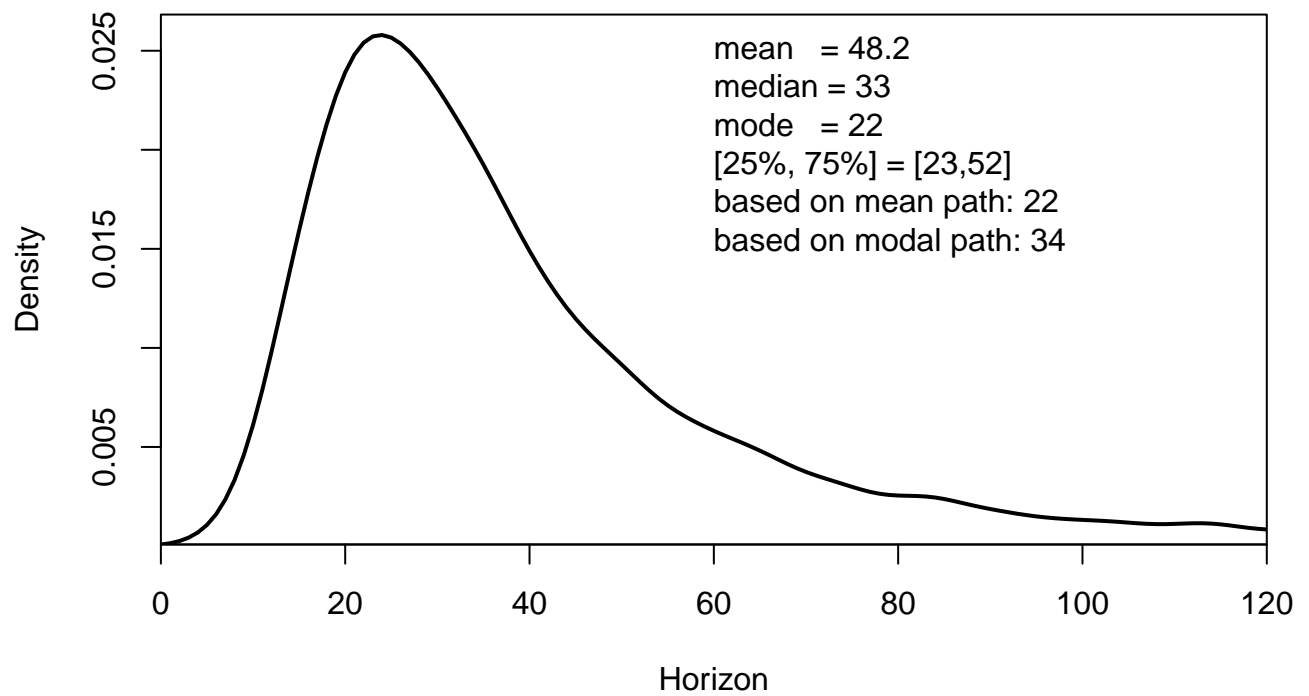
Notes: Density (under the risk-neutral probability measure) for the distribution of the liftoff horizon (in months), on December 31, 2012. Model: $MZ(2)$.

Figure 9: Liftoff
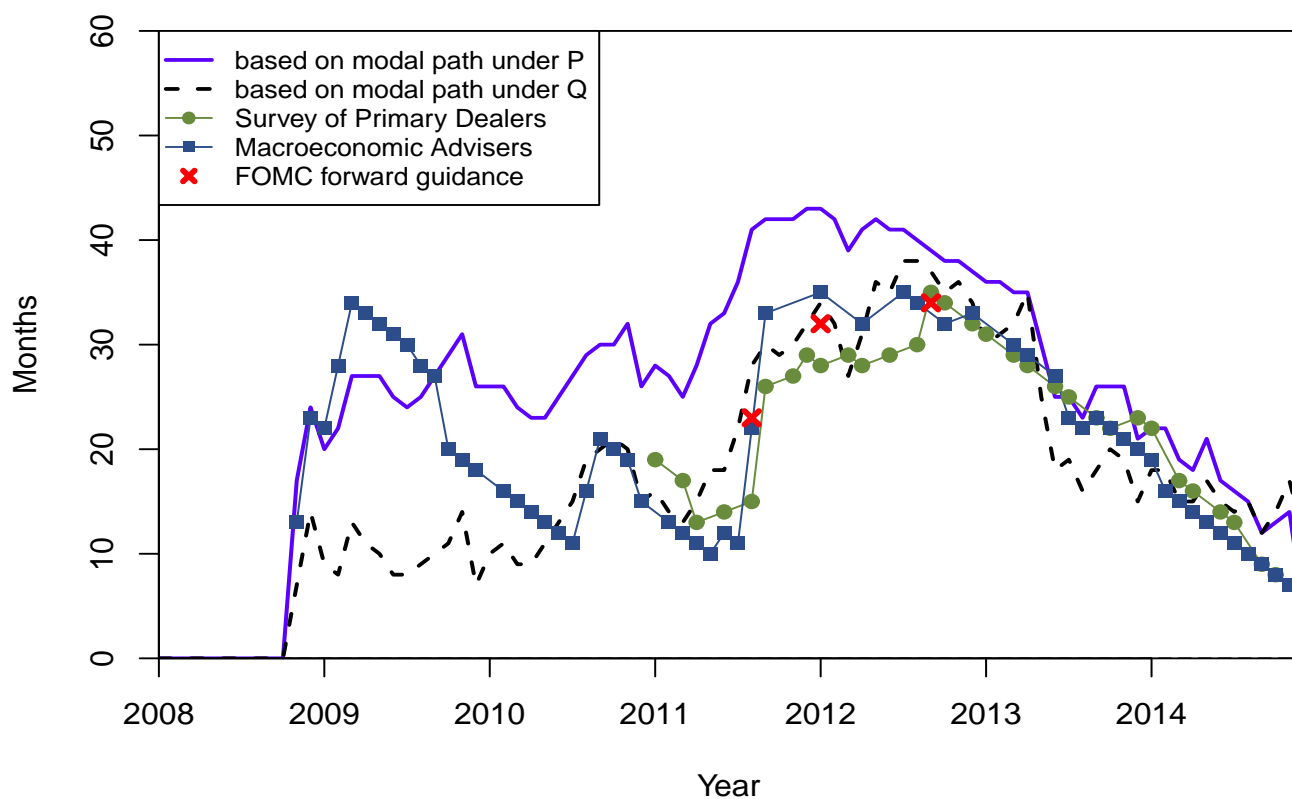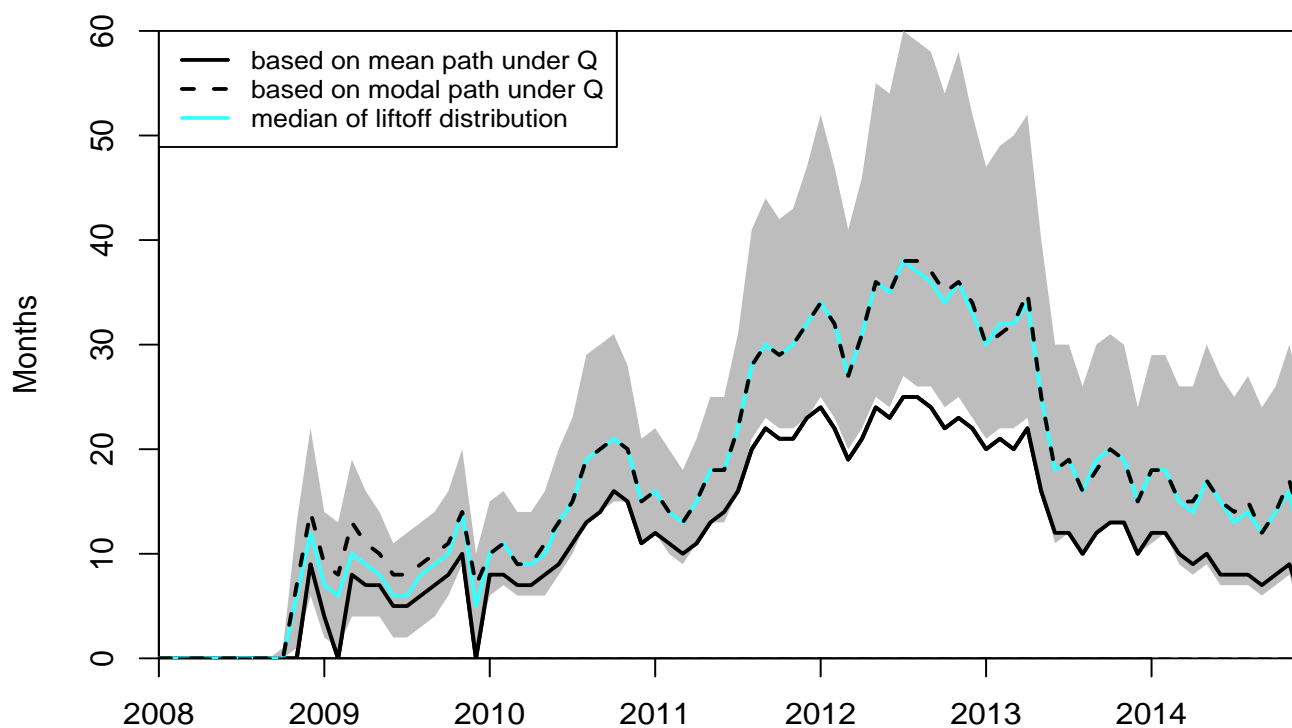
Notes: The top panel compares liftoff estimates from $MZ(2)$ based on the modal and mean paths, and the median and interquartile rang (shaded area) of the liftoff distribution (all under $\mathbb{Q}$). The bottom panel compares modal-path estimates under $\mathbb{Q}$ and $\mathbb{P}$ to liftoff estimates from the Survey of Primary Dealers (median response) and from Macroeconomic Advisers, and to the FOMC's calendar-based forward guidance. Period: January 2008 to December 2014.

Figure 10: Pace of tightening
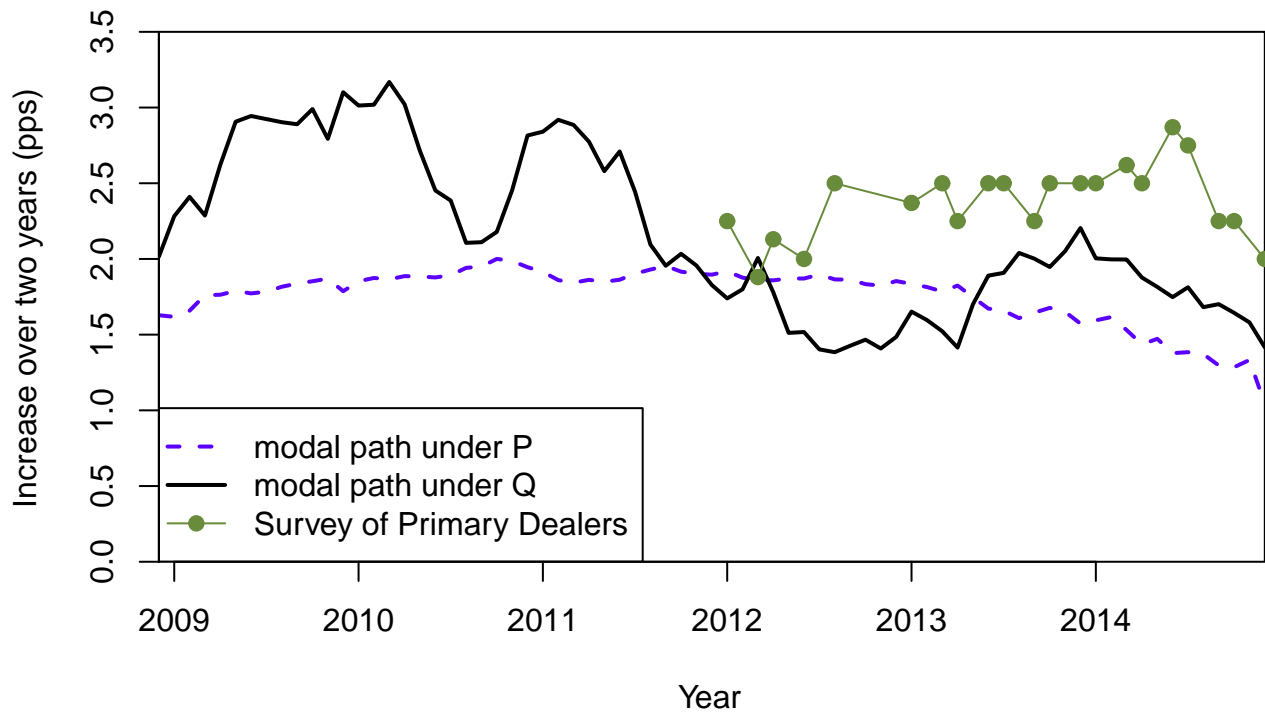
Notes: Model-based estimates of the pace of tightening after policy liftoff, measured as the increase in the modal path under the $\mathbb{Q}$- and the $\mathbb{P}$-measure during the subsequent two years, in percentage points. Also shown is the pace of tightening implied by the policy path reported in the Survey of Primary Dealers (median response). Model: $MZ(2)$. Period: December 2008 to December 2014.