TERM STRUCTURE FORECASTING AND SCENARIO GENERATION

Robert Engle and Emil Siriwardane[1]

Prepared for a Special Issue of the Journal of Econometrics in Honor of Christian Gourieroux.

7/6/2015

ABSTRACT

A statistical model is developed to generate probability based scenarios for forecasting and risk management of long term risks. The model forecasts transformed daily forward interest rates over a 10 year horizon. The model is a reduced rank vector autoregression with time varying volatilities and correlations. A quasi-differencing version reduces the impact of autocorrelated measurement errors.

Risk managers with a large book of fixed income assets are typically challenged to estimate the risk over both short and long run horizons. Familiar measures of market risk such as Value at Risk, VaR, are based on factor exposures to extreme events and will be time varying. This literature has been nicely surveyed by Gourieroux and Jasiak(2010) who also focus on longer term credit risks. Following this discussion, longer run risks involve estimates of what the term structure might look like in a year or in ten years. Most term structure models are not adequate to this task, so scenario analysis is used. However, designing scenarios is an art form which is particularly complicated in the fixed income asset class as the scenario should be internally consistent and arbitrage free and must have a probability assessment that motivates the response to risk. It furthermore will be useless if it does not stress the assets that are held in the firms portfolio.

An alternative way to generate scenarios is to use probability based scenarios as described by Christensen, Lopez and Rudebusch(2014). The strategy is to construct a large number of equally likely scenarios which are drawn from the joint predictive distribution of the term structure. The value of the portfolio of assets can be calculated from each scenario over time and a Profit and Loss distribution computed at different horizons. Because each scenario has a known probability of occurring, the risks at different horizons in the future can be assessed. Their application is to the management of the FED's enormous fixed income portfolio, but it could similarly be applied jointly to any financial firm's assets and liabilities.

This is essentially the proposal embodied in the Solvency II Directive, a new regulatory framework for the European insurance industry that adopts a more dynamic risk-based approach and

[1] The authors appreciate financial and intellectual support from A.I.G in developing this model. Particular thanks go to Sid Dalal, David Li and Kevin Chen. We also would like to thank Glenn Rudebusch, Jim Hamilton, Michael Bauer and Cynthia Wu for useful conversations on this research.

implements a non-zero failure regime. The directive mandates market valuation of assets and liabilities and a maximum 1 in 200 probability of failure over a long horizon. Both assets and liabilities incur risk and all risks may be correlated.

This paper proposes an econometric approach to generating scenarios for the US treasury term structure. These scenarios can be interpreted as a predictive distribution for the term structure in the near or distant future. These predictions should be arbitrage free and should have sufficient range that the scenarios stress all realistic outcomes.

2. Literature Survey

There is a vast literature on the term structure on non-defaultable securities. It is the main subject of countless books, courses and careers. The purpose of most of this literature is however different from the purpose of this paper. Most of the literature is concerned with estimating the fair market prices of securities that don't have an observable price in terms of others that do. This is extremely important as fixed income securities and their derivatives trade only occasionally but must be evaluated based on other securities that do trade. Thus the goal is estimating the yield curve at a moment of time based on other prices recorded at that moment. Part of this analysis requires estimating the price at which a bond would trade and how this is decomposed into term and risk premiums. This analysis formally constrains the yield curve to be arbitrage free. Early models of this form treat the short rate as the state variable and by postulating a dynamic relation, derive the entire term structure. These include the Vasicek(1977), Cox, Ingersoll and Ross(1985), Ho and Lee(1986), and Hull and White(1990) among many others. Today short rate models are particularly problematic with the extended period of nearly zero short rates. A popular multi-factor specification is the affine structure which allows closed form expressions for these objects. See for example Duffie and Kan(1996) and a recent survey by Piazzesi(2010).

Two features of these models are generally in conflict. To avoid arbitrage opportunities, nominal interest rates must be non-negative, at least if cash can be stored costlessly. However, to ensure that there are no arbitrage opportunities, it is sufficient to construct a risk neutral measure that prices all assets. Bjork(2009) shows that the affine family is the only tractable family that satisfies both conditions. Yet it is a very restricted specification which is claimed to be unable to model important features of the data such as the very low short interest rates we see today. See Duffee(2002) and Christensen and Rudebusch(2013) for example of these criticisms.

To compute term premiums, it is necessary to compute the expectation of the future instantaneous short rate, r(s)

$$term\ premium = P_{t,T} - E\left[e^{-\int_t^T r(s)ds} \right]. \tag{1}$$

The difference between this expectation and the current bond price is interpreted as a term or risk premium. The standard approach to estimating bond prices is to specify a risk neutral measure, Q which satisfies

$$P_{t,T} = E_t^Q \left[e^{-\int_t^T r(s)ds} \right] \tag{2}$$

The estimation of a term structure involves estimating both an empirical measure often called P as in (1), and a risk neutral measure, Q. The risk premium embedded in some security such as a forward rate, is the difference between the risk neutral and empirical specification.

$$Risk\ \text{Pr}\, emium = E^Q\left(f_{t=h}^i\right) - E^P\left(f_{t=h}^i\right) \tag{3}$$

This can be done tractably with positive rates, only for members of the affine class. Yet affine processes are widely recognized as being restricted representations of the data which require some adjustment for the zero lower bound on rates. A recent extension of the affine models by Monfort, Pegararo, Renne and Roussellet(2014) is a promising approach.

Forecasting applications such as Diebold and Rudebusch(2013), and Diebold and Li(2006) and risk management applications like Christensen, Lopez and Rudebusch(2013) and Muller, Dacorogna and Blum(2010) focus on good statistical representations of the data without requiring that the term premia be computable. The term structures from these methods are approximately arbitrage free but do not allow closed form solutions for the pricing kernel. The forecasting applications are focused on the mean of the future distribution of yields, while the risk management applications are more focused on the range and probability of extreme outcomes.

We follow this route. We develop an econometric model that is suitable for forecasting and scenario generation in the risk management context. Like the Heath Jarrow and Morton(1992) model, we focus on forecasting models for the forward yields that restrict them to be non-negative and mean reverting. The model implies term premia but these cannot be calculated in closed form.

3. Econometric Specification

Yields of zero coupon bonds are carefully estimated by Gurkaynak, Sack and Wright(2006) using a generalization of the Nelson Siegel term structure model due to Svenson to infer zero rates from detailed historical quotes from off the run treasuries. These yield curves are constructed daily and updated weekly for a range of maturities greater than a year. We supplement their data with a six month series from the FED's form H-15. We focus on the widely reported maturities that are most actively traded although the model can easily be extended to cover more maturities. We analyze nine maturities from 1993 to May 2015 which is roughly 5000 observations per series. These are six months, one year, two years, three years, five years, 7 years, 10 years, 20 years and 30 years[i]. We denote these zero coupon yields as r_t^i , the yield to maturity observed on day t for maturity τ_i at the set of maturities

$$\tau = \{\tau_1, ..., \tau_9\} = \{0.5, 1, 2, 3, 5, 7, 10, 20, 30\}. \tag{4}$$

The data are transformed into forward rates. These rates are the natural forwards with maturities given by the difference between successive yields. Thus there is a six month, forward which starts in six months and a one year forward starting in one year and one starting in two years. These will be labeled for the ending date. Thus f_t^i is the forward rate observed on day t for a contract starting at $t + \tau_{i-1}$ and ending at the next higher yield maturity, $t + \tau_i$.

Forward rates are computed assuming continuous compounding. These match the cash flows from investing in a zero coupon bond of maturity τ_{i-1} and then putting the proceeds into a forward of maturity τ_i, τ_{i-1}, with the cash flow from investing in a single zero coupon bond of maturity τ_i.

$$f_t^i = \frac{r_t^i \tau_i - r_t^{i-1} \tau_{i-1}}{\tau_i - \tau_{i-1}} \tag{5}$$

The same relation is easily inverted to compute the yields as a function of predicted forwards.

An arbitrage free term structure model must have a zero probability that a forward rate is negative. [ii] To ensure this condition, the forward rates are transformed before modeling. The common transformation is the log transform which is simply to model the natural log of the forward rates. Thus

$$y_t^i = \log\left(f_t^i\right) \tag{6}$$

An alternative is to use a linear transform for higher rates and a log transform for low rates, constraining the transform so that it is continuously differentiable at the switch point. This transformation has the advantage that it only affects very low rates and does not magnify the volatility of high rates by exponentiating the predicted values. Thus if the switch is given by $\overline{f}$, then

$$y_t^i = \begin{cases} f_t^i & \text{if } \; f_t^i > \overline{f} \\ a + b\log\left(f_t^i\right) & \text{if } \; f_t^i \le \overline{f} \end{cases} \tag{7}$$

and

$$a = \overline{f}\left(1 - \log\left(\overline{f}\right)\right), \; b = \overline{f} \tag{8}$$

which will make the function continuously differentiable and strictly monotonic for all positive f. The function is easily inverted so that simulations of y can be transformed into simulations of f. Because these transformations are non-linear, expectations will not transform directly but quantiles will. Closed form solutions for term premiums will not be available but confidence intervals and medians will be.

This transformation serves the same function as the popular shadow rate model. The process that is modeled, in this case y, can go negative but the effective forward rate cannot. Future rates can go very close to zero and because the process is very persistent, they will remain there for substantial time as y becomes quite negative. However, the process is ultimately mean reverting so rates will adjust to historical averages.

In the empirical work, a switching point for the transformation must be assumed. With very little experimentation, we will use 1% so that all rates above 1% are used linearly and rates below 1% are transformed into logs. In this case we simply have a=b=1 in (8).

The transformed forward rates are modeled as a reduced rank vector autoregression with errors that are allowed to have asymmetric volatility processes and time varying correlations. The vector autoregression has reduced rank to reflect the factor structure. The model is specified as follows:

$$y_t - \mu = A\left(y_{t-1} - \mu\right) + H_t^{1/2}\varepsilon_t, \quad A = \phi\theta', \; H_t = D_t R_t D_t, \; \varepsilon_i \sim i.i.d.\; N\left(0, I\right) \tag{9}$$

The matrices $\left(\phi, \theta\right)$ are 9x3 matrices that reflect the factor loadings of the three factors and ensure the factor structure of the model. The theta matrix just reflects the Nelson-Siegel regression. The three factors are interpreted as level, slope and curvature in the NS framework and are estimated by a cross sectional regression which has the same regressors every day. To be precise, consider the model

$$y_t^i = y_t^{\tau_i} = L_t + S_t\left(\exp\left(-\lambda\tau_i\right)\right) + C_t\left(\lambda\tau_i \exp\left(-\lambda\tau_i\right)\right) + \varepsilon_t^\tau \tag{10}$$

Consequently, each factor is just a time invariant linear combination of the y's which we denote by theta. We use the value $\lambda = 0.0609$ following Diebold and Li(2006) and others. An alternative to the NS factor structure is to use principle components. The performance is very similar and is reported in some cases below. The matrix phi is estimated in equation (9). The matrix D is diagonal with asymmetric GARCH volatilities on the diagonal, each specified to be GJR-GARCH. The correlation matrix is specified to be a DCC correlation process with just two parameters determining the 9x9 correlation matrix which gets updated every day.

The DCC specification makes it easy to estimate this model equation by equation. Defining the factors as $F_t = \theta' y_t$, there are nine equations that look as follows:

$$\begin{aligned} y_t^i &= \mu_i + \phi_i F_{t-1} + e_{i,t} \qquad e_{i,t} = \sqrt{h_t^i}\,\varepsilon_{i,t}, \\ h_t^i &= \omega_i + \alpha_i e_{i,t-1}^2 + \gamma_i e_{i,t-1}^2 I_{e_{i,t-1<0}} + \beta_i h_{t-1}^i \end{aligned} \tag{11}$$

These are simply univariate GARCH models with lagged factors as regressors. The standardized residuals, $\left\{\varepsilon_{i,t}\right\}$ are then fitted with a standard DCC to estimate the conditional correlations.

Once the model has been estimated, it is straight forward to simulate it. The shocks can be drawn from the historical distribution or simulated with a standardized distribution such as the Gaussian. If they are drawn from the historical distribution, then it may be natural to allow them to be generated by a non-parametric copula as in Brownlees and Engle(2010). The results presented below use normal random variables for the simulations. For some sample periods, the VAR in (9), has a maximum root that is greater than or equal to one. This would make the entire process non-stationary and lead to confidence intervals that grow with the horizon of the simulation. Consequently the A matrix will be slightly adjusted as follows

$$A^* = A - \xi I \tag{12}$$

The value of ξ is chosen so that the maximal eigenvalue of A* is equal to .999 or else is just its estimated value. In the implementation of equation (11), the mean of y is computed initially and the entire analysis is done on demeaned data. The means are then added back after the simulation. This insures that the shrinkage in (12) does not also affect the mean of the forecasts, only the rate of decay.

4. Results

The model described in equation (11) is estimated using daily data from Gurkaynak, Sack and Wright(2006) from October 1993 through March 2015. Forward rates are constructed from their zero coupon bond yields. The forwards are then transformed with either log or log-linear transformations. The factors are computed from a cross sectional Nelson Siegel regression and then are used as lagged explanatory variables in the main equation. This is estimated allowing for asymmetric heteroskedasticity and cross equation dynamic conditional correlation. A Gaussian quasi likelihood function is used in each case.

The parameter estimates for the factor loadings in the log forward reduced rank VAR are given in Table 1A below. The GJR GARCH parameters are presented in table 1B and the DCC parameters are given in Table 1C.

Table 1. about here

These parameter estimates are generally highly significant. The factor loadings are almost all very significantly different from zero. The GARCH parameters are somewhat different from those expected in equity markets. The asymmetric term is often not significant and the alpha is larger usual. The persistence is generally quite high and when combined with these larger alphas, implies a high volatility of volatility. Correlations are estimated to be changing dynamically but with more volatility than is typical of equity market data.

An important feature of this model – and many popular models in the literature, is that the residuals have substantial autocorrelation. The first lag of the residuals from these 9 equations for the forward maturities is given in Table 2.

Table 2 about here

Similar results are found for the log linear transformation. The parameters are apparently well estimated but the residual autocorrelation is very substantial.

A natural solution to this problem is to add lagged factors or lagged dependent variables. Unfortunately, adding lagged factors makes little difference and adding lagged dependent variables reduced the factor structure as the lagged dependent variable becomes by far the most significant variable.

In order to develop a better model, it is essential to determine the source of the autocorrelation.

5. Measurement Errors

A more conventional approach to forecasting the yield curve is to forecast the factors and then compute the yields assuming no errors. Thus the model is specified as

$$\begin{aligned} y_t &= \mu + \delta' F_t \\ F_t &= A F_{t-1} + H_t^{1/2} \varepsilon_t \end{aligned} \tag{13}$$

The first equation in (13) models the relation between the contemporaneous value of the factors and the forwards. Whether the factors are generated by a Nelson-Siegel procedure, by principle components or by a specific economic assumption such as the short rate, this equation maps factor forecasts onto all the yields. The yields are assumed to be exact linear combinations of the factors. In this case the forecast errors observed in equation (11), are simply the errors in forecasting the factors in (13).

However, there may be errors in the first equation of (13) which we shall call measurement errors or in the language of equity analysis, idiosyncratic errors. Then the residuals from (11) are a combination of measurement errors and factor forecast errors. We define the vector u as the measurement error in

$$y_t - \delta F_t - \mu = u_t \tag{14}$$

A regression of transformed forwards on the factors has residuals that are estimates of the measurement errors. These measurement errors are generally assumed to be small and are ignored in the simulation and forecasting of the model. See for example Diebold and Li(2006) and Christensen Lopez, and Rudebusch(2012).

With our data, the R^2 of the time series regressions on the first equation, whether estimated with NS or PC, is typically greater than 95% and often above 99%. However, the errors are still significant. These errors affect not only the yields on date t but also affect the factors in t+1 and later, which would affect the yields in the future. Furthermore, the measurement errors have persistence that further influences multi-period forecasting.

To assess the importance of the measurement errors, we now compare the specifications in (11) and (13). The forecast errors in the reduced rank model are composed of both forecasting errors for the factors and measurement errors. Thus it is interesting to compare the variance of the measurement errors with the variance of the reduced rank forecast errors. In Table 3 below, we examine these ratios by comparing the square of the standard error of the measurement errors with the average of the GARCH variance of the reduced rank equation. The results are similar but even more striking when the equation in (11) is estimated with least squares.

Table 3 about here

Not only are the measurement errors big relative to the total forecast errors, but they are also highly persistent. The first order autocorrelation from the 3PC model is given in Table 5 below. The column AR(1) is estimated jointly with the factor loadings using 3 Principle Components. The AUTOCOR(1)is simply the first order autocorrelation of the residuals from the least squares fit. Both estimates clearly show how persistent these measurement errors really are.

Table 4 about here

6. Quasi-Differenced Factor Models

A solution to this modeling problem is to quasi difference the transformed forward rates before projecting them onto the Factors. This simple procedure reduces the serial correlation dramatically and yet allows the factors to shape the yield curve and its dynamics.

Let u be the measurement error and ν be the factor forecast error. Then (9) becomes

$$y_t = \phi F_{t-1} + \mu + u_t + \phi \nu_t \tag{15}$$

Assume that u is autocorrelated but that ν is not. Then

$$u_t = \rho u_{t-1} + \eta_t \tag{16}$$

and consequently, from (14) (15) and (16),

$$\begin{aligned} y_t &= \phi F_{t-1} + \mu + \rho u_{t-1} + \eta_t + \delta \nu_t \\ y_t &= \phi F_{t-1} + \mu + \rho\left(y_{t-1} - \delta F_{t-1} - \mu\right) + \eta_t + \delta \nu_t \\ y_t - \rho y_{t-1} &= \mu(1-\rho) + (\phi - \rho\delta) F_{t-1} + \eta_t + \delta \nu_t \end{aligned} \tag{17}$$

Thus by simply quasi-differencing the transformed forwards, we can eliminate the impact of the measurement errors. The estimated factor loadings will have a somewhat different interpretation; however the model can be estimated and simulated as before.

Thus we estimate the following model as the quasi-difference model

$$y_t - \mu - \rho\left(y_{t-1} - \mu\right) = \phi^* F_{t-1} + H_t^{1/2}\varepsilon_t \tag{18}$$

where H is a GARCH-DCC process and $\theta^* = \phi - \rho\delta$. The key feature of this model is that ρ is the same for each equation and is specified in advance. The quasi differencing of each forward reduces the importance of the measurement errors but retains the impact of the factors. This model is completely easy to estimate and to simulate. Remembering that the factors are defined by $F_t = \theta' y_t$, it is clear that the evolution of the factors is governed by

$$F_t = \left(\rho I + \theta'\phi^*\right) F_{t-1} + \theta' H_t^{1/2}\varepsilon_t \tag{19}$$

Thus the factors will be stationary if the maximum eigenvalue of $\left(\rho I + \theta'\phi^*\right) = \left(\rho I + \theta'\phi - \rho\theta'\delta\right)$ is less than one and this is sufficient for all the yields to be stationary.

In the implementation here, rho=.9. This value could easily be improved but it is a compromise between whitening the residuals and forecasting the factors. It is clear that if rho=1 and delta is very similar to phi, then the effect of the factors is eliminated and the pricing model will just be a set of independent random walks. For the same data set and the log linear transformation, the estimated parameters are given in table 5. The three parts again are the factor loadings in 5A, the GARCH parameters in 5B and the DCC in 5C.

Table 5 about here

The same model has been estimated for the log but the parameter estimates will not be reported. These models have much better residual properties than the base models without quasi-differencing. This can be seen by comparing the autocorrelations and squared autocorrelations in Table 6. Although some of these correlations are significant statistically because of the large sample size, they are economically small.

Table 6 about here

7. Simulated Forwards

For each forward rate, 1000 simulations are carried out for a 10 year horizon. The simulation is done with the estimated reduced rank VAR with DCC-GARCH errors in equation (11). The data are demeaned so that each series has mean zero. The process is as follows:

1) Gaussian shocks for day 1 of the simulation are drawn for each maturity.

2) To match the correlation structure, these shocks are multiplied by the Cholesky decomposition of the predicted correlation matrix.

3) The shocks are multiplied by the conditional standard deviation

4) The errors are then added to the predicted mean based on the factors and factor loadings to give the day 1 values of the transformed forwards.

5) These values are used to update the factors, the volatilities and the correlations.

6) The cycle continues with step 1 again.

Then

a) The mean is added back to get predicted transformed forwards.

b) The inverse transform is invoked to obtain scenarios of forwards

c) The forwards are combined to get yields

d) The 5%, 50% and 95% quantiles of the yields and forwards are recorded

Figure 1, shows the simulation of the 6month forward rate from the quasi-differenced log-linear forward reduced rank VAR model. Over the next 10 years this model expects the short rate to rise to a median value of about 2.2% by 2024. However the confidence interval between the 5% and 95% probabilities is quite wide. There remains a sizeable probability of very low short rates over the next ten years and the upper level stabilizes at about 5.2%. After 10 years it is unlikely that the distribution of short rates will depend upon anything except the historical data. In fact the historical 5% and 95% quantiles are (.06, 6.08).

In Figure 2, the forward rate from 7 to 10 years, is simulated for the next ten years based on the quasi-differenced, log linear transformation and the RRVAR. The rate today is under 3% but the steady state is about 5.5% with a confidence interval of approximately (3%,7%). In the data, this interval was (3.4%, 8.1%). Again, the steady state appears to be reached by 2018. This however is not true of all maturities or models. For this model, the short rate does not reach its steady state until 2025.

Combining the forwards together to generate yields allows us to look at a confidence interval for the whole yield curve at a given horizon. In figures 3-6 the ten year yield curve forecasts are presented for the four models considered.   These are the Quasi-Differenced Log-Linear Transform of forwards with RRVAR, and the Quasi-Differenced Log forward model as well as the base Log Linear Transform and the base log transform. All four of these plots include the unconditional quantiles of the data. At the 10 year horizon, one might expect that the model quantiles should match the data quantiles. If the model quantiles are too tight, then it is likely that the risk management role will be to underestimate risks. However if the intervals are too wide, then it may be that the risks will be overestimated.

From these four plots it appears that the model intervals are generally smaller than the data but this downward bias is smallest for the Quasi-Differenced models and for the log-linear transformation. Thus from this point of view the best model is the Log Linear Transformed Quasi-Differenced model.

8. Model Validation

Model validation can be done be examining the fit of individual pieces of the model or by examining the historical performance through a backtest. Ideally both would be employed in choosing the superior model. As discussed above, we have done many diagnostic tests of the residuals from the model and find that they are not as well behaved as they should be if the model is correctly specified. There is apparently a moderate amount of serial correlation at a daily level in the yield curve, conditional on the yield factors.

Because we have estimated the model recursively, we can examine whether future confidence intervals formed at a time in the past are accurate representations of the uncertainty in the model forecasts. For example, 10 years ago, what would be the confidence interval for the short rate? Would it include zero or a value close to zero. Five years ago this would still be a challenge.

As there are many forecast dates and many horizons, it may be useful to have a systematic approach to backtesting this model. Let the forecast of quantile q of yield maturity i made at date t for date h be denoted $r_{h/t}^{i,q}$. Then define two indicators of whether the realization is contained within a specific confidence interval.

$$
\begin{aligned}
I_{h/t}^{i} &= \begin{Bmatrix} 1 \; if \;\; r_h^i \in \left(r_{h/t}^{i,.95}, r_{h/t}^{i,.05} \right) \\ 0 \qquad otherwise \end{Bmatrix} \\
B_{h,t}^{i} &= \begin{Bmatrix} 1 \; if \;\; r_h^i \le \left(r_{h/t}^{i,.50} \right) \\ -1 \qquad otherwise \end{Bmatrix}
\end{aligned} \tag{20}
$$

If forecasts are initiated at T dates and evaluated at future values of these same dates, there will be NT(T-1)/2 values of these indicators. For a perfect model, the mean of indicator, I should average .9. If the mean is above .9 then the confidence intervals are too wide and if it is below .9, they are too small. Examination of the combinations of t,h and i would indicate where the model is deficient. If the mean of I is less than .9 when t-h is large, then the model is too confident at long horizon forecasts and if it is less than .9 for short maturities, then it is too confident at short maturities.

The mean of B will indicate whether the data are above the median or below. If the mean of B is positive, it means that the predicted median is greater than the realized data and conversely.

The values of I and B will not be independent and therefore the sampling properties are not simple to assess. For a particular i and a particular t-h there will be non-overlapping forecasts that could be considered independent. In most cases, there will be only a very small number of independent forecasts. However, for different maturities and overlapping forecasts, the dependence structure is complex to derive and estimate. Consequently we will use the simple mean of I and B to assess the backtest performance of competing models.

One important source of bias for these models is the failure to account for other features of the macroeconomy. The fact that interest rates are generally lower than they were forecast to be 10 years ago, is due to the slowing of the US economy and the financial crisis.   This information is not in the term structure model and its inclusion would improve forecast performance. For this reason, the main focus of the backtests is on the width of the confidence intervals rather than the level. In future versions of the model, macro variables will be incorporated as in Ang and Piazzessi among others.

TABLES

Table 1A.

Coefficients and T-Statistics for Reduced Rank VAR on Log Forwards

Factors constructed from Nelson Siegel on

Daily data from 1993 to 2015 May

Factor Loadings

Maturity	F1	F2	F3
6m	1.0111	0.7141	0.2662
	(1305.50)	(1460.20)	(796.50)
1yr	0.9766	0.4368	0.3309
	(700.50)	(483.30)	(535.40)
2yr	1.0045	0.2477	0.323
	(964.20)	(372.90)	(744.70)
3yr	1.0232	0.1222	0.2545
	(875.30)	(157.30)	(462.00)
5yr	0.9796	0.0355	0.1483
	(772.40)	(45.20)	(272.20)
7yr	0.9762	0.0176	0.0397
	(686.50)	(22.20)	(67.60)
10yr	0.9189	-0.0266	0.0084
	(602.60)	-(29.20)	(14.60)
20yr	1.0055	0.045	-0.0459
	(685.20)	(52.00)	-(89.30)
30yr	0.7802	-0.0424	-0.0075
	(489.80)	-(49.20)	-(11.90)

Table 1B

GARCH Parameters

Maturity	Omega	Alpha	Gamma	Beta
6m	0.0002	0.1943	0.0153	0.7927
	(14.06)	(20.32)	(3.10)	(91.96)
1yr	0.0005	0.2341	0.0061	0.7445
	(17.62)	(26.66)	(1.40)	(97.27)
2yr	0.0001	0.0645	0.0066	0.9189
	(12.89)	(17.77)	(3.54)	(235.31)
3yr	0.0001	0.0859	-0.0038	0.907
	(16.29)	(24.97)	-(2.75)	(266.72)
5yr	0.0003	0.1703	-0.0044	0.8132
	(28.35)	(43.54)	-(2.26)	(225.43)
7yr	0.0003	0.2149	0.0027	0.7759
	(28.96)	(56.20)	(1.26)	(258.84)
10yr	0.0006	0.2871	0.0086	0.6908
	(27.41)	(51.26)	(2.53)	(140.25)
20yr	0.001	0.3703	0.0143	0.5965
	(22.36)	(35.52)	(2.54)	(68.67)
30yr	0.0009	0.4364	0.008	0.5596
	(16.42)	(33.73)	(1.31)	(53.10)

Table 1C

DCC Parameters

Alpha	0.1378
	(91.68)
Beta	0.8563
	(557.32)

Table 2

First Order Residual Autocorrelation of log forward RR VAR

Mat	Auto
6M	0.5198
1YR	0.6080
2YR	0.2877
3YR	0.3745
5YR	0.5911
7YR	0.7207
10YR	0.6893
20YR	0.7264
30YR	0.8227

Table 3

Measurement Error Variance Relative to Total Forecast Error Variance[iii]

NAME	LS_PC_LOG	PC_LOG	PC_TRANS	NS_LOG	NS_TRANS
6m	0.995	0.635	0.914	0.373	0.567
1yr	0.941	0.789	0.682	0.659	0.768
2yr	0.897	0.867	0.705	0.445	0.407
3yr	0.952	0.914	0.908	0.567	0.540
5yr	0.933	0.849	0.859	0.756	0.763
7yr	0.570	0.492	0.545	0.812	0.898
10yr	0.848	0.792	0.713	0.780	0.893
20yr	0.938	0.623	0.694	0.872	0.891
30yr	0.758	0.637	0.535	0.850	0.745

Table 4

Autocorrelation of measurement Errors

NAME	AR(1)	AUTOCOR(1)
6m	0.9973	0.9819
1yr	0.9694	0.9684
2yr	0.9973	0.9862
3yr	0.9990	0.9965
5yr	0.9991	0.9973
7yr	0.9947	0.9809
10yr	0.9990	0.9970
20yr	0.9982	0.9916
30yr	0.9984	0.9926

Table 5A

Parameter Estimates for Quasi-differenced Log-Linearly Transformed Forwards

Maturity	F1	F2	F3
6m	0.102	0.070	0.027
	(220.07)	(224.83)	(140.91)
1yr	0.097	0.045	0.032
	(109.00)	(83.88)	(92.92)
2yr	0.100	0.025	0.032
	(130.92)	(51.66)	(100.99)
3yr	0.101	0.012	0.026
	(124.88)	(23.33)	(75.17)
5yr	0.102	0.003	0.016
	(125.29)	(6.59)	(47.47)
7yr	0.100	0.001	0.005
	(121.69)	(2.87)	(14.62)
10yr	0.098	0.002	-0.002
	(116.99)	(3.77)	-(5.28)
20yr	0.099	0.004	-0.005
	(141.97)	(8.71)	-(17.41)
30yr	0.085	-0.002	-0.002
	(89.60)	-(3.31)	-(5.34)

Table 5B

GARCH Estimates for Quasi Differenced Log Linearly Transformed Model

Maturity	Omega	Alpha	Gamma	Beta
6m	0.0000472	0.0728	0.1173	0.8685
	(14.17)	(6.98)	(9.14)	(149.23)
1yr	0.0000534	0.0399	0.0175	0.9441
	(7.47)	(8.01)	(3.65)	(187.85)
2yr	0.000055	0.0292	0.0197	0.9518
	(11.25)	(9.49)	(5.79)	(300.98)
3yr	0.0000461	0.0405	-0.0001	0.9519
	(11.99)	(13.32)	-(.05)	(315.52)
5yr	0.0000354	0.0396	0.0015	0.9536
	(9.72)	(12.79)	(1.50)	(295.74)
7yr	0.0000242	0.04	0.0028	0.9546
	(6.76)	(13.83)	(2.64)	(299.18)
10yr	0.0000397	0.0443	0.0007	0.9474
	(5.82)	(14.73)	(.62)	(244.37)
20yr	0.0000359	0.0365	0.0081	0.9494
	(5.49)	(10.64)	(4.03)	(211.84)
30yr	0.0001654	0.091	-0.022	0.905
	(9.53)	(6.44)	-(2.94)	(79.98)

Table 5C

DCC Parameter Estimates

Alpha	0.0386
	(12.65)
Beta	0.9592
	(328.97)

Table 6

Autocorrelation of Residuals in Quasi-Difference Log Lin Model

MATURITY	AUTO1	AUTOSQ1
6_m	0.012	0.004
1_yr	-0.043	-0.013
2_yr	0.037	0.001
3_yr	0.034	-0.008
5_yr	0.056	-0.016
7_yr	0.108	0.022
10_yr	0.132	0.030
20_yr	0.102	-0.015
30_yr	0.157	-0.002

FIGURES

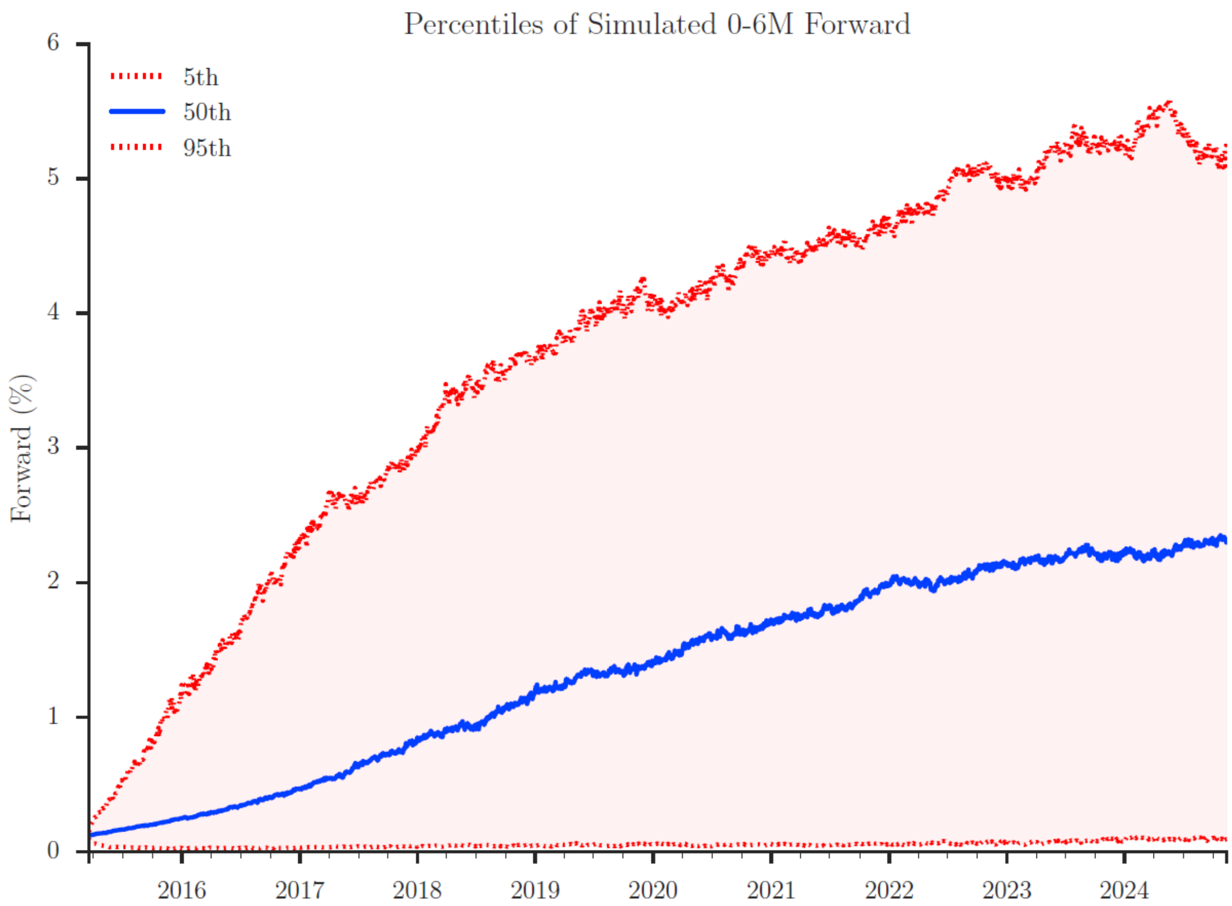

Figure 1

Simulated Short Rate from quasi-differenced log-linear transformed forward in Reduced Rank VAR

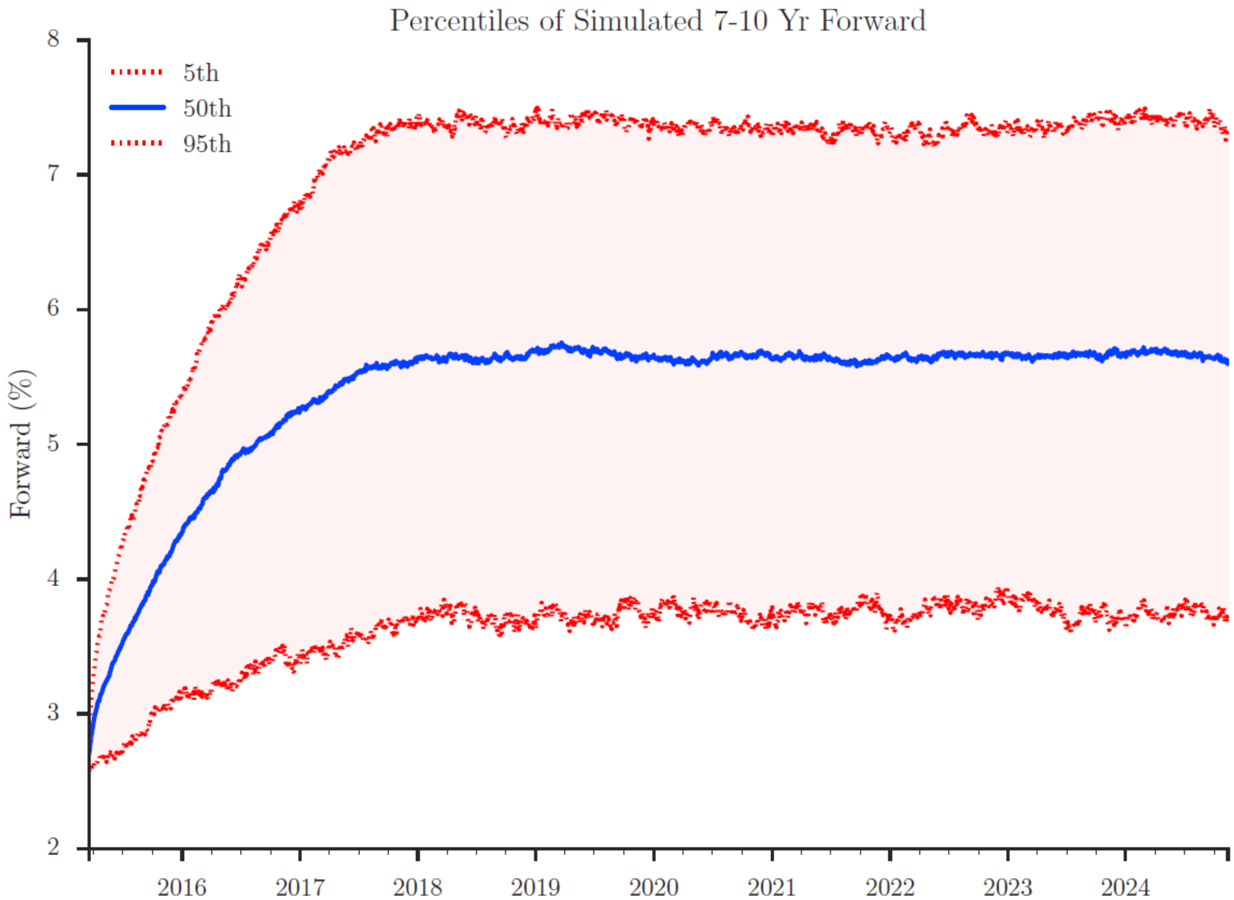

Figure 2

Simulated 10 year forwards from quasi-differenced log-linear transformed forward in Reduced Rank VAR

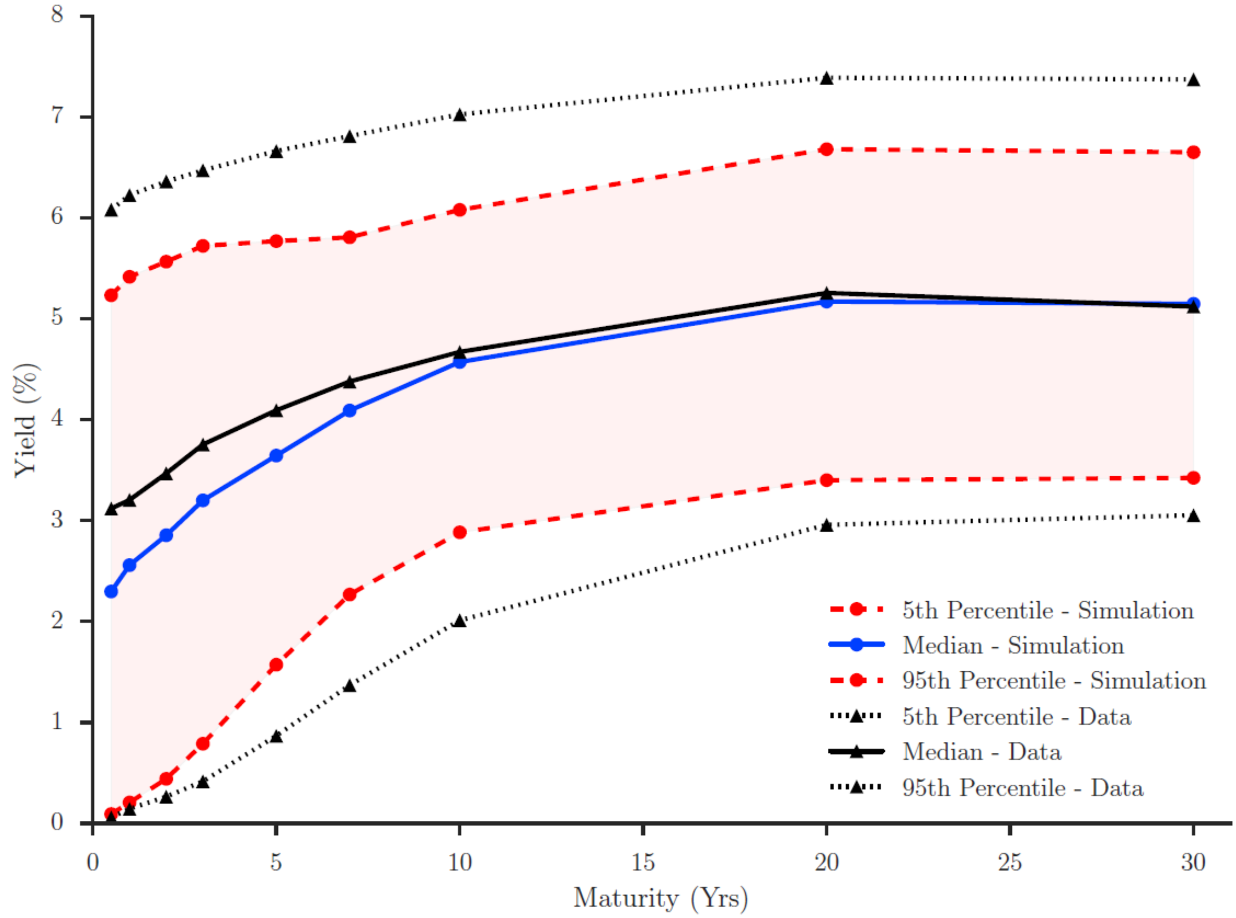

Figure 3

Simulated yield curve after 10 years from quasi-differenced log-lin transformed forwards in RRVAR

Plot includes quantiles of the original data as well

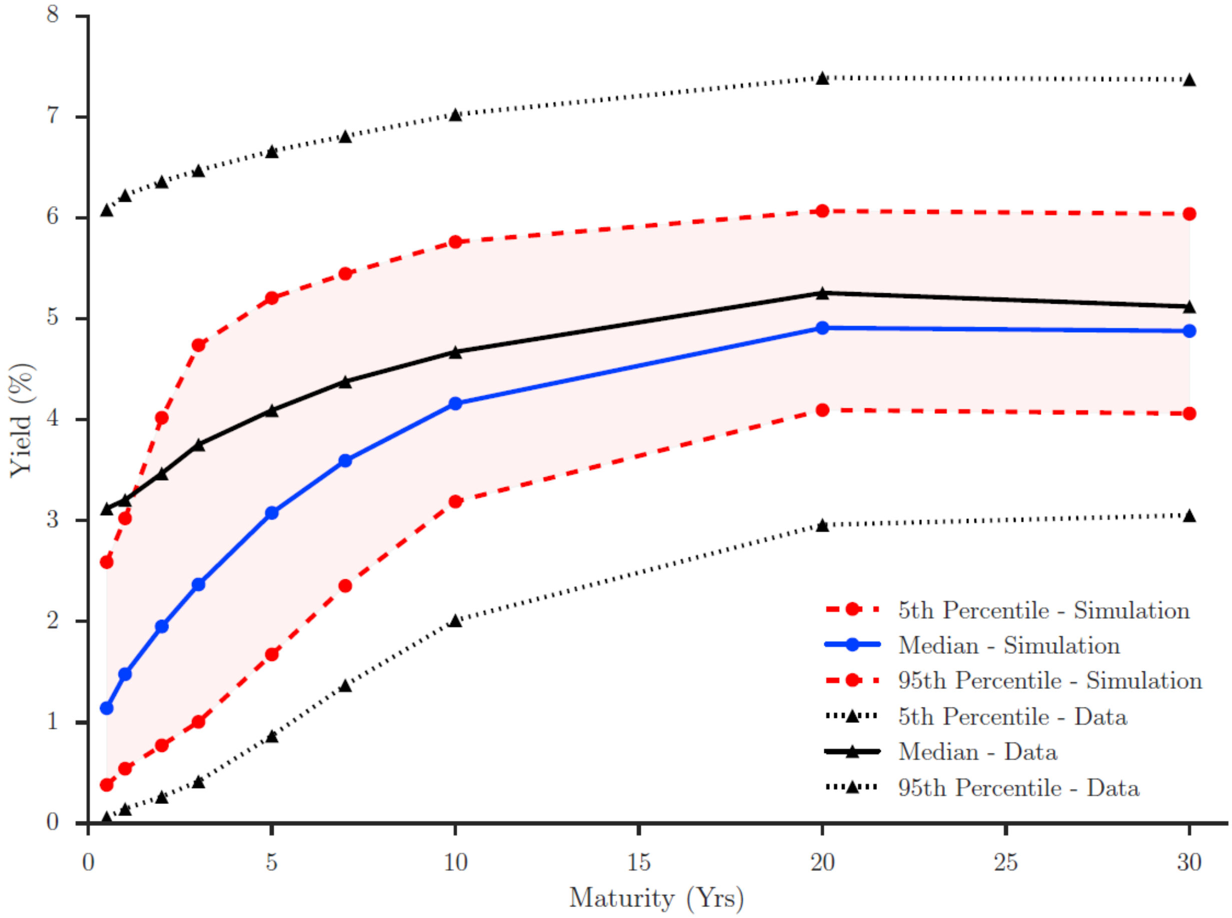

Figure 4

Simulated yield curve after 10 years from quasi-differenced log forward RRVAR

Plot includes quantiles of the original data as well

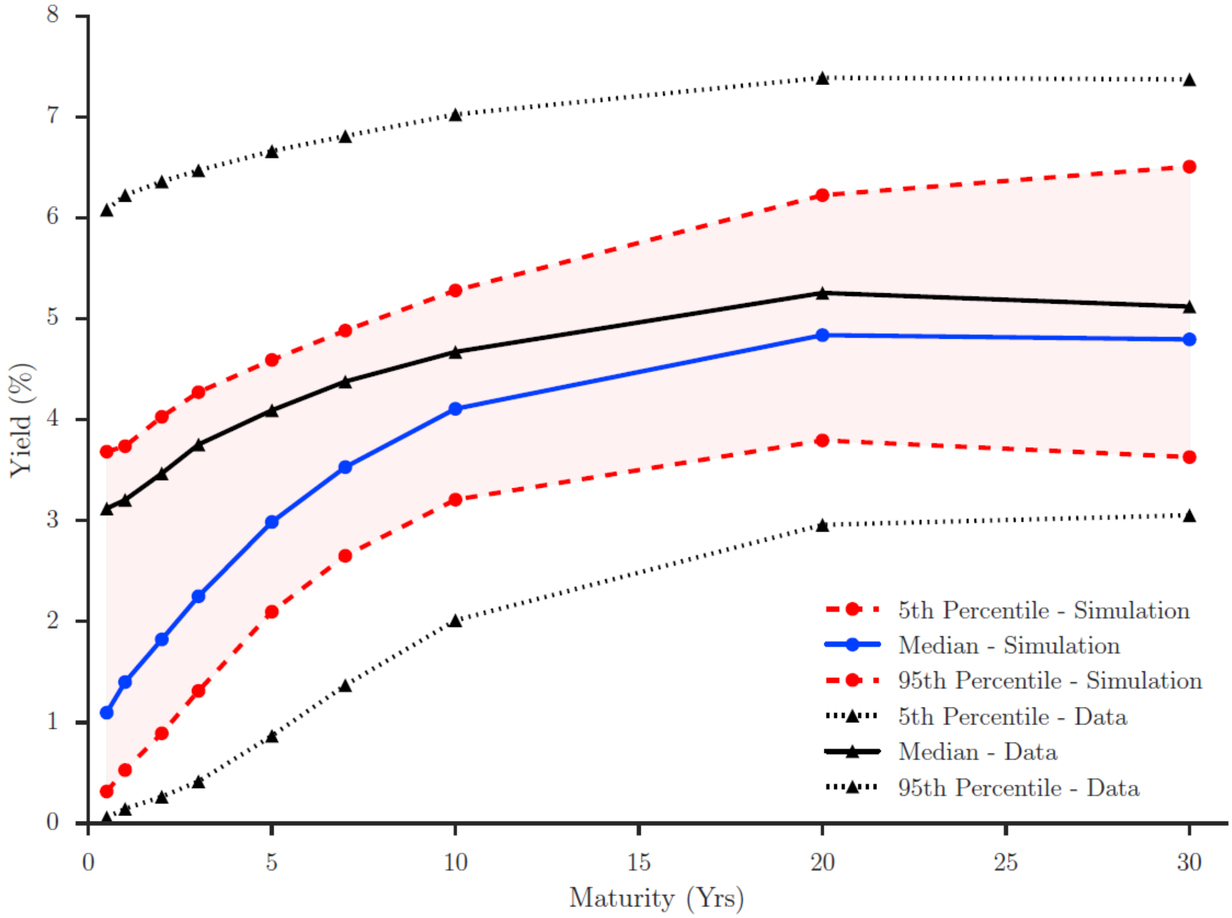

Figure 5

Simulated yield curve after 10 years from log forward RRVAR

Plot includes quantiles of the original data as well

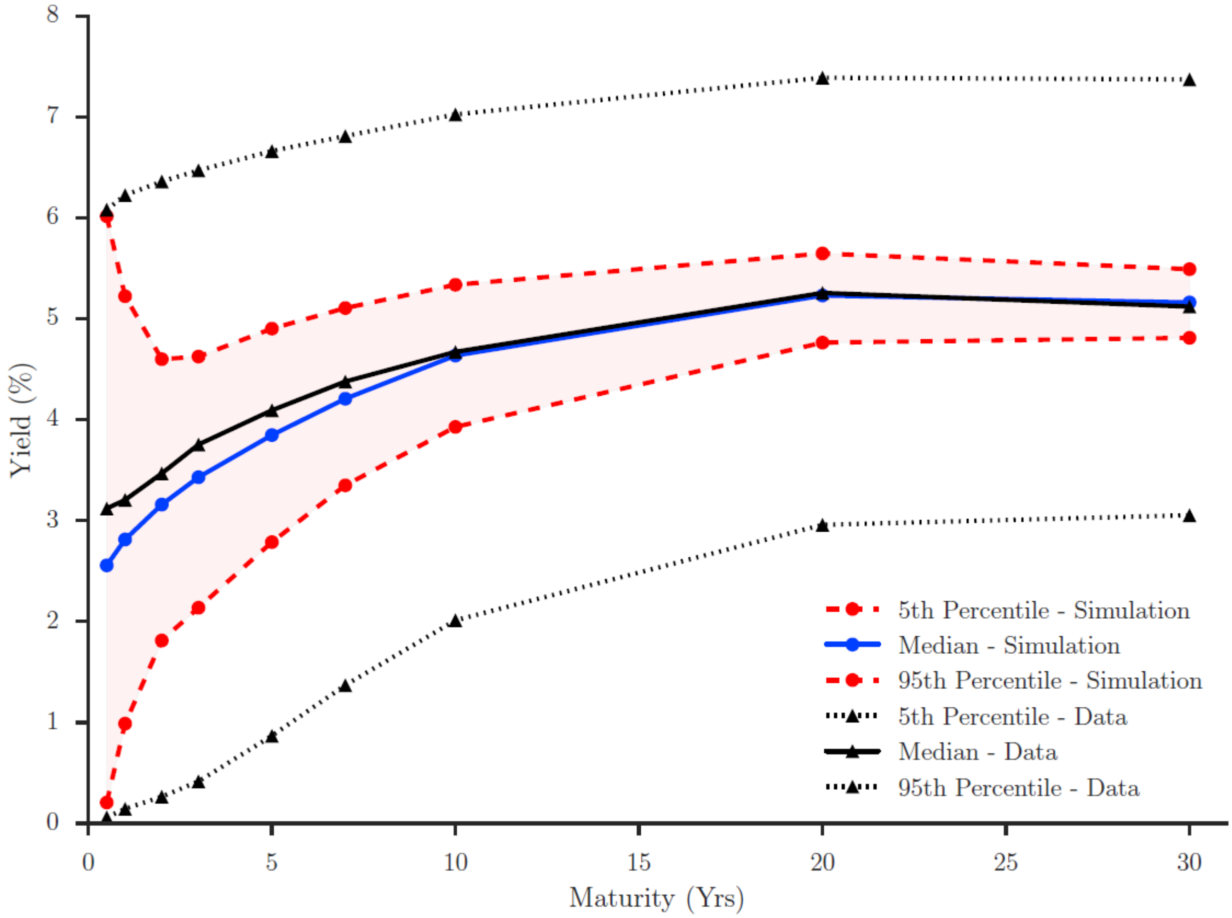

Figure 6

Simulated yield curve after 10 years from log-lin transformed forwards in RRVAR

Plot includes quantiles of the original data as well

REFERENCES

Bjork, Thomas,(2009), "An overview of interest rate theory", in *Handbook of Financial Time Series, (*ed. Andersen, Davis, Kreiss, Mikosch), Springer, 615-651

Brownlees, Christian and R. Engle(2010), "Volatilities, Correlations and Tails for Systemic Risk Modeling", NYU Stern Finance working paper

Christianson, Jens and Glenn Rudebusch,(2013) "Modeling yields at the zero lower bound: are shadow rates the solution?", manuscript, Federal Reserve Bank of San Francisco

Christensen, Jens, Jose Lopez and Glenn Rudebusch,(2013), "A Probability Based Stress Test of Federal Reserve Assets and Income", working paper, Federal Reserve Bank of San Francisco

Cox, J., J.Ingersoll, and S. Ross,(1985),"A theory of the term structure of interest rates", *Econometrica,* 53,385-407

Diebold, Francis and Glen Rudebusch,(2013), *Yield Curve Modeling and Forecasting: The Dynamic Nelson Siegel Approach,* Princeton University Press

Diebold, Francis and Canlin Li, (2006)"Forecasting the term structure of government bond yields", *Journal of Econometrics,* 130,337-364

Duffee, Greg (2002) "Term premia and interest rate forecasts in affine models", *Journal of Finance,* 57,405-443

Duffie, Darrell and Rui Kan,(1996), "A yield-factor model of interest rates," *Mathematical Finance,* 6, 379-406

Gourieroux, Christian and Joanna Jasiak, (2010) "Value at Risk" in *Handbook of Financial Econometrics,* ed (Ait-Sahalia and Hansen), Elsevier, 554-609

Heath, D. , R. Jarrow, and A. Morton,(1992), "Bond pricing and the term structure of interest rates: a new methodology for contingent claims valuation", *Econometrica,*60, 77-105

Ho, T. and S. Lee(1986), "Term structure movements and pricing interest rate contingent claims", *Journal of Finance,*41,1011-1029

Hull, J. and A. White,(1990), "Pricing interest-rate-derivative securities", *Review of Financial Studies",*3,573-592

Monfort, Alain, Fulvio Pegararo, Jean-Paul Renne and Guillaume Roussellet(2014), "Staying at Zero with Affine Processes: A New Dynamic Term Structure Model", manuscript, Banque de France

Muller, Ulrich, Michele Dacorogna and Peter Blum,(2010) "Bootstrapping the economy - a non-parametric method of generating consistent future scenarios", manuscript SCOR

Nelson, Charles and Andrew Siegel(1987) "Parsimonious modeling of yield curves", *Journal of Business,*60,473-489

Piazzesi, Monika, (2010), "Affine Term Structure Models", in *Handbook of Financial Econometrics,* ed(Ait-Sahalia and Hansen), North Holland

Vasicek, O. (1977), "An equilibrium characterization of the term structure", *Journal of Financial Economics,*5,177-188

---

[i] As the thirty year bond was discontinued for part of the sample period, we use the Treasury's interpolation based on the 20 year bond prices.

[ii] This assumes that the nominal yield on holding cash is zero. If indeed, this is negative because of storage costs, then this lower bound would be slightly negative.

[iii] These models use three principle components as the factors. Results with Nelson Siegel are similar or slightly worse as the Nelson Siegel factors are not quite as effective at reducing measurement errors as principle components.