Estimation of Multi-Factor Shadow-Rate Term Structure Models

PRELIMINARY DRAFT

Don H. Kim   Marcel Priebsch*

October 9, 2013

Abstract

Using recently developed yield approximation methodology, we estimate a three-factor Gaussian shadow-rate term structure model of the U.S. Treasury yield curve on a sample of yields and survey forecasts from 1990 to 2013. For comparison, we also estimate a standard Gaussian term structure model, both for the pre-crisis period 1990 to 2008, and for the full sample through 2013. We find that the shadow-rate model performs better along a number of dimensions, though it still shows some patterns indicative of misspecification.

*Don H. Kim (`don.h.kim@frb.gov`) and Marcel Priebsch (`marcel.a.priebsch@frb.gov`) are with the Board of Governors of the Federal Reserve System, Washington, D.C. The analysis and conclusions set forth in this paper are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors of the Federal Reserve System.

1 Introduction

Affine-Gaussian term structure models are workhorse models that have proved useful in many applications, ranging from a simple decomposition of yields into expectations and term premium components (e.g., Kim and Wright, 2005) to analysis of the role of macroeconomic variables in the yield curve (e.g., Ang and Piazzesi, 2003; Joslin, Priebsch, and Singleton, 2013). U.S. short-term interest rates have been effectively zero since late 2008, and this has made the use of affine-Gaussian models—as well as other models that do not respect the zero lower bound (ZLB) on nominal yields—potentially problematic.

Among "ZLB models," i.e., models that respect the zero lower bound, Black's (1995) model with a Gaussian shadow-rate process has several conceptually attractive features. For example, during periods when the short rate is sufficiently above the zero bound, the shadow-rate model behaves approximately like an affine-Gaussian model, whose empirical properties are well understood. Furthermore, in many scenarios the FRB/US model used by the Federal Reserve Board for policy rate simulation assumes a version of the Taylor rule with truncation at the zero bound,[1] which intuitively corresponds to Black-model-type dynamics of the short rate.[2]

However, there have been relatively few empirical studies with Black's shadow-rate model. Ichiue and Ueno (2007) and Kim and Singleton (2012) have analyzed Japanese yield data with two-factor shadow-rate models and obtained encouraging results, but the case of three factors (which has been a typical dimension of affine-Gaussian models hitherto studied) have only been investigated in recent efforts by Christensen and Rudebusch (2013), Bauer and Rudebusch (2013) and Xia (2013). Part of the

[1]See, for example, Chung, Laforte, Reifschneider, and Williams (2012).

[2]However, as discussed below, this does not necessarily imply the reverse, namely that the model-implied shadow rate can be used to gauge the accommodativeness of monetary policy.

difficulty has been the computational hurdle in obtaining fast and accurate formulas for bond prices in multi-factor shadow-rate models. Using a new method (proposed in Priebsch, 2013) that allows accurate and (relatively) fast computation of bond prices in general N-factor Gaussian shadow-rate models, in this paper we investigate the empirical performance of a Gaussian shadow-rate model and an affine-Gaussian model as applied to U.S. data, including the post-2008 period. We are particularly interested in addressing the following questions:

1. Clearly the affine-Gaussian model fails to capture the time-varying volatility of short- and intermediate-term yields induced by the zero bound. Despite this limitation with the *second* moment, could it still do a reasonable job of capturing *first* moment properties, such as the expected path of the short rate embedded in the yield curve? This question is of practical interest in view of the affine-Gaussian models' tractability and continued widespread use during the ZLB period (e.g., Li and Wei, 2013).

2. Along what key dimensions does the shadow-rate model improve upon affine-Gaussian models? Does the shadow-rate model perform sufficiently well, or is there evidence pointing to the need for alternative models of the ZLB, such as regime-switching models?[3] In particular, we discuss the apparent disappearance of the "level" principal component (more precisely, the collapse of "level" and "slope" principal components into a single factor) during the ZLB period, and ask whether this represents a nonlinearity that the shadow-rate model can capture, or a more nontrivial kind of structural break that calls for a different class of models.

The remainder of this paper is organized as follows. In Section 2 we briefly describe

[3]For example, Dai, Singleton, and Yang (2007).

the model. In Section 3, we discuss our data. As in Kim and Orphanides (2012), we supplement yield data with survey forecasts of a short-term interest rate in order to help overcome the small sample problems in estimation with yield data alone. Principal components analysis of yield data is also discussed there. Section 4 describes the QML estimation method based on the unscented Kalman filter used in this paper, and Section 5 describes our empirical results. Key empirical findings are as follows:

1. The standard latent-factor affine-Gaussian model (non-ZLB model) with three factors can fit most of the survey data on short-term rate forecasts reasonably well. However, in trying to fit the zero-bound related patterns in the near-term forecasts of the short-term rate, the estimated model loses lot of persistence; as a result, the model-implied long-horizon forecast of the short-term rate can be unrealistic. Furthermore, the model can produce unreasonable implications for variables that are not explicitly fitted, such as the expected path of the 10-year yield.

2. The estimated shadow-rate (ZLB) model does well in capturing key first- and second-moment properties. It matches survey forecasts of the short-term rate well, and also captures zero-bound-induced volatility compression features. Furthermore, the ZLB model produces sensible forecasts for longer-maturity yields, whose survey forecasts were not used in estimation. In addition, the ZLB model outperforms the affine-Gaussian model in one-month-ahead (out-of-sample) forecasts during the ZLB period.

3. If the shadow-rate model is well-specified, the implied VAR(1) dynamics of the factors driving the shadow rate process should not exhibit a structural break, even with the ZLB period included. In order to investigate this, we analyze the filtered state variables and examine whether they are consistent with the

estimated dynamics. We find that the innovation vectors implied by the filtered state variables in the shadow-rate model display non-i.i.d. patterns during the ZLB period, suggesting that potential misspecification remains in the shadow-rate model.

2 Methodology

2.1 Model

We work within the standard, continuous-time setup with N latent Gaussian factors, though we consider two different specifications for the short rate (the instantaneous interest rate): The usual affine-Gaussian specification, and the shadow-rate specification that respects the ZLB.

Let $W_t^{\mathbb{P}}$ be N-dimensional standard Brownian motion on a complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with canonical filtration $\{\mathcal{F}_t\}_{t\geq 0}$. Assume there is a pricing measure $\mathbb{Q}$ on $(\Omega, \mathcal{F})$ that is equivalent to $\mathbb{P}$, and denote by $W_t^{\mathbb{Q}}$ Brownian motion under $\mathbb{Q}$ as derived from Girsanov's Theorem (Karatzas and Shreve, 1991). Suppose the factors (or states) representing uncertainty underlying term-structure securities follow the multivariate Ornstein-Uhlenbeck process

$$dX_t = (K_0^{\mu} + K_1^{\mu} X_t)dt + \Sigma dW_t^{\mu} \tag{1}$$

were $\mu \in \{\mathbb{P}, \mathbb{Q}\}$. Let the Gaussian short rate be

$$s_t = \rho_0 + \rho_1 \cdot X_t. \tag{2}$$

Since X_t is a Gaussian process (Karatzas and Shreve, 1991), it follows from (2) that the short rate s_t takes on strictly negative values with strictly positive probability. To modify the model in a way that accounts for the zero lower bound on nominal yields, Black (1995) proposes to think of s_t as a *shadow* short rate and define the observed ZLB short rate as the shadow rate censored at zero:[4]

$$r_t = \max\{\rho_0 + \rho_1 \cdot X_t, 0\} = \max\{s_t, 0\}. \tag{3}$$

The arbitrage-free time t price of a zero-coupon bond maturing at time T is given by

$$P_t^T = \mathbf{E}_t^{\mathbb{Q}}\left[\exp\left(-\int_t^T q_u\, du\right)\right] \tag{4}$$

where q is either the Gaussian short rate s or the ZLB short rate r. The associated zero-coupon bond yield is defined as

$$y_t^T = -\frac{\log P_t^T}{T-t}. \tag{5}$$

Bond prices (and hence yields) can equivalently be expressed in terms of forward rates:

$$P_t^T = \exp\left(-\int_t^T f_t^s\, ds\right) \qquad \Leftrightarrow \qquad f_t^T = -\frac{d}{dT}\log P_t^T \tag{6}$$

where f_t^T denotes the instantaneous time T forward rate effective at time t.

[4]For simplicity of exposition, in this section we follow Black (1995) and set the lower bound *equal* to zero. In practice, the empirical lower bound may be a small (negative or positive number), say $r_{\min}$. All derivations below are easily modified to accommodate a non-zero lower bound (see Priebsch, 2013), and in Section 4 we treat $r_{\min}$ as a free parameter.

2.2 Bond Pricing

In the Gaussian model, zero-coupon bond prices take on the standard exponential affine form,

$$P_t^T = e^{A(T-t)+B(T-t)\cdot X_t}$$

where A and B follow a system of ordinary differential equations in terms of the model parameters (see Duffie and Kan, 1996). By (5), Gaussian yields are affine functions of the states X_t, with loadings depending only on time to maturity $T-t$ and the model parameters.

In the ZLB shadow-rate model, yields are nonlinear functions of X_t, and no equally convenient expressions for bond prices and yields exist. Several approximation schemes have been proposed: Gorovoi and Linetsky (2004) show that in a one-factor model, yields can be computed by an eigenfunction expansion, but this method does not generalize to multiple factors. Ichiue and Ueno (2007) approximate bond prices by a variant of the binomial tree familiar from option pricing, and Kim and Singleton (2012) numerically solve a partial differential equation in terms of t and $x \in \mathbb{R}^N$, but both approaches are subject to the curse of dimensionality so that these studies consider models with no more than $N = 2$ factors. Christensen and Rudebusch (2013) estimate three-factor Nelson-Siegel models using a yield formula proposed by Krippner (2012) based on an approximate forward rate and (6), but Priebsch (2013) shows that this method can give yields that deviate from arbitrage-free yields by more than five basis points in a realistic empirical setting, and the approximation error is largest precisely when the ZLB is a binding constraint. Priebsch (2013) proposes a method to approximate arbitrage-free yields in the Gaussian shadow-rate model by a second-order cumulant-generating-function expansion.[5] He

[5]Independently, Ichiue and Ueno (2013) propose a method equivalent to a first-order variant of

demonstrates that this method is accurate to approximately one half of a basis point, both during normal times and when the ZLB is binding, and that it is sufficiently fast to be computationally feasible in estimation. We use this method to approximate ZLB yields below.

3 Data

3.1 Yields

We use end-of-month zero-coupon U.S. Treasury yields from January 1990 through June 2013, for maturities of six months, one to five, seven, and 10 years. Zero yields are extracted from Treasury bills and coupon-bearing notes and bonds in the CRSP U.S. Treasury Database, using the unsmoothed Fama and Bliss (1987) methodology.[6]

We do not include yields at the very short end of the yield curve as these tend to be heavily influenced by idiosyncratic money-market factors (see Duffee, 1996). For illustration, Figure 1 plots raw quotes on off-the-run Treasury securities for August 16, 2013. The blue line corresponds to the level of the ZLB estimated by our shadow-rate model below (see Table 3), which falls roughly in the middle of the current federal funds rate target range of 0 to 25 basis points. Yields at maturities exceeding six months appear to line up along a smooth curve roughly asymptoting towards the level of the ZLB we estimate. In contrast, yields at the very short end appear disconnected from this curve and notably more dispersed.

Table 1 displays the loadings on yields in the construction of their first three principal components (PCs), as well as the percentage of total sample variation in yields explained by each principal component. Principal components are constructed sep-

the same approach.

[6] We are grateful to Anh Le for providing the code for this procedure.

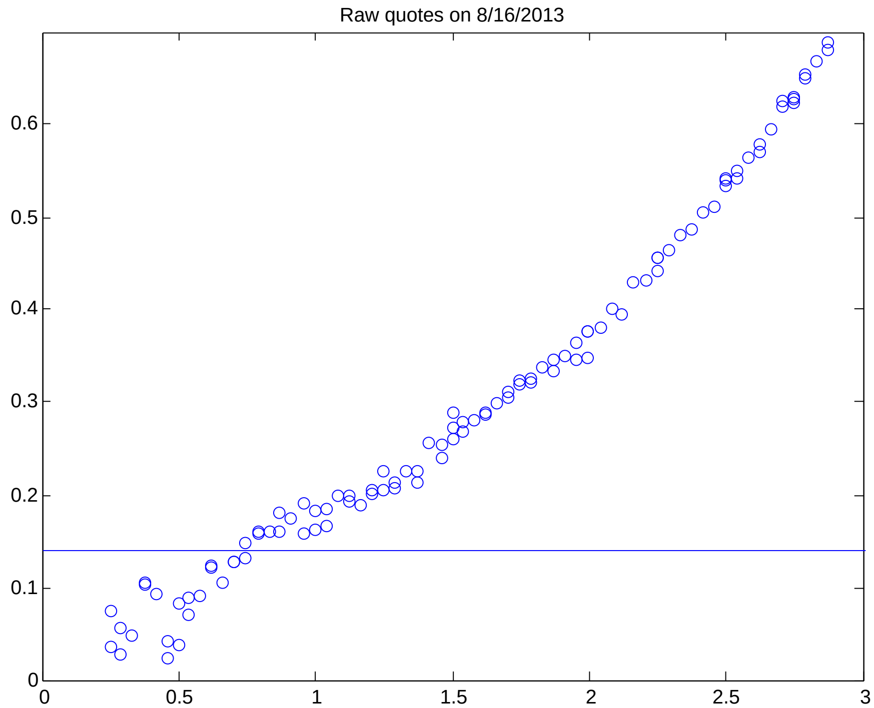

Figure 1: Raw quotes on off-the-run Treasury securities for August 16, 2013. The blue line corresponds to the level of the ZLB estimated by our shadow-rate model.

	PC1	PC2	PC3
6m	0.3817	−0.5075	0.5046
1y	0.3937	−0.4240	0.1291
2y	0.3900	−0.1580	−0.3475
3y	0.3715	0.0189	−0.3885
4y	0.3556	0.1601	−0.3032
5y	0.3381	0.2688	−0.1608
7y	0.3079	0.4071	0.1484
10y	0.2715	0.5231	0.5632
% Var	93.4877	6.1775	0.2587

(a) Sub-sample January 1990 to November 2008

	PC1	PC2	PC3
6m	0.0346	−0.1225	0.3665
1y	0.0729	−0.1948	0.4811
2y	0.1743	−0.4953	0.1813
3y	0.2831	−0.4484	0.2605
4y	0.3671	−0.3357	−0.1925
5y	0.4313	−0.1227	−0.4484
7y	0.5122	0.1690	−0.2807
10y	0.5476	0.5868	0.4653
% Var	97.3202	1.5236	0.9203

(b) Sub-sample December 2008 to June 2013

Table 1: Loadings on yields in the construction of their first three principal components (PCs), as well as the the percentage of the total yield variation explained by each principal component.

arately for two sub-sample periods: January 1990 to November 2008 (a period with interest rates at "normal" levels), and December 2008 to June 2013 (when short-term yields were constrained by the ZLB). As Table 1a shows, in normal times the first three principal component loadings take their usual "level," "slope," and "curvature" form (see Litterman and Scheinkman, 1991). The three principal components explain successively smaller fractions of the total variance in yields, and together account for over 99.9 percent of yield variation. In contrast, the loadings in the ZLB period (Table 1b) look notably different. The first principal component looks more like a slope factor, whereas the second and third principal components both have curvature qualities, with humps at different maturities. At first sight, the apparent disappearance of a level factor in the ZLB period might suggest that yield variation can be adequately captured by only two factors (see Ichiue and Ueno (2013) for an argument along these lines). However, Table 1b also shows that the first two principal components account for only 98.8 percent of total yield variation, and that all three principal components are still needed to explain a similar fraction as in normal times (in fact, in the ZLB period the third principal component explains a notably larger share of yield variation than in normal times).

In addition, the substantial shift in the pattern of principal component loadings in Table 1 suggests that a linear factor model (such as an affine-Gaussian term structure model) will have a hard time adequately capturing yield behavior in the normal and ZLB periods simultaneously. Table 2 shows the sub-sample correlation matrix for the ZLB period of the first three principal component scores constructed over the entire sample period. By construction, principal component scores are uncorrelated over their full sample. However, as Table 2 shows, over the ZLB sub-sample period, there is substantial correlation between principal component scores, especially between the first and second PCs. This suggests that the level and slope factors indeed largely

	PC1	PC2	PC3
PC1	1.0000	0.9780	0.4345
PC2	0.9780	1.0000	0.5220
PC3	0.4345	0.5220	1.0000

Table 2: Sub-sample correlation for December 2008 to June 2013 of the first three principal component scores, constructed over the entire sample (January 1990 to June 2013).

collapse into a single factor during the ZLB period (although recall from the discussion above that this does not imply that fewer factors are sufficient to adequately capture yield variation near the ZLB).

3.2 Surveys

We augment the yield data with survey forecasts from Blue Chip, interpolated to constant horizons of one to four quarters (available monthly), as well as annually out to five years and for five-to-10 years (available every six months). As discussed by Kim and Orphanides (2012), this potentially leads to more precise estimates of the parameters governing the data-generating distribution $\mathbb{P}$. Survey forecasts are subject to the same lower-bound constraint as yields,[7] but model-implied survey forecasts are substantially simpler to compute than yields, as forecasters report their expectation of the *arithmetic mean* of future observed short rates, while yields are computed as a geometric mean due to the effect of compounding. For details, we refer to Priebsch (2013).

[7]This follows from equivalence of the measures $\mathbb{P}$ and $\mathbb{Q}$, and more fundamentally from the absence of arbitrage.

4 Estimation

While the statistical properties of the term structure model laid out in Section 2 are formulated in terms of the latent state vector X_t, the data actually observed by the econometrician consist of yields, y_t, and survey expectations, z_t (see Section 3). To bridge this gap and obtain estimates of the model's parameters $\theta = (K_0^{\mathbb{P}}, K_1^{\mathbb{P}}, K_0^{\mathbb{Q}}, K_1^{\mathbb{Q}}, \rho_0, \rho_1, \Sigma)$, the term-structure literature has pursued two conceptually related but computationally distinct estimation strategies.

The first, attributed to Chen and Scott (1993), is to make X_t effectively observable by assuming that an invertible function of yields and survey forecasts can be observed by the econometrician without error. Commonly, a subset of yields, or certain linear combinations of yields (such as low-order principal components), are assumed to be perfectly observable. Under this assumption, we can back out the implied state vector X_t, and compute the model-implied conditional distribution of all yields and survey expectations through a change of variables. The model parameters θ can then be estimated by maximum likelihood or a similar method.

The second estimation strategy, is to assume all yields and survey expectations are subject to some measurement error. As a result, X_t cannot be perfectly inferred from observables but must be filtered. The parameters θ are then estimated as part of a joint estimation and filtering problem.[8]

Selecting between the two estimation strategies is often a matter of convenience. The first method—inverting the mapping between state vector and observables—has computational advantages when the term structure model is affine, and a linear combination of yields is assumed perfectly observable. In this case, only a single matrix needs to be inverted. Moreover, linear combinations of yields (such as principal

[8] An early reference discussing this approach is Duan and Simonato (1999).

components) may effectively "diversify away" most measurement error, so that the assumption of perfect observability is empirically tenable (see Joslin, Le, and Singleton, 2013). On the other hand, when the mapping between states and observables is nonlinear—as in the case of a shadow-rate model—, numerical inversion is computationally more costly. Moreover, the assumption that a given linear combination of yields can be observed without error is more difficult to justify. For instance, when the shadow rate is close to zero or negative, model-implied short-term yields are substantially less sensitive to changes in X_t than longer-term yields. Consequently, intuition suggests that shorter-term yields have a lower signal-to-noise ratio and might be less informative about X_t. The small loadings on short-term yields in the first principal component in Table 1b is an empirical manifestation of this phenomenon.

Both arguments set forth in the previous paragraph—computational complexity and signal-to-noise ratio—favor the filtering approach for estimation during the ZLB period. We therefore pursue this approach below.

When discretely sampled at intervals $\Delta t > 0$, the state vector X follows a first-order Gaussian vector autoregression,

$$X_{t+\Delta t} = m_{0,\Delta t} + m_{1,\Delta t} X_t + \varepsilon_t \tag{7}$$

where $\varepsilon_t \sim N(0, \Omega_{\Delta t})$, and $m_{0,\Delta t}$, $m_{1,\Delta t}$, and $\Omega_{\Delta t}$ can be computed analytically as functions of the model parameters. Equation (7) represents the *transition equation* of the filtering problem.

Next, denote by $H_y : \mathbb{R}^N \times \Theta \mapsto \mathbb{R}_+^{M_Y}$ the mapping from states X and parameters θ to model-implied yields y, and by $H_z : \mathbb{R}^N \times \Theta \mapsto \mathbb{R}_+^{M_Z}$ the analogous mapping from states and parameters to model-implied survey forecasts z. To simplify notation, denote the stacked mapping $\left(H_y^\top, H_z^\top\right)^\top$ by H. In the Gaussian model, the mapping

H is linear, while in the ZLB shadow-rate model, it is nonlinear (see Section 2). If we assume that all yields and survey expectations are observed with i.i.d. additive Gaussian errors, we obtain the *observation equation*

$$\begin{pmatrix} y_t \\ z_t \end{pmatrix} = H(X_t) + e_t. \tag{8}$$

Together, equations (7) and (8) form a non-linear filtering problem.

The simple (linear) Kalman filter—optimal when measurement and observation equation are linear and all shocks are Gaussian—has been modified in a number of ways to accommodate nonlinearity. Adapted to the present problem, the challenge lies in efficiently computing a forecast and forecast error for (y_t, z_t) given a forecast of X_t.[9] Previous studies of zero-bound term structure models (Kim and Singleton, 2012; Christensen and Rudebusch, 2013; Ichiue and Ueno, 2007) have relied on the *extended Kalman filter*, in which the observation equation (8) is linearized by a first-order Taylor expansion around the conditional mean of X_t:[10]

$$\begin{pmatrix} y_t \\ z_t \end{pmatrix} \approx H(\mathbf{E}_{t-1}[X_t]) + H'(\mathbf{E}_{t-1}[X_t]) \cdot (X_t - \mathbf{E}_{t-1}[X_t]) + e_t. \tag{9}$$

The conditional moments of (y_t, z_t) can then be approximated easily based on (9) and the conditional moments of X_t. However, due to the linear approximation (which effectively treats (y_t, z_t) as conditionally Gaussian random vectors), the extended Kalman filter can be numerically unstable and may fail to converge.

The *unscented Kalman filter*, proposed by Julier, Uhlmann, and Durrant-Whyte

[9] The forecast of (y_t, z_t) is then compared to the observed value, and this information is used to "update" the filtered value of X_t.

[10] Note this method is exact, and reduces to the standard linear Kalman filter, when H is a linear transformation.

(1995), aims to deliver improved accuracy and numerical stability relative to the extended Kalman filter, without substantially increasing the computational burden.[11] Instead of linearizing equation (8) as in (9), the unscented Kalman filter uses the "unscented transformation" (Julier and Uhlmann, 1996) to approximate the conditional moments of $H(X_t)$ directly. Once the conditional moments of $H(X_t)$ are known, computing conditional moments of (y_t, z_t) is trivial (since e_t is assumed to be i.i.d.). The unscented transformation involves a general three-step procedure for computing the moments of an arbitrary nonlinear transformation H of a random variable X: First, a set of $2N+1$ sample points (called "sigma points") around the mean of X is selected. Second, each sigma point is transformed under H. Third, the moments of $H(X)$ are computed as weighted sample moments of the transformed sigma points. The sigma points as well as weights are chosen carefully to ensure that the approximate moments of $H(X)$ are accurate to third order when X is Gaussian, and accurate to second order otherwise.[12] The order of accuracy does not depend on the nature of the nonlinearity in the transformation H.

The numerical complexity of the extended Kalman filter and the unscented Kalman filter is indeed comparable: The unscented filter requires evaluating H at a number of sigma points that is $O(N)$, while computation of $H'(X)$ by finite differences in the extended Kalman filter also requires evaluation of H at a number of points that is $O(N)$.[13,14]

In light of its superior accuracy at similar computational cost, we use the unscented

[11] A detailed treatment of the unscented Kalman filter, and a comparison to the extended Kalman filter, can be found in Wan and van der Merwe (2001).

[12] Note in this regard the unscented transformation is closely related to the numerical evaluation of an integral by Gaussian quadrature.

[13] The complexity of the extended filter may be lower when H' is known analytically.

[14] This argument also establishes that filtering states in a nonlinear setup (whether by extended or unscented Kalman filter) is typically computationally *simpler* than inverting the mapping H under the assumption that some pricing errors are zero. This is because numerical inversion of H by a method such as Newton-Raphson requires *repeated* computation of H'.

Kalman filter rather than the extended Kalman filter to estimate our shadow-rate model.[15] The algorithm is described in detail in Wan and van der Merwe (2001). As a by-product of the filtering procedure, it conveniently produces estimates of the mean and covariance matrix of (y_t, z_t) conditional on the econometrician's information set as of time $t-1$. We use these to set up a quasi–maximum likelihood function based on (8),[16] which we then maximize numerically to obtain estimates of the parameters θ as well as their asymptotic standard errors (following Bollerslev and Wooldridge, 1992). For the affine-Gaussian model, the unscented Kalman filter reduces to the usual linear Kalman filter, and the QML problem becomes regular maximum-likelihood estimation.

5 Estimation Results

Without further restrictions, the parameters θ are not econometrically identified. Invariant transformations can be applied to the latent state vector X_t, resulting in observationally equivalent models with different parameters (Dai and Singleton, 2000; Joslin, Singleton, and Zhu, 2011). To achieve identification, we impose the normalizations $\rho_1 \geq 0$, $\Sigma = 0.1 I_N$, $K_0^{\mathbb{P}} = 0$, $K_1^{\mathbb{P}}$ is lower triangular.

We estimate the Gaussian and shadow-rate models on the data set described in Section 3, using the (quasi-)maximum likelihood procedure discussed in Section 4.

Table 3 displays the estimated model parameters $\hat{\theta}$ for the shadow-rate model, as well as their asymptotic standard errors. Note the estimate of $r_{\min}$, the lower bound on the observed short rate and hence nominal yields, is 14 basis points. [17]

[15]Christoffersen, Dorion, Jacobs, and Karoui (2012) and Wu (2010) confirm that the unscented Kalman filter performs better than the extended Kalman filter in the specific setting of term structure model estimation.

[16]This estimation approach is described and analyzed in Lund (1997).

[17]This does not rule out observed yields below 14 basis points, but the model would attribute

ρ_0	0.0455	$r_{\min}$	0.0014		
	(0.0015)		(0.0001)		
ρ_1	0.0304	$K_1^{\mathbb{P}}$	-0.1597		
	(0.0098)		(0.0520)		
	0.0461		0.6437	-0.9358	
	(0.0083)		(0.1845)	(0.0619)	
	0.0665		-1.1104	1.8746	-0.9358
	(0.0051)		(0.4253)	(0.2979)	(0.0619)
$K_0^{\mathbb{Q}}$	0.0075	$K_1^{\mathbb{Q}}$	-0.4930	0.8345	-0.4706
	(0.0073)		(0.2026)	(0.2169)	(0.1208)
	-0.0099		-0.0308	0.4340	-0.4814
	(0.0070)		(0.1446)	(0.1251)	(0.0799)
	-0.0389		-0.7702	1.9466	-1.3221
	(0.0167)		(0.3271)	(0.3168)	(0.1325)

Maturity	σ_Y
6m	0.0015
1y	0.0013
2y	0.0005
3y	0.0003
4y	0.0004
5y	0.0003
7y	0.0006
10y	0.0015
Average	0.0008

Maturity	σ_Z
1q	0.0013
2q	0.0005
3q	0.0010
4q	0.0015
2y	0.0027
3y	0.0025
4y	0.0026
5y	0.0030
5y–10y	0.0038
Average	0.0021

Table 3: Quasi–maximum likelihood parameter estimates (asymptotic standard errors) for the three-factor ZLB shadow-rate model (top table), and estimated standard deviations of observation errors in yields, σ_Y, and survey forecasts, σ_Z (bottom tables).

ρ_0	0.0416				
	(0.0053)				
ρ_1	0.0419	$K_1^{\mathbb{P}}$	-0.0615		
	(0.0106)		(0.0295)		
	0.0142		0.3623	-1.1091	
	(0.0083)		(0.1277)	(0.0774)	
	0.0825		-0.7910	2.3179	-1.1091
	(0.0061)		(0.3267)	(0.3075)	(0.0774)
$K_0^{\mathbb{Q}}$	-0.0322	$K_1^{\mathbb{Q}}$	-0.3920	1.5884	-1.1900
	(0.0252)		(0.2521)	(0.2896)	(0.2062)
	-0.0419		-0.0258	0.5843	-0.7268
	(0.0317)		(0.1421)	(0.1510)	(0.1274)
	-0.0976		-0.4598	2.3148	-1.8670
	(0.0565)		(0.3407)	(0.3313)	(0.1747)

Maturity	σ_Y
6m	0.0015
1y	0.0014
2y	0.0005
3y	0.0002
4y	0.0003
5y	0.0003
7y	0.0006
10y	0.0011
Average	0.0007

Maturity	σ_Z
1q	0.0013
2q	0.0005
3q	0.0011
4q	0.0015
2y	0.0027
3y	0.0026
4y	0.0028
5y	0.0029
5y–10y	0.0030
Average	0.0020

Table 4: Maximum likelihood parameter estimates (asymptotic standard errors) for the three-factor Gaussian model estimated through October 2008 (top table) and estimated standard deviations of observation errors in yields, σ_Y, and survey forecasts, σ_Z (bottom tables).

ρ_0	0.0467				
	(0.0010)				
ρ_1	0.0020	$K_1^{\mathbb{P}}$	-0.9396		
	(0.0074)		(0.1016)		
	0.0181		-1.4181	-0.9396	
	(0.0099)		(0.2320)	(0.1016)	
	0.0713		1.2217	1.3721	-0.4303
	(0.0040)		(0.1972)	(0.1647)	(0.1066)
$K_0^{\mathbb{Q}}$	0.034822	$K_1^{\mathbb{Q}}$	-0.1704	0.1381	0.1096
	(0.0416)		(0.1239)	(0.1081)	(0.0795)
	-0.5224		-1.5047	-1.5580	-1.0965
	(0.1731)		(1.0370)	(0.5770)	(0.3525)
	0.2305		-0.3141	0.5786	0.3501
	(0.0889)		(0.3661)	(0.2789)	(0.1439)

Maturity	σ_Y
6m	0.0009
1y	0.0010
2y	0.0012
3y	0.0010
4y	0.0006
5y	0.0000
7y	0.0011
10y	0.0023
Average	0.0010

Maturity	σ_Z
1q	0.0012
2q	0.0005
3q	0.0005
4q	0.0005
2y	0.0026
3y	0.0037
4y	0.0040
5y	0.0041
5y–10y	0.0056
Average	0.0025

Table 5: Maximum likelihood parameter estimates (asymptotic standard errors) for the three-factor Gaussian model estimated through June 2013 (top table) and estimated standard deviations of observation errors in yields, σ_Y, and survey forecasts, σ_Z (bottom tables).

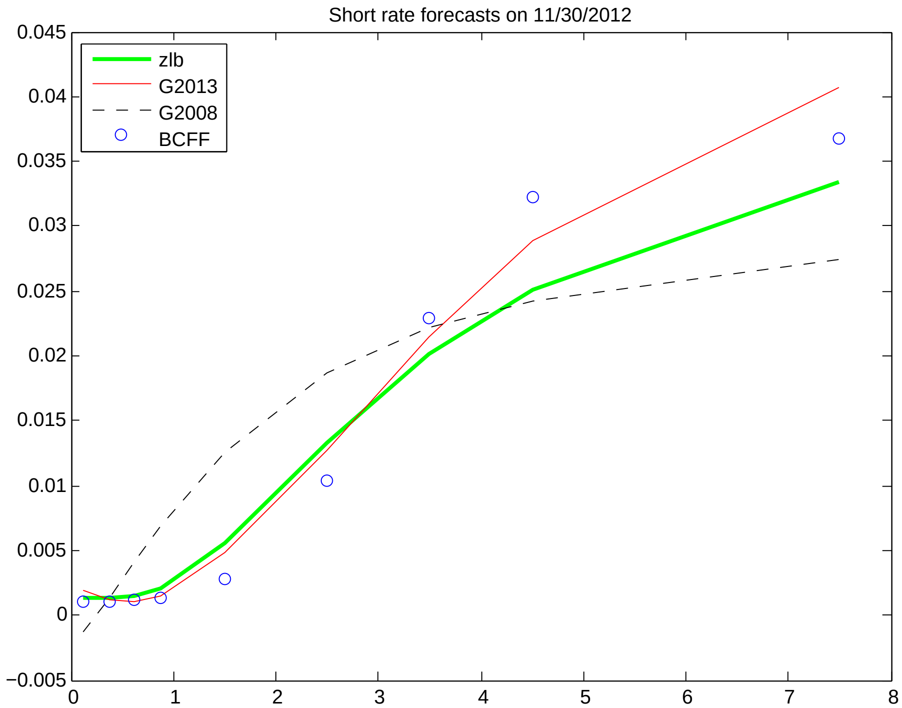

Figure 2: Model-implied and survey forecasts for the short rate as of November 2012.

The bottom tables in Table 3 show the QML-estimated standard deviations of the measurement errors in yields and survey variables (e_t in equation (8)). The average yield error is 8 basis points, and the average error in surveys is 21 basis points. For both yields and surveys, errors follow a roughly U-shaped pattern, being largest at the short and long ends.

We estimate the affine-Gaussian model both for pre-crisis sub-sample (G2008, Table 4), and for the full sample (G2013, Table 5). One notable difference between the two estimates is the much higher degree of mean reversion of the states X_t under $\mathbb{P}$ in G2013. To fit the ZLB features of the data (reflected, for instance, in near-term survey forecasts), the model is sacrificing persistence to a point at which the half life of the most persistent factor is less than two years.

---

any violation of this lower bound to observation error as opposed to fundamental drivers.

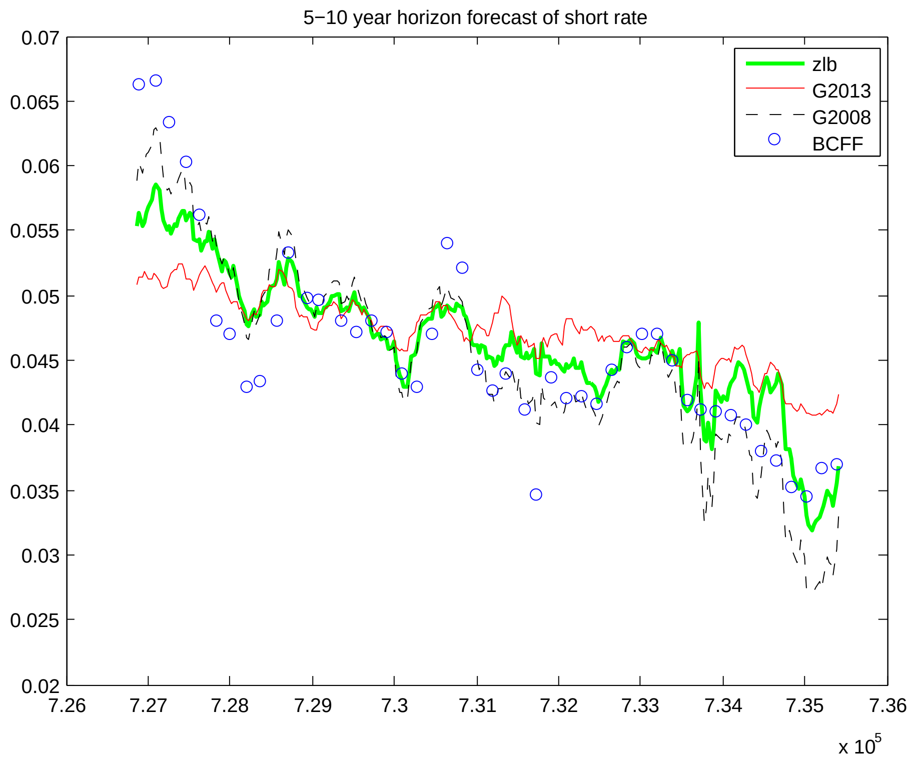

Figure 3: Model-implied and survey forecasts for the short rate five to 10 years ahead.

Figure 2 shows short-rate forecasts implied by all three of our estimated models, as well as surveys, as of November 2012. The pre-ZLB Gaussian model G2008 fits these out-of-sample surveys poorly. In particular, it predicts too fast a return to higher interest rates. In contrast, both the full-sample Gaussian model G2013 and the shadow-rate model capture the shape and level of the term structure of survey forecasts well.

However, when we look at the time series of long-term short-rate forecasts (Figure 3), it becomes apparent that G2013 is not able to generate the same degree of variability as G2008, the shadow-rate model, and surveys. This translates into a notably larger estimated fitting error as shown in the bottom right table of Table 5.

5.1 Term Premiums

The poor performance of the G2008 model in predicting the near-term path of short rates during the ZLB period (Figure 2) means that the short-horizon term premium estimates based on this model are unreliable. This can be seen from the fact that in the time series of two-year yield term premiums (top panel of Figure 4), the G2008 version is showing unreasonably negative two-year term premiums in the last two years of the sample. Conversely, the 10-year yield term premium estimates based on all three models—shadow-rate, G2013, and G2008—are quite similar, especially so for the shadow-rate and G2008 models. This suggests that long-horizon yield term premium estimates for the ZLB period generated by a Gaussian model estimated on a pre-ZLB sample of yields (such as the model estimated in Kim and Wright (2005)) appear to be more reliable than similarly generated short-horizon yield term premium estimates.

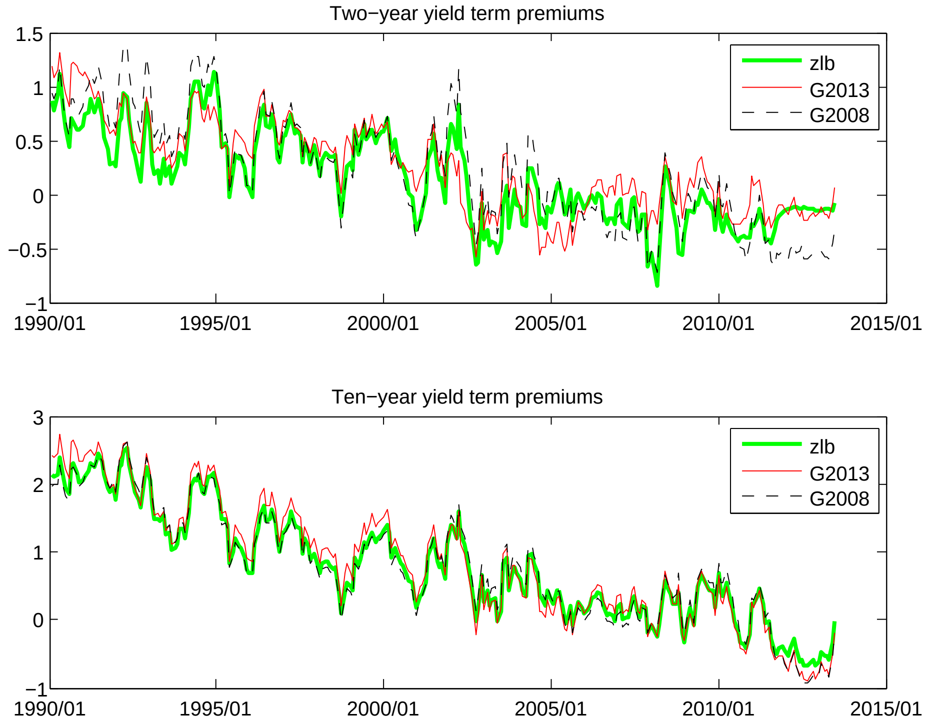

Figure 4: Model-implied yield term premiums.

5.2 Yield Forecasts

Although the Gaussian model G2013 produces yield term premiums that agree reasonably well with those from the shadow-rate model, this is partly due to the discipline imposed on the model by the survey forecasts of the short rate. If we look at other model implications not similarly disciplined, we find substantial differences between the shadow-rate model and the Gaussian model. For example, the forecast of the 10-year yield as of June 2013 (our last sample date) are quite different between the two models. In the recent episode of rising long-term rates, the G2013 model produces an unrealistic near-term projection of the 10-year rate, and a very large rise in the 10-year yield over the longer term, as can be seen in figure (Figure 5). Similarly, the G2008 model also implies a very unrealistically flat path of the 10-year yield. The difference between the models also show up in the time series of 10-year yield forecasts

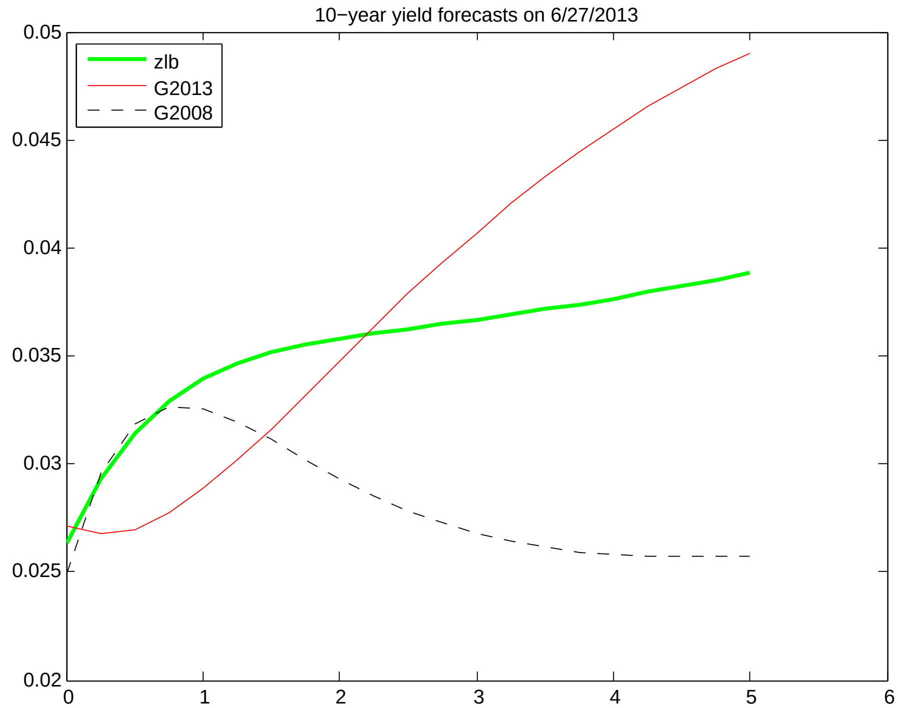

Figure 5: Model-implied forecasts of the 10-year yield as of June 2013.

during the ZLB period, as can be seen in Figure 6.

Table 6 allows more systematic comparison of the forecasting performance during the ZLB period of Gaussian and shadow-rate models. It suggests that in-sample, the Gaussian model is at least as good as the shadow-rate model, and sometimes even better. However, the out-of-sample forecasts (based on estimation up to Oct 2008) show that the shadow-rate model produces smaller forecasting errors for short- and intermediate maturity yields.

5.3 Results Specific to Shadow-Rate Model

The shadow-rate model produces quantities of interest with no (or only trivial) counterpart in the Gaussian model.

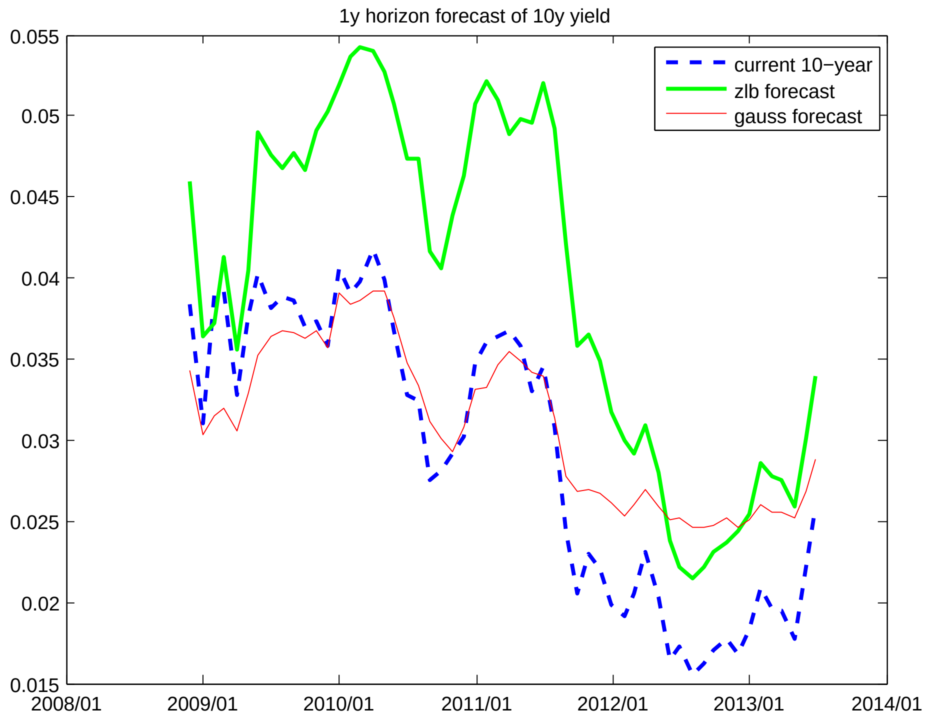

Figure 6: Model-implied forecasts of the 10-year yield one year ahead.

	6m	1y	2y	3y	4y	5y	7y	10y
Shadow-Rate	8	11	20	28	33	34	37	38
Gaussian	7	12	17	22	23	25	30	36

(a) In-sample

	6m	1y	2y	3y	4y	5y	7y	10y
Shadow-Rate	8	11	21	31	37	39	39	38
Gaussian	16	14	26	38	43	42	40	39

(b) Out-of-sample

Table 6: Root-mean-square forecasting errors (in basis points) for different yields at a one-month forecasting horizon.

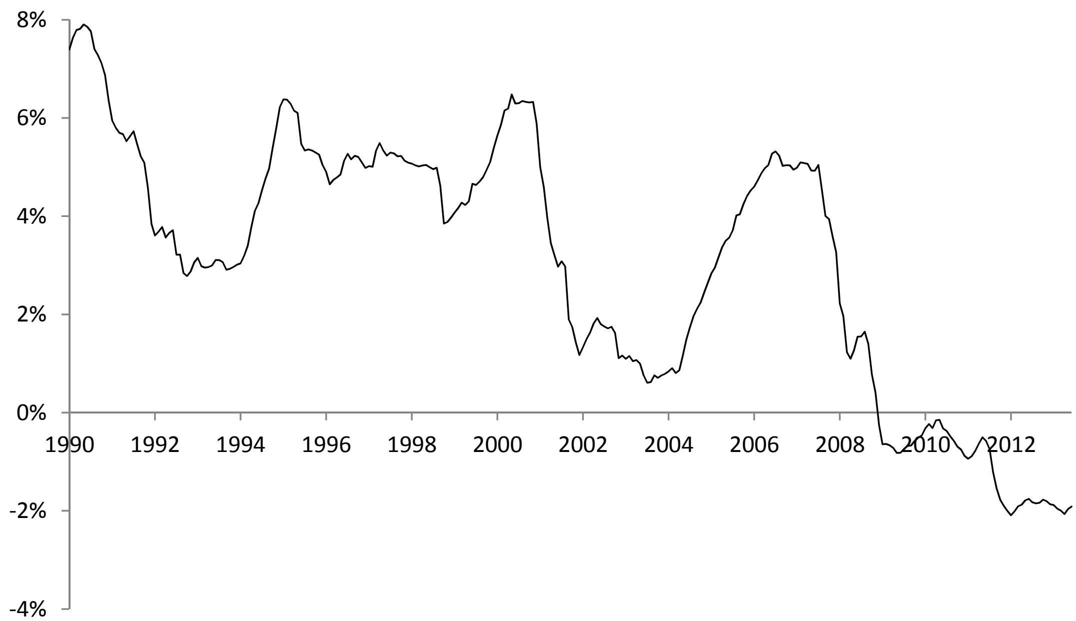

Figure 7: ZLB model-implied shadow short rate s_t based on smoothed states $X_{t|T}$.

5.3.1 Shadow Rate

Figure 7 plots the model-implied shadow short rate s_t over the sample period, based on the states implied by the Kalman smoother (that is, incorporating all information up to June 2013, the end of the sample). The shadow rate turned negative in December 2008, after the FOMC established a target federal funds rate range of 0 to 0.25 percent and the effective lower bound became binding, and has stayed negative through the end of the sample.

While, by definition, the shadow rate is equal to the observed short rate when the ZLB is not binding, we do not give a structural interpretation to the shadow rate when it is negative. Since it is unobservable, the level of the shadow rate is model specific, and under different modeling assumptions, different levels of the shadow rate may imply (close to) observationally equivalent yield curves.

Figure 8 projects the path of the short rate going forward as of June 2013, under

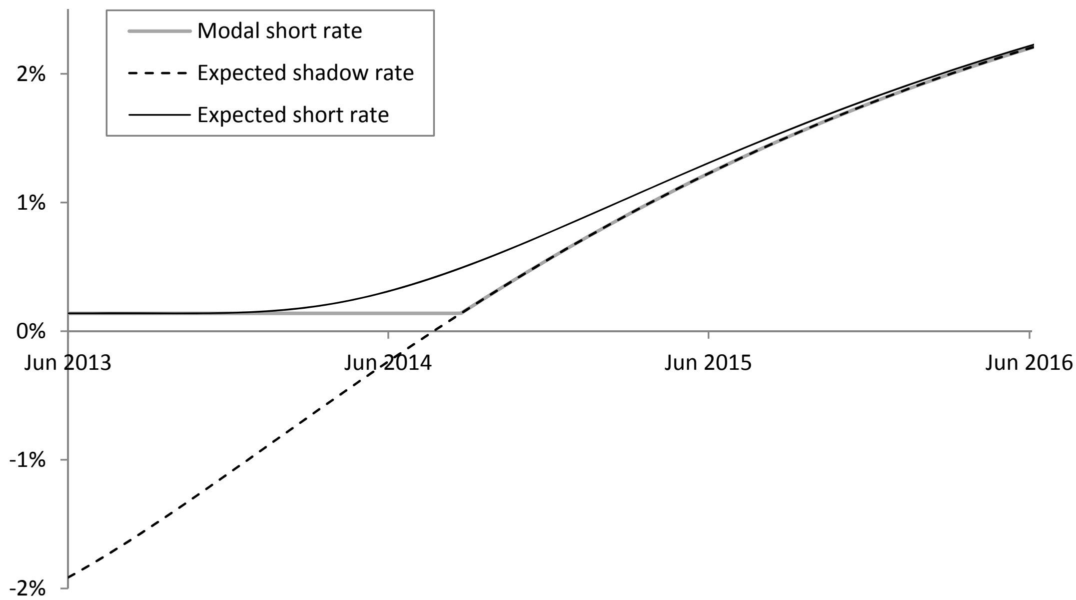

Figure 8: ZLB model-implied short rate as of June 2013. The figure shows the expected path of the shadow rate, as well as the expected and most likely paths of the observed short rate.

the shadow-rate model. The dashed line depicts the expected future shadow rate, $\mathbf{E}_T[s_{T+u}]$. Since the shadow rate is Gaussian, this is also the most likely (or modal) path of the shadow rate. The solid black line shows the expected path of the *observed* future short rate, $\mathbf{E}_T[r_{T+u}]$. Finally, the solid gray line represents the most likely path of the observed short rate. Initially, there is little difference between the expected and the most likely projected paths of the observed short rate. Eventually, as the forecast horizon increases, uncertainty drives a wedge between the two—at the ZLB, all risk is upside risk. Therefore, even though the most likely short rate path remains at the ZLB for more than a year, the expected path slowly increases, and exceeds 25 basis points by June 2014.

5.3.2 Second Moments

While (1) implies that the latent factors X_t are conditionally homoskedastic, and the linear mapping from factors to yields in the Gaussian model preserves this property, the nonlinear relationship between factors and yields in the ZLB shadow-rate model endogenously generates conditional heteroskedasticity in yields. To see this, note that

$$\frac{X_{t+\Delta t} - X_t}{\sqrt{\Delta t}} \to N(0, \Sigma\Sigma^\top)$$

in distribution, as $\Delta t \to 0$, conditional on $\mathcal{F}_t$. Therefore, by an application of the Delta method,

$$\frac{y_{t+\Delta t}^{T+\Delta t} - y_t^T}{\sqrt{\Delta t}} \to N\left(0, \frac{dy_t^T}{dX_t^\top}\Sigma\Sigma^\top \frac{dy_t^T}{dX_t}\right). \tag{10}$$

Since dy_t^T/dX_t is not constant in the shadow-rate model, the instantaneous model-implied yield volatility varies depending on the current state vector X_t. Figure 9 plots the model-implied instantaneous volatility for changes in the two-year yield, as well as an empirical counterpart, the annualized realized daily volatility over the following month, from 2000 onwards. Prior to 2008, the model-implied volatility is constant with the exception of a minor dip in 2003. This is another manifestation of the fact that the ZLB model behaves like a Gaussian model (in which conditional yield volatility is constant) when the ZLB is not binding. On the other hand, from late 2008 onwards, the model-implied yield volatility shows substantial time variation due to proximity to the ZLB, and the realized volatility follows the model-implied pattern closely.

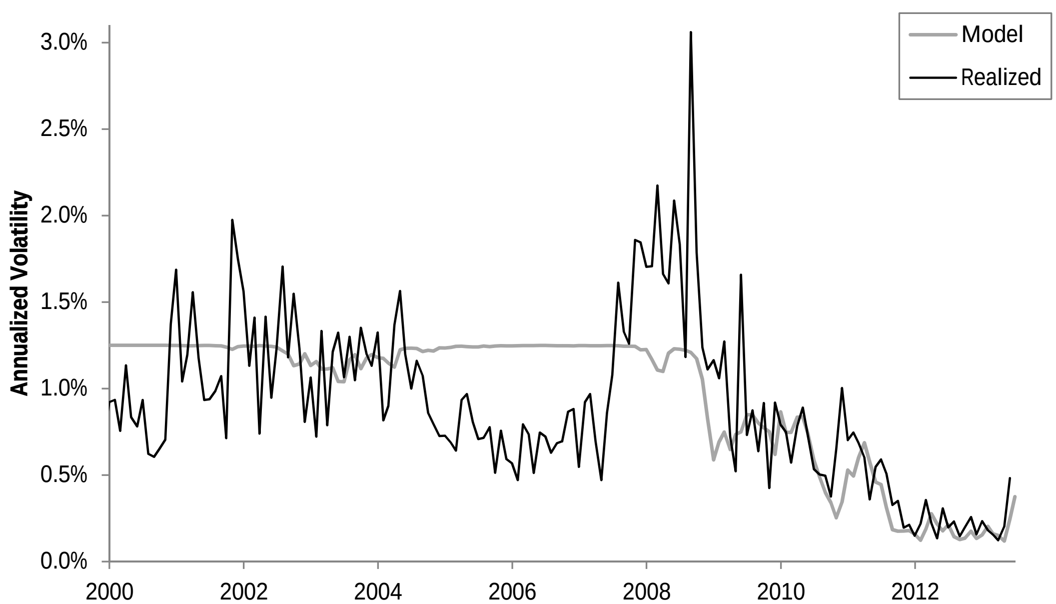

Figure 9: Instantaneous ZLB model-implied volatility for changes in the two-year yield y_t^{t+2}, and annualized realized volatility over the following month, using daily yield data.

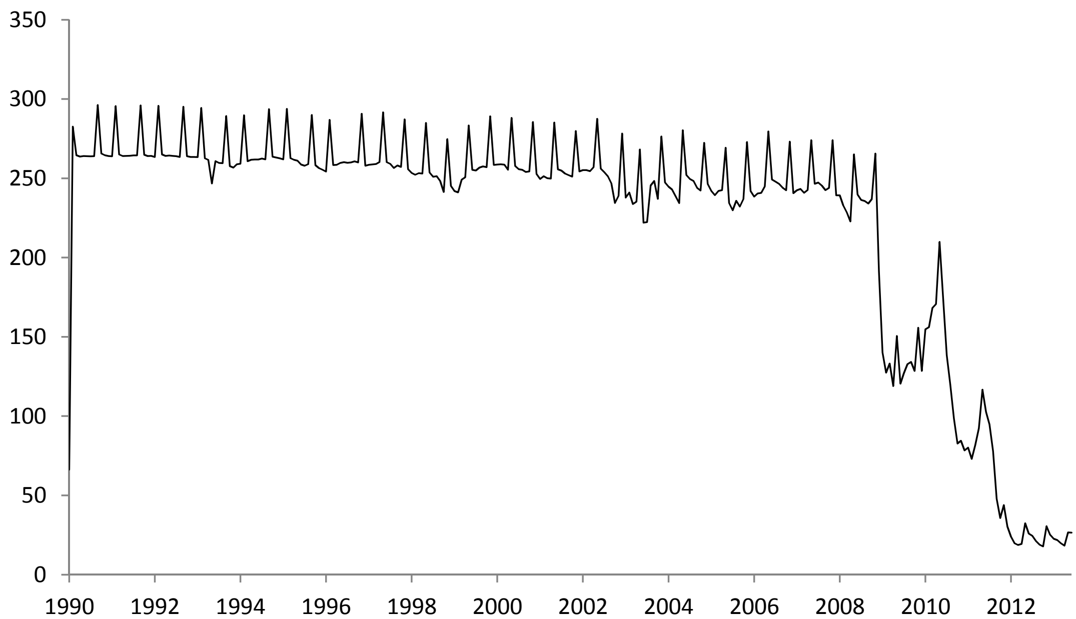

Figure 10: Posterior precision of filtered state variables, as implied by the unscented Kalman Filter.

5.3.3 Precision of Filtered State Variables

In the shadow-rate model, the mapping between states X_t and observables (yields y_t and survey forecasts z_t) is nonlinear. As discussed in Section 4, this leads to time variation in the signal-to-noise ratio of observables, and by implication in the econometrician's ability to infer states from observables.

Figure 10 plots a measure of relative posterior precision of filtered states, defined as

$$\frac{\|\mathrm{Var}^{\mathbb{P}}(X_t)\|_F}{\|\mathrm{Var}^{\mathbb{P}}(X_t|y_1, z_1, \ldots, y_t, z_t)\|_F}.$$

Relative precision will be equal to 1 if observables up to time t contain no relevant information whatsoever, so that the posterior variance of X_t is equal to the unconditional variance. On the other hand, as the posterior variance of X_t goes to zero, relative precision will go to infinity. The figure shows near-constant precision through the end of 2008 (the pronounced seasonality pattern is driven by the availability of long-term surveys). In late-2008, the zero lower bound becomes binding, and yields and survey forecasts become dramatically less informative about the state variables.

5.4 Diagnostics

In Gaussian shadow-rate models, the yields are nonlinear functions of the state variables, but the state variables themselves follow a relatively simple process, namely the multivariate Ornstein-Uhlenbeck process, which, sampled in discrete time, is simply a VAR(1) process (see (7)). Checking whether the filtered state variables are consistent with VAR(1)-dynamics may therefore be a useful diagnostic. An analogue in the literature is Duffie and Singleton (1997), who examine a two-factor CIR model for swap yields. According to the model, the factors are supposed to be independent, but the actual factors implied by the model and the data turn out to be highly correlated, pointing to misspecification. We also made a similar argument in Section 3 above, when we showed in Table 2 the high degree of correlation of yield PCs during the ZLB period.

While in our normalization, factors are conditionally uncorrelated (Σ is diagonal), feedback through $K_1^{\mathbb{P}}$ accommodates a general unconditional correlation structure. Thus, nonzero correlation between model-implied factors is not in itself an indication of misspecification. However, a "collapsing" of factors during the ZLB period, as described in Section 3, might be cause for concern. Table 7 shows the sample correlation of filtered states for models G2013 and the shadow-rate model, both for the entire sample period and for the ZLB sub-sample.

In both models, the correlation structure between states seems to change during

Full Sample			ZLB Period		
1.0000	−0.5062	0.0406	1.0000	−0.9167	0.8467
−0.5062	1.0000	0.6929	−0.9167	1.0000	−0.9682
0.0406	0.6929	1.0000	0.8467	−0.9682	1.0000

(a) Gaussian model G2013.

Full Sample			ZLB Period		
1.0000	0.7925	0.5633	1.0000	0.7519	−0.2262
0.7925	1.0000	0.8994	0.7519	1.0000	−0.2136
0.5633	0.8994	1.0000	−0.2262	−0.2136	1.0000

(b) Shadow-rate model.

Table 7: Sample correlation of filtered states $X_{t|t}$.

the ZLB period, suggestive of a potential structural break, even after accounting for ZLB effects. Furthermore, in the Gaussian model, the states become close to perfectly correlated during the ZLB period, resulting in reduced flexibility in fitting the shape of the yield curve.

For a cleaner reading, note that the standardized *innovation vector*

$$\eta_t = \Omega_{\Delta t}^{-1/2} \varepsilon_t \tag{11}$$

implied by the model and data should be independent. That is, using the estimated model parameters, filtered states, and (7) as well as (11), we can compute a time series of implied innovations that should be i.i.d. $N(0, I_N)$. In particular, if our shadow-rate model is well specified, this should hold even when the sample includes a period of zero interest rates. Therefore, the examination of the innovation vectors η_t's can reveal how structurally stable the specification is.

In general, if η_t are i.i.d. normal with zero mean and identity covariance matrix, it follows from the Central Limit Theorem for random vectors that

$$T\left(\frac{1}{2}\left\|\frac{1}{T}\sum\nolimits_t \eta_t\eta_t^\top - I\right\|_F^2 + \sum\nolimits_{j=1}^J \left\|\frac{1}{T}\sum\nolimits_t \eta_t\eta_{t-j}^\top\right\|_F^2\right) \to \chi^2_{[(2J+1)N^2+N]/2}$$

in distribution as $T \to \infty$, so that this can serve as a test statistic against deviations in scale, location, and correlation structure (between elements of η_t as well as across time). This statistic also has the desirable property that it is invariant to orthogonal transformations applied to η_t (since the Frobenius norm is invariant to orthogonal transformations), and hence to invariant transformations of the underlying state vector.

While our analysis of the properties of the implied innovation vectors is ongoing, preliminary results indicate a departure from independence during the ZLB period both for the Gaussian model G2013 and the shadow-rate model. Hence, this suggests that even after accounting for ZLB effects, the recent episode is potentially characterized by a structural break.

References

Ang, A., and Piazzesi, M. (2003), "A No-Arbitrage Vector Autoregression of Term Structure Dynamics with Macroeconomic and Latent Variables," *Journal of Monetary Economics*, 50:745–787

Bauer, M., and Rudebusch, G. (2013), "Monetary Policy Expectations at the Zero Lower Bound," Working Paper, Federal Reserve Bank of San Francisco

Black, F. (1995), "Interest Rates as Options," *Journal of Finance*, 50(5):1371–1376

Bollerslev, T., and Wooldridge, J. (1992), "Quasi–Maximum Likelihood Estimation and Inference in Dynamic Models with Time-Varying Covariances," *Econometric Reviews*, 11(2):143–172

Chen, R., and Scott, L. (1993), "Maximum Likelihood Estimation for a Multifactor Equilibrium Model of the Term Structure of Interest Rates," *Journal of Fixed Income*, pp. 14–31

Christensen, J., and Rudebusch, G. (2013), "Estimating Shadow-Rate Term Structure Models with Near-Zero Yields," Working Paper, Federal Reserve Bank of San Francisco

Christoffersen, P., Dorion, C., Jacobs, K., and Karoui, L. (2012), "Nonlinear Kalman Filtering in Affine Term Structure Models," CREATES Research Paper 2012-49, Aarhus University

Chung, H., Laforte, J.-P., Reifschneider, D., and Williams, J. C. (2012), "Have We Underestimated the Likelihood and Severity of Zero Lower Bound Events?," *Journal of Money, Credit and Banking*, 44:47–82

Dai, Q., and Singleton, K. (2000), "Specification Analysis of Affine Term Structure Models," *Journal of Finance*, 60(5):1943–1978

Dai, Q., Singleton, K., and Yang, W. (2007), "Regime Shifts in a Dynamic Term Structure Model of U.S. Treasury Bond Yields," *Review of Financial Studies*, 20:1669–1706

Duan, J.-C., and Simonato, J.-G. (1999), "Estimating and Testing Exponential-Affine Term Structure Models by Kalman Filter," *Review of Quantitative Finance and Accounting*, 13:111–135

Duffee, G. (1996), "Idiosyncratic Variation in Treasury Bill Yields," *Journal of Finance*, 51:527–552

Duffie, D., and Kan, R. (1996), "A Yield-Factor Model of Interest Rates," *Mathematical Finance*, 6:369–406

Duffie, D., and Singleton, K. (1997), "An Econometric Model of the Term Structure of Interest Rate Swap Yields," *Journal of Finance*, 52:1287–1321

Fama, E., and Bliss, R. (1987), "The Information in Long-Maturity Forward Rates," *American Economic Review*, 77(4):680–692

Gorovoi, V., and Linetsky, V. (2004), "Black's Model of Interest Rates as Options, Eigenfunction Expansions and Japanese Interest Rates," *Mathematical Finance*, 14(1):49–78

Ichiue, H., and Ueno, Y. (2007), "Equilibrium Interest Rate and the Yield Curve in a Low Interest Environment," Bank of Japan Working Paper

——— **(2013)**, "Estimating Term Premia at the Zero Bound: An Analysis of Japanese, U.S., and U.K. Yields," Bank of Japan Working Paper No. 13–E–8

Joslin, S., Le, A., and Singleton, K. (2013), "Why Gaussian Macro-Finance Term Structure Models Are (Nearly) Unconstrained Factor-VARs," *Journal of Financial Economics*, forthcoming

Joslin, S., Priebsch, M., and Singleton, K. J. (2013), "Risk Premiums in Dynamic Term Structure Models with Unspanned Macro Risks," forthcoming, Journal of Finance

Joslin, S., Singleton, K., and Zhu, H. (2011), "A New Perspective on Gaussian Dynamic Term Structure Models," *Review of Financial Studies*, 24(3):926–970

Julier, S., and Uhlmann, J. (1996), "A General Method for Approximating Nonlinear Transformations of Probability Distributions," Technical Report, RRG, Department of Engineering Science, University of Oxford

Julier, S., Uhlmann, J., and Durrant-Whyte, H. (1995), "A New Approach for Filtering Nonlinear Systems," in *Proceedings of the American Control Conference*, pp. 1628–1632

Karatzas, I., and Shreve, S. (1991), *Brownian Motion and Stochastic Calculus*, Graduate Texts in Mathematics Series, Springer, London

Kim, D., and Singleton, K. (2012), "Term Structure Models and the Zero Bound: An Empirical Investigation of Japanese Yields," *Journal of Econometrics*, 170(1):32–49

Kim, D., and Wright, J. (2005), "An Arbitrage-Free Three-Factor Term Structure Model and the Recent Behavior of Long-Term Yields and Distant-Horizon Forward Rates," Finance and Economics Discussion Series 2005–33, Federal Reserve Board, Washington, D.C.

Kim, D. H., and Orphanides, A. (2012), "Term Structure Estimation with Survey Data on Interest Rate Forecasts," *Journal of Financial and Quantitative Analysis*, 47(1):241–272

Krippner, L. (2012), "Modifying Gaussian Term Structure Models When Interest Rates are Near the Zero Lower Bound," Reserve Bank of New Zealand Discussion Paper 2012/02

Li, C., and Wei, M. (2013), "Term Structure Modeling with Supply Factors and the Federal Reserve's Large-Scale Asset Purchase Programs," *International Journal of Central Banking*, 9(1):3–39

Litterman, R., and Scheinkman, J. (1991), "Common Factors Affecting Bond Returns," *Journal of Fixed Income*, 1:54–61

Lund, J. (1997), "Non-Linear Kalman Filtering Techniques for Term Structure Models," Working Paper, Aarhus School of Business

Priebsch, M. (2013), "Computing Arbitrage-Free Yields in Multi-Factor Gaussian Shadow-Rate Term Structure Models," FEDS Working Paper 2013–63, Federal Reserve Board, Washington, D.C.

Wan, E., and van der Merwe, R. (2001), "The Unscented Kalman Filter," in *Kalman Filtering and Neural Networks*, ed. by S. Haykin, pp. 221–280

Wu, S. (2010), "Non-Linear Filtering in the Estimation of a Term Structure Model of Interest Rates," *WSEAS Transactions on Systems*, 9(7):724–733

Xia, J. C. W. F. D. (2013), "Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound," Working Paper, Chicago Booth and UC San Diego