

NBER WORKING PAPER SERIES

DO BONDS SPAN VOLATILITY RISK IN THE U.S. TREASURY MARKET? A
SPECIFICATION TEST FOR AFFINE TERM STRUCTURE MODELS

Torben G. Andersen
Luca Benzoni

Working Paper 12962
<http://www.nber.org/papers/w12962>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2007

We are grateful to Darrell Duffie, Michael Fleming, Bob Goldstein, Mike Johannes, Chris Jones, Rick Nelson, Jun Pan, Sam Thompson, and seminar participants at the Chicago Fed, St. Louis Fed, the Third T.N. Thiele Symposium on Stochastic Volatility, the 2005 International Conference on "Capital Markets, Corporate Finance, Money and Banking" at the Cass Business School, London, the 2006 Econometric Society Winter Meeting, the 2006 CIREQ-CIRANO-MITACS Financial Econometrics Conference, the 2006 Bank of Canada Conference on Fixed Income Markets, the Multivariate Modeling and Risk Management Conference, Sandbjerg, Denmark, the NBER Asset Pricing Meeting, March 2006, as well as Indiana University, the University of Chicago, the Federal Reserve Board of Governors, D.C., the University of Illinois at Chicago, and Carnegie Mellon University for helpful comments and suggestions. Further, we thank Mitch Haviv of GovPX for providing useful information on their data. Of course, all errors remain our sole responsibility. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Chicago, the Federal Reserve System, or the National Bureau of Economic Research.

© 2007 by Torben G. Andersen and Luca Benzoni. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Do Bonds Span Volatility Risk in the U.S. Treasury Market? A Specification test for Affine
Term Structure Models

Torben G. Andersen and Luca Benzoni

NBER Working Paper No. 12962

March 2007

JEL No. C14,C32,G12

ABSTRACT

We investigate whether bonds span the volatility risk in the U.S. Treasury market, as predicted by most 'affine' term structure models. To this end, we construct powerful and model-free empirical measures of the quadratic yield variation for a cross-section of fixed-maturity zero-coupon bonds ("realized yield volatility") through the use of high-frequency data. We find that the yield curve fails to span yield volatility, as the systematic volatility factors are largely unrelated to the cross-section of yields. We conclude that a broad class of affine diffusive, Gaussian-quadratic and affine jump-diffusive models is incapable of accommodating the observed yield volatility dynamics. An important implication is that the bond markets per se are incomplete and yield volatility risk cannot be hedged by taking positions solely in the Treasury bond market. We also advocate using the empirical realized yield volatility measures more broadly as a basis for specification testing and (parametric) model selection within the term structure literature.

Torben G. Andersen

Kellogg School of Management

Northwestern University

2001 Sheridan Road

Evanston, IL 60208

and NBER and CREATES

t-andersen@kellogg.northwestern.edu

Luca Benzoni

Federal Reserve Bank of Chicago

230 S. LaSalle St.

Chicago, IL 60604

and Carlson School of Management

University of Minnesota

lbenzoni@umn.edu

1 Introduction

The secondary U.S. Treasury market is among the largest, most liquid and important financial markets worldwide. In the third quarter of 2005, daily trading volume averaged \$539 billion, about tenfold the volume at the NYSE. The market is open round-the-clock, with trading taking place in New York as well as overseas (Fleming (1997)). Competition among dealers and brokers typically results in low bid-ask spreads, low brokerage fees, and fast order execution. The Federal Reserve System uses the market to implement its monetary policy through open market interventions. Due to their low risk, U.S. Treasuries are widely purchased by money managers as well as U.S. and foreign investors. Finally, these securities serve as both input and benchmark for the pricing of other financial instruments. As such, the pricing and hedging of U.S. Treasuries (and their derivatives) has been the focus of much attention.

Several years of academic research have fostered considerable progress in our understanding of the properties of the term structure of interest rates. Litterman and Scheinkman (1991) demonstrate that virtually all variation in U.S. Treasury rates is captured by three factors, interpreted as changes in ‘level,’ ‘steepness,’ and ‘curvature.’ This evidence has motivated much work on reduced-form term structure models, in which bond yields are expressed as an affine (or quadratic) function of a state vector (see, e.g., Duffie and Kan (1996), Duffie et al. (2000), and Piazzesi (2003)). These models have proven quite successful at capturing the cross-sectional properties of bond yields (see, e.g., Ahn et al. (2002, 2003), Brandt and Chapman (2002), and Dai and Singleton (2000)). However, some of their implications are still controversial. A major concern among market participants is how to hedge their positions in Treasury securities. In particular, Litterman, Scheinkman and Weiss (1991) note that the relative appeal of bonds with different maturities depends not only on the expected movement in future interest rates, but also on the uncertainty surrounding these moves. A key implication of most affine term structure models is that the quadratic variation of bond yields at any maturity is a linear combination of the concurrent term structure of yields. Consequently, according to these models, interest rate volatility risk is spanned and may be hedged by trading in a portfolio of bonds. In this paper, we empirically examine this prediction.

Previous studies have investigated this issue using data on the London Interbank Offer Rate (LIBOR) yields or swap rates as well as the associated derivatives, finding conflicting evidence. Collin-Dufresne and Goldstein (2002) conclude that swap rates have limited explanatory power for returns on at-the-money ‘straddles,’ i.e., portfolios mainly exposed to volatility risk. Motivated by this evidence, they propose an affine term structure model in which bond prices are unaffected by changes in volatility, referring to this feature as an ‘unspanned stochastic volatility’ (USV) restriction. Similarly, Li and Zhao (2005) find that some of the most sophisticated multi-factor dynamic term structure models have serious difficulties in hedging caps and cap straddles, even though they capture bond yields well. In stark contrast, Fan et al. (2003) find that swaptions and even swaption straddles can be well

hedged with LIBOR bonds alone, supporting the notion that bond markets are complete. The latter results are consistent with Litterman, Scheinkman, and Weiss (1991) who find the yield spreads of certain ‘butterfly’ combinations (which are very sensitive to volatility) to be highly correlated with the curvature factor.

More recently, other studies have investigated the properties of term structure models that embed the USV restriction and, again, the evidence is mixed. For instance, Collin-Dufresne, Goldstein, and Jones (2004, CDGJ) show that the LIBOR volatility implied by an affine multi-factor specification from the swap rate curve can be *negatively* correlated with the time series of volatility obtained from a standard GARCH approach. In response, they argue that an affine four-factor USV model delivers both realistic volatility estimates and a good cross-sectional fit. Thompson (2004) proposes a new class of specification tests that he applies to affine models of the LIBOR swap curve. Consistent with CDGJ, he detects problems with the unrestricted affine model at the short end of the yield curve. In contrast to CDGJ, however, he finds that the USV restriction is strongly rejected due to excessive pricing errors produced by the USV model. Jagannathan et al. (2003) find that an affine three-factor model can fit the LIBOR swap curve rather well. However, they identify significant shortcomings when confronting the model with data on caps and swaptions, thus concluding that derivatives must be used when evaluating term structure models. Building on this insight, Bikbov and Chernov (2004) investigate different versions of an affine three-factor model using data on Eurodollar futures and options. Consistent with CDGJ, they find that the volatility state variable implied by a USV model is more highly correlated with other volatility measures (e.g., options implied volatilities) than the volatility factor implied by unrestricted affine models. Like Thompson (2004) and in stark contrast to CDGJ, however, they reject the USV restriction. Remarkably, this happens not only when the model is confronted jointly with futures and options data, but also in the special case in which only futures data are used for estimation. An intriguing conjecture, inspired by such findings, is that affine models may be able to accommodate the dynamic structure of yield volatility after all, but that derivatives prices are necessary for efficient inference since measurement errors may render the theoretical link between yield levels and volatility elusive from observed bond prices alone.

We argue that the preceding literature has not focused on the fundamental yield volatility implications that characterize the affine model class. The basic prediction is that the instantaneous yield volatility is spanned by the contemporaneous cross-section of yields. Within the diffusive model class, a natural test of this property is to directly relate measures of realized quadratic variation to corresponding movements in the term structure of yields over short, say, daily, weekly or monthly horizons. From this perspective, the difficulty in gauging model adequacy stems more from the unobserved, or latent, nature of yield volatility than from the measurement errors associated with the extraction of yields from observed bond prices. However, recent contributions in the volatility modeling literature have documented, both theoretically and empirically, that realized volatility may be measured with

good precision at the daily level from intraday price data (see, e.g., Andersen et al. (2001, 2003b) and Barndorff-Nielsen and Shephard (2002b, 2004)).¹ These measures provide direct data-driven estimates of the underlying realized quadratic variation and therefore endow the notion of realized daily variation with concrete measurable content, independent of any modeling assumptions. Hence, we use a sample of high-frequency data on U.S. Treasuries covering more than a decade to construct yield volatility estimates. More specifically, we form series of intra-daily yields on Treasuries with three- and six-month, as well as one-, two-, five-, and ten-year maturity, and then estimate the quadratic variation by summing the squared intra-daily yield changes. We may consequently test directly whether bonds span volatility risk by relating our model-free realized volatility measures to the cross-section of daily bond yields.

The advantages of our approach are manifold. First, we test the generic affine yield volatility spanning condition directly. Hence, the analysis is independent of any particular specification of the underlying model. In contrast, Bikbov and Chernov (2004), CDGJ, and Thompson (2004) rely on specific affine term structure representations. As such, their analysis is a joint test of the USV restriction *and* a certain interest rate model. If the latter model is misspecified, the findings from such tests must be interpreted with caution. In addition, the main restrictions we test are based exclusively on the affine structure under the so-called ‘equivalent martingale’ or ‘pricing’ measure, so they are independent of whether the representation under the ‘actual’ measure is non-affine (Duarte (2004)). Second, we have access to a sequence of market prices for any given day which allows us to control for, and minimize, the impact of measurement errors on the extracted zero-coupon bond yield series. Third, we only use data for the specific fixed-income market we analyze. This sidesteps potentially serious concerns regarding the reliability of derivatives prices available from secondary over-the-counter markets due to liquidity and market microstructure issues. Fourth, we do not need to specify a truly compelling time series model for the conditional yield variance process. This is one approach previously adopted to gauge the coherence between the volatility dynamics implied by the model and the data; e.g., CDGJ and Dai and Singleton (2003). On the other hand, the availability of the high-frequency based realized volatility series allows us, if necessary or convenient, to construct simple, yet efficient, time series forecasts of future quadratic yield variation. In fact, such forecasts typically outperform those based on standard volatility models estimated from daily or lower frequency data (e.g. Andersen et al. (2003)). This facilitates direct comparison of our approach to prior contributions along this dimension. Fifth, we obtain realized yield variation measures for multiple maturities so that we can study the volatility dynamics across the term structure. This enhances the power of the empirical analysis as the spanning condition should hold for each individual maturity. Sixth, the latter point enables us to consider the more specialized predictions that stem

¹Occasionally, direct measures of low-frequency return variation have been obtained from cumulative higher-frequency squared returns, but without articulation of the theoretical basis. Early studies include French, Schwert, and Stambaugh (1987), Hsieh (1991), Poterba and Summers (1986), and Schwert (1989, 1990).

from particular popular models. For instance, it is common to describe the term structure of interest rates by using a multi-factor affine model in which a single factor determines the conditional variance of the state variables. Dai and Singleton (2000) refer to the ‘maximal’ version of this model as the $A_1(N)$ specification, where N is a positive integer equal to the number of latent factors. Bikbov and Chernov (2004), CDGJ, and Thompson (2004) use this model with $N = 3$ and/or $N = 4$ in their studies. A key implication of the $A_1(N)$ model is that the innovations to the quadratic variations of any pair of bond yields are perfectly correlated. By using our measures of realized volatility, we can examine in a fully non-parametric setting whether this condition is consistent with the evidence. Seventh, we have full flexibility in testing the affine spanning restriction at an arbitrary horizon, say, daily, weekly or monthly, or, theoretically, even at an intraday level.

We are able to expand our specification analysis beyond the traditional affine diffusive model class. The ‘quadratic term structure model’ studied by Ahn et al. (2002) has volatility spanning restrictions effectively identical to those of the affine diffusive class, so they are covered by our analysis. In contrast, the presence of jumps does change the spanning restriction qualitatively. For the affine jump-diffusion class, the direct spanning condition fails while the conditionally expected future quadratic variation, as before, is spanned by the yield cross-section. Hence, no alternative volatility forecast, based on current information, should possess explanatory power beyond what is captured by the yield cross-section. Thus, we conduct a second set of specification tests where we exploit past realized volatility measures as additional forecast variables. The null hypothesis of an affine jump-diffusion implies that such auxiliary forecasts should be ineffective and provide no significant improvement to the explanatory power of forecasts based solely on the yield cross-section.

Our analysis hinges critically on the quality of our nonparametric realized yield volatility measures. Consequently, we perform a variety of robustness checks to assess the reliability of these empirical quadratic variation proxies. Most significantly, we estimate an EGARCH-type semi-nonparametric (SNP) model (see, e.g., Gallant and Nychka (1987)) for the daily three-month maturity yield. This model is used to compute one-day-ahead volatility forecasts, which we contrast to the corresponding realized volatility series. We confirm that the properties of these series are consistent with the typical relationships between volatility forecasts and the subsequent volatility realizations. Moreover, the quantitative properties of the realized yield volatility series, both in terms of their general dynamic properties and their ‘unconditional’ term structure features, are shown to be similar to comparable evidence from the literature. We conclude that our realized yield variation measures are unlikely to be subject to any systematic measurement errors that may render the interpretation of our results problematic.

We use our realized volatility measures to test the affine yield spanning conditions. To this end, we estimate linear regressions in which the dependent variable is the yields’ realized volatility. At each date, we compute average daily bond yields and we extract orthogonal principal components from

these series. We use the yields' principal components as explanatory variables in our regressions. In stark contrast with the notion that the yields' quadratic variation is a linear combination of the bond yields, the explanatory power of these regressions is, in most cases, nearly zero. For instance, the R^2 coefficient is less than 0.6% when the dependent variable is the realized volatility of yields with maturity of two or more years. When the dependent variable is the realized volatility of yields with maturity of one year or less, the R^2 coefficient shows little improvement, ranging from 1% to approximately 4%. Interestingly, we find that the first three principal components (i.e., level, slope, and curvature) typically have insignificant coefficients in these regressions. Higher-order principal components often enter significantly, although with limited explanatory power. This finding is strikingly at odds with the notion that curvature is related to interest rate volatility. We confirm that these results apply also when the spanning condition is tested using weekly and monthly realized volatility measures. Moreover, an analysis of sub-samples shows that the volatility spanning condition is violated consistently across the sample period. Likewise, we reject the auxiliary implications of those affine multi-factor term structure models in which a single factor governs the conditional variance of the state variables. Finally, we document that there is a substantial degree of predictable yield variation which is captured by simple time series models but is largely unrelated to the yield curve. This strongly repudiates the conditional volatility spanning condition that applies within the more general affine jump-diffusion setting.

In conclusion, we find compelling evidence indicating that interest rate volatility cannot be extracted from the cross section of bond yields in the U.S. Treasury market. This finding underscores the importance of adapting some variant of USV—within or outside the affine setting—to term structure modeling, and in particular to applications that require a good fit to the yield volatility dynamics like the hedging of interest rate volatility risk or the pricing of fixed-income derivatives. Furthermore, at a general level, the pricing of yield risk across the business cycle is of great interest in financial and monetary economics. Affine term structure models provide such estimates by relating empirical measures of the yield risk premium to the corresponding model implied factor exposure. This risk exposure is proportional to the yield factor volatility. Systematic errors in model implied yield factor volatility are thus likely to be absorbed into the inferred market prices of risk, potentially producing misleading evidence regarding the market pricing of yield factor risk.

The remainder of the paper is organized as follows. In Section 2, we discuss the link between Treasury yields and their quadratic variation in the context of affine diffusive term structure models. We clarify how this relation is affected by the presence of jumps and we introduce our realized volatility measures of the yields' quadratic variation. In Section 3, we describe the U.S. Treasury market data and we document the salient features of our volatility series and forecasts relative to the existing literature. Section 4 contains our main empirical findings. In Section 4.1, we focus on affine diffusive models, while Section 4.2 extends the analysis to the presence of jumps. Concluding remarks are in Section 5.

2 Affine Term Structure Models

This section discusses the empirical implications of the general continuous-time affine model class for the yield volatility of zero-coupon bonds. These models provide testable restrictions that apply not only to standard affine multi-factor diffusions but also to the recently popular quadratic-Gaussian models. Moreover, the testable restrictions arise directly from the affine specification of the diffusion coefficient which is invariant across the equivalent martingale (risk-neutral) and the physical (actual) probability measures. Hence, the restrictions remain valid for the generalization to the ‘completely affine’ class proposed by Duffee (2002) and they also cover models which allow for a more general non-affine drift under the physical measure, as proposed by Duarte (2004) and further analyzed in Cheridito et al. (2005). There are also interesting predictions for the affine jump-diffusion representations of the term structure that we delineate from the pure diffusion case. The explicit linkages between yield levels and yield variation that we develop in detail below form the basis for our specification analysis of the entire model class through spanning conditions involving nonparametric realized volatility measures.

2.1 Bond Yields and Yield Volatility in Affine Diffusion Models

We first review features of affine diffusive term structure models that are relevant for our empirical inquiry. Following Duffee and Kan (1996) and Dai and Singleton (2000), the short term interest rate, $y_0(t)$, is an affine (i.e., linear-plus-constant) function of a vector of state variables, $X(t) = \{x_i(t), i = 1, \dots, N\}$:

$$y_0(t) = \delta_0 + \sum_{i=1}^N \delta_i x_i(t) = \delta_0 + \delta'_X X(t), \quad (1)$$

where the state-vector X has risk-neutral dynamics

$$dX(t) = \mathcal{K}(\Theta - X(t))dt + \Sigma\sqrt{S(t)}dW^Q(t). \quad (2)$$

In equation (2), W^Q is an N -dimensional Brownian motion under the so-called pricing (Q) measure, \mathcal{K} and Σ are $N \times N$ matrices, Θ is an $N \times 1$ vector, and $S(t)$ is an $N \times N$ diagonal matrix with the i th diagonal element given by $[S(t)]_{ii} = \alpha_i + \beta'_i X(t)$, where α_i is a scalar and β_i is an $N \times 1$ vector.

Within this setting, one can find (effectively) closed-form expressions for the time- t price of a zero-coupon bond with time-to-maturity τ :

$$P(t, \tau) = e^{A(\tau) - B(\tau)'X(t)}, \quad (3)$$

where the scalar function $A(\tau)$ and the $N \times 1$ vector of functions $B(\tau) = \{B_k(\tau), k = 1, \dots, N\}$ solve a system of ordinary differential equations.

This result establishes a fundamental link between the state-vector $X(t)$ and the term structure of bond yields. Specifically, the time- t yield, $y_\tau(t)$, on a zero-coupon bond with time-to-maturity τ is

defined by the relationship, $P(t, \tau) = e^{-\tau y_\tau(t)}$. Thus, in view of equation (3) we have,

$$y_\tau(t) = -\frac{A(\tau)}{\tau} + \frac{B(\tau)'}{\tau} X(t). \quad (4)$$

Next, we write the corresponding $J \times 1$ vector of (observed) zero-coupon bond yields as $Y(t) = \{y_{\tau_j}(t), j = 1, \dots, J\}$, where we assume we observe more yields than there are state variables, i.e., $J \geq N$, and we define, for notational simplicity, the $N \times 1$ vector $A = \left\{ \frac{A(\tau_j)}{\tau_j}, j = 1, \dots, J \right\}$ and the $N \times J$ matrix $B = \left\{ \frac{B(\tau_j)}{\tau_j}, j = 1, \dots, J \right\}$. We may then write the above relation as a system of equations,

$$Y(t) = -A + B'X(t). \quad (5)$$

Typically, it is assumed that the B matrix has full row rank, N , which implies that all state variables affect the contemporaneous cross-section of bond prices and yields. In this case we may invert the system to express the state vector as an affine function of the yields. We have, for C a $J \times 1$ vector,

$$X(t) = (B'B)^{-1} B' (Y(t) + A) = C + (B'B)^{-1} B' Y(t). \quad (6)$$

Meanwhile, by Itô's Lemma the yield, y_τ , in equation (4) follows a diffusion process,

$$dy_\tau(t) = \mu_{y_\tau}(X(t), t) dt + \frac{B(\tau)'}{\tau} \Sigma \sqrt{S(t)} dW^Q(t). \quad (7)$$

Consequently, the (instantaneous) variation of the yield, given as the squared yield volatility, is

$$V_{y_\tau}(t) = \frac{B(\tau)'}{\tau} \Sigma S(t) \Sigma' \frac{B(\tau)}{\tau}. \quad (8)$$

Since the elements of the $S(t)$ matrix, by definition, are affine in the state vector $X(t)$, as stated below equation (2), and $X(t)$ in turn is an affine function of $Y(t)$ from equation (6), it follows, for any τ , that we can find a set of constants $a_{\tau,j}$, $j = 0, \dots, J$, so that

$$V_{y_\tau}(t) = a_{\tau,0} + \sum_{j=1}^J a_{\tau,j} y_{\tau_j}(t). \quad (9)$$

Hence, the instantaneous variation of (constant maturity) yields is tied to the contemporaneous level of the yields and thus to the cross-section of bond prices through the affine mapping in equation (9). Since the yield variation also is directly related to the time series properties of the yields, it plays a dual role in standard affine diffusive term structure models. CDGJ highlight the implied link between the bond yields and the short rate variation. The above shows that this relationship remains valid for *any* fixed maturity yield, implying a range of simultaneous constraints across the yield volatility spectrum.² A noteworthy implication is that an investor can use a portfolio of zero-coupon bonds to hedge volatility risk in the Treasury market.

²As noted, an exception is the USV models of Casassus et al. (2004), Collin-Dufresne and Goldstein (2002) and CDGJ where the full rank condition on the B matrix in equation (5) is violated and the inversion argument behind equation (6) does not apply. A related class of models is explored in Kimmel (2004).

We now derive the implications of equation (9) for discrete-time data to enable testing on the basis of bond price data observed at a reasonably high frequency. First, we recall the definition of the quadratic variation process for the constant maturity yield y_τ initiated at time $t_0 = 0$,

$$QV_{y_\tau}(t) \equiv \int_0^t V_{y_\tau}(s) ds. \quad (10)$$

The quadratic variation process is positive and strictly increasing in t , $t > 0$, as long as the volatility coefficient remains bounded away from zero. It represents the cumulative or integrated yield variation over the given interval. The affine model restrictions relate naturally to the increments in the yield quadratic variation process over daily or intraday periods $[t, t+h]$, $h > 0$ which we denote by

$$QV_{y_\tau}(t+h, h) \equiv QV_{y_\tau}(t+h) - QV_{y_\tau}(t) = \int_t^{t+h} V_{y_\tau}(s) ds. \quad (11)$$

Next, observe that equation (9) implies,

$$\int_t^{t+h} V_{y_\tau}(s) ds = a_{\tau,0} + \sum_{j=1}^J a_{\tau,j} \int_t^{t+h} y_{\tau_j}(s) ds. \quad (12)$$

We may rewrite equation (12) by defining the average yield of y_{τ_j} over $[t, t+h]$ as $\bar{y}_{\tau_j}(t+h, h)$. This term corresponds to the integral on the extreme right of equation (12). Then, also exploiting equation (11), we obtain the following affine diffusion restriction,

$$QV_{y_\tau}(t+h, h) = a_{\tau,0} + \sum_{j=1}^J a_{\tau,j} \bar{y}_{\tau_j}(t+h, h). \quad (13)$$

We term this expression *the fundamental affine yield variation spanning condition*. The yield levels on the right hand side are readily approximated through empirical observations on the intraday yields—or more crudely the yields at the close of trading—across the maturity spectrum. The quadratic variation increment on the left-hand-side is slightly more delicate, as it cannot be measured with precision from daily bond price data. Perhaps as a consequence, the quadratic yield variation has not been the focus of direct measurement or testing within the term structure literature. Instead, most existing studies rely on parametric conditional yield variance estimates or implied volatility measures backed out from derivatives prices. Although this approach can be rigorously justified by theory, as exemplified below, the substitution of an alternative volatility proxy in place of the quadratic variation is not innocuous. It inevitably entails a loss of power in terms of testing the affine spanning condition.

2.2 Spanning Restrictions for the Conditional Yield Variance

The pure diffusive semi-martingale representation of bond prices implies that the *predictable* yield changes over short daily or intraday periods are negligible (order dt^2) in terms of their contribution to the quadratic yield variation compared to the yield innovations (order $dW^Q(t)^2 = dt$). Practically,

this means we may ignore the conditional mean of the yield changes in computing the conditional yield variance over short horizons. Letting the integer $n \geq 1$ denote the number of equidistant yield changes sampled over the (short) interval $[t, t + h]$, we thus invoke the following approximate martingale relation,

$$E_t^P \left[y_\tau \left(t + \frac{i h}{n} \right) \right] = y_\tau(t), \quad i = 1, \dots, n, \quad (14)$$

where the subscript t indicates that the expectation is evaluated conditional on time t information, and the superscript P indicates the so-called actual or ‘physical’ probability measure as opposed to the equivalent martingale pricing measure, Q .

It follows that the conditional yield variance is given by,

$$Var_t^P [y_\tau(t + h)] = E_t^P \left[\sum_{i=1, \dots, n} \left(y_\tau \left(t + \frac{i h}{n} \right) - y_\tau \left(t + \frac{(i-1)h}{n} \right) \right)^2 \right]. \quad (15)$$

Equation (15) holds for an arbitrary n , so by letting n increase towards infinity we have, by basic properties of the quadratic variation process, that

$$Var_t^P [y_\tau(t + h)] = E_t^P [QV_{y_\tau}(t + h, h)]. \quad (16)$$

This relation highlights important distinctions between these two concepts of yield volatility. The conditional variance is a forward looking expectation of the future sample path variation, and thus fundamentally an ex-ante concept. In contrast, the quadratic variation denotes the actual realized sample path variation, so it is an ex-post (realization) measure. If the volatility is (conditionally) deterministic as when volatility is constant, the two notions of yield variation coincide. In general, however, the yield variation has a large, genuinely unpredictable, innovation component which renders volatility stochastic. As such, the sample variability of the quadratic yield variation process inevitably will exceed that of the conditional yield variance process because sample realizations, by construction, fluctuate more than their a priori expectations.

The spanning condition in equation (13) ties the contemporaneous yield level to the yield variation in terms of realizations. This is a much more powerful and stringent requirement than the corresponding prediction based on the variation in the conditional yield forecasts. In order to formally derive the latter implication of the affine term structure models, we first substitute equation (13) into equation (16) to obtain,

$$Var_t^P [y_\tau(t + h)] = a_{\tau,0} + \sum_{j=1}^J a_{\tau,j} E_t^P [\bar{y}_{\tau_j}(t + h, h)]. \quad (17)$$

This prediction is valid only under the P measure, as it is related directly to the observed time series variation of the yields. The conditional moments over discrete (non-infinitesimal) horizons will differ across the measures due to the differential volatility drift specifications, even if the instantaneous volatility (and quadratic variation) is identical under P and Q . Now, assuming that the diffusion model

is also affine under the physical measure, which still allows for the ‘essentially’ affine model of Duffee (2002) but excludes the extension by Duarte (2004), the future expected yields will be given as a linear combination of the current cross-section of yields, so that

$$Var_t^P [y_\tau(t+h)] = b_{\tau,0} + \sum_{j=1}^J b_{\tau,j} y_{\tau_j}(t). \quad (18)$$

This is a formal specification of the type of affine spanning condition that has been investigated in prior empirical studies. It is inherently less powerful in terms of testing the underlying model than equation (13) and it requires the model to be affine under both the P and Q measures. A final caveat is that it is valid only if the true conditional variance process that appears on the left hand side of (18). Using an ad hoc specified time series model to generate these volatility forecasts will inevitably induce a degree of measurement error into any testing procedure based on this relationship.

A similar logic applies if we use an implied volatility forecast extracted from derivatives prices in lieu of the time series model based forecast, except that the forecasts now are formed under the pricing measure, Q . Of course, this approach assumes that the derivatives pricing model is correctly specified and quality data on derivatives are available. In that case, the implied conditional variance forecasts (under Q) should also be spanned by the cross-section of yields. As for the fundamental spanning condition (13) this hinges only on the model being affine under Q , so it applies also for Duarte style extensions of the basic affine model. On the other hand, the forecast horizon must equal the maturity of the derivatives contracts, typically necessitating monthly volatility predictions rather than daily or weekly forecasts, thus reducing the forecast comparison sample and lowering test power correspondingly.

We may also explicitly relate daily changes in the conditional yield variance to the evolution of the yield cross-section. Letting, generically, $\Delta c(t) = c(t) - c(t-h)$, we have from equation (18) that

$$\Delta Var_t^P [y_\tau(t+h)] = \sum_{j=1}^J b_{\tau,j} \Delta y_{\tau_j}(t). \quad (19)$$

Of course, we can derive an equivalent expression for changes in the implied volatility forecasts under the Q measure. Several studies employ such a specification of the affine spanning condition. This approach is obviously closely related to the specification in equation (18) so we focus on the latter in the empirical sections below. We have, however, confirmed that the findings are qualitatively similar, albeit even less flattering, for the basic affine model restrictions when tested using the representation in (19).

2.3 A Framework for Testing the Fundamental Affine Spanning Condition

The previous sections have outlined the basic volatility spanning conditions implied by standard affine diffusive term structure models. We now explicitly compare and contrast different approaches that may

be adopted for testing these restrictions. The fundamental affine yield variation spanning condition in equation (13) directly motivates the regression,

$$QV_{y_\tau}(t+h, h) = a_{\tau,0} + \sum_{j=1}^J a_{\tau,j} \bar{y}_{\tau_j}(t+h, h) + \eta(t+h, h). \quad (20)$$

In the ideal scenario where the increments to the quadratic yield variation and the contemporaneous mean yield levels are all measured perfectly, the error term of the regression is identically zero and the associated R^2 coefficient should be unity. In practice, as we discuss in detail later on, there is noise in the extraction of the average zero-coupon yields for different maturities and some measurement error in the estimation of the ex-post realized quadratic yield variation, but these sources of noise are relatively small compared to the sample variation in the yield volatility and interest levels. In other words, under the null hypothesis of a standard affine diffusive model, we would expect the yield cross-section to have a very high degree of explanatory power for the observed yield variation.

In contrast to the above specification, the preceding literature has invoked largely diagnostic procedures based on the predictive relation in equations (18) and (19). One way to proceed is to explore the correlation between the cross-sectional yield based volatility forecast embedded in equation (18) and common time series oriented volatility forecasts based on, e.g., a GARCH model or the implied volatility inherent in derivatives prices. Denoting one such volatility forecast for time $t+h$, constructed on the basis of time- t information, by $F(t+h, t)$, one may directly study the coherence between the alternative forecasts through the regression,

$$F(t+h, t) = c_{\tau,0} + \sum_{j=1}^J c_{\tau,j} y_{\tau_j}(t) + \eta_f(t, h). \quad (21)$$

Under the null hypothesis that the standard affine diffusive model is valid, one would expect a high degree of correlation between the alternative forecasts and a correspondingly high R^2 coefficient from the regression. Obviously, there are measurement errors in the yields and, more importantly, misspecification and estimation errors associated with the auxiliary time series or derivative price based volatility measure so there is no sharp prediction concerning the actual degree of correlation but rather an expectation of qualitative coherence.

A formal, yet practical approach to the question may be developed by casting the issue within the framework of equation (20). Specifically, under the null hypothesis, the conditional variance is spanned by the yield cross-section so any auxiliary forecast variable should not have any predictive power for the future yield variation over-and-above the information conveyed by the contemporaneous yields. Taking advantage of the relations in equations (16) and (18) we may test this property via the regression,

$$QV_{y_\tau}(t+h, h) = d_{\tau,0} + \sum_{j=1}^J d_{\tau,j} y_{\tau_j}(t) + d_F F(t+h, t) + \eta_{af}(t+h, t). \quad (22)$$

The coefficient d_F should be insignificant and the associated R^2 should not be significantly higher than when the $F(t+h, t)$ variable is excluded from the regression. However, in contrast to the regression (20), there is no presumption that the explanatory power should be close to perfect. This is a forecast regression for the future yield variation which is likely to possess a genuinely stochastic (unpredictable) component, so the explanatory power may well be relatively low. Equation (20), on the other hand, states that changes in the future yield variation, whether predictable or not, go hand in hand with corresponding shifts in the future yield curve, thus producing a (near) perfect association.

In summary, the regression approach developed above provides a unifying framework for testing the affine yield variation spanning property. The specifications all explore the extent to which the observed yield variation may be rationalized by a set of explanatory variables including the yield cross-section. In equation (20) the relation is contemporaneous and should, in principle, be perfect. If this fundamental property is rejected, then the forecast regressions (22) speak directly to the amount of predictability provided by the yield curve relative to alternative forecasts.

2.4 Cross-Maturity Correlation in Yield Volatility

The volatility spanning condition in equation (13) applies to any affine model of the form (1)-(2). Naturally, additional restrictions may be operative in more specific representations. The literature has documented a trade-off in the ability of affine models to capture the yield cross-section and the yield volatility simultaneously. The more factors are allowed to drive the volatility dynamics the less flexibility is allowed in specification of the risk premia and the yield correlation structure, which hampers the cross-sectional fit. Recent empirical studies favor models with a single factor determining the conditional variance of the state variables. Dai and Singleton (2000) refer to the ‘maximal’ version of this model as the $A_1(N)$ specification, where N , as before, represents the number of state variables. For example, Bikbov and Chernov (2004), CDGJ, and Thompson (2004) use this model with $N = 3$ and/or $N = 4$ in their studies.

Because of the prominence of this specification, we review the specific yield variation restriction that applies within this model class. Specifically, it is straightforward to show for an $A_1(N)$ model that the quadratic variations of any pair of yields y_{τ_1} and y_{τ_2} are perfectly correlated,

$$\text{corr}(V_{y_{\tau_1}}, V_{y_{\tau_2}}) = 1, \quad (23)$$

and a similar condition not surprisingly also applies to the *innovations* in the quadratic variation of any pair of yields,

$$\text{corr}(dV_{y_{\tau_1}}^{\text{stochastic}}, dV_{y_{\tau_2}}^{\text{stochastic}}) = 1. \quad (24)$$

This is, of course, also not a novel insight. However, as for the volatility spanning condition, these relations have not previously been subjected to direct empirical scrutiny based on nonparametric or model-free measures of quadratic yield variation.

2.5 Natural Extensions

2.5.1 Quadratic Term Structure Models

The yield-volatility spanning condition is readily extended to cover the so-called Quadratic Term Structure Model (QTSM) introduced by Ahn et al. (2002). In fact, as noted by Ahn et al. (2003), the QTSM is isomorphic to the ATSM in its mechanism for generating volatility as the yield variation remains proportional to the level of the state variables. Ahn et al. (2002) and Cheng and Scaillet (2005) formally show how the quadratic models may be embedded in an affine model with an extended state vector. Hence, as long as we allow for a sufficiently large dimensional state vector our analysis automatically covers the quadratic models as well. This is of particular interest as one of the prime motivations behind these models is that they should be better at replicating the yield volatility than the regular affine models.³

2.5.2 Affine Jump-Diffusion Term Structure Models

A modification of the yield spanning condition is required to accommodate the possibility of jumps in the state variables and yields. This is empirically relevant as the evidence strongly suggests that, e.g., macroeconomic announcements induce instantaneous jumps in the yields upon release.⁴ Following Duffie et al. (2000), the state vector X in affine jump-diffusion models has Q -dynamics

$$dX(t) = \mathcal{K}(\Theta - X(t))dt + \Sigma\sqrt{S(t)}dW^Q(t) + Zdq^Q(t), \quad (25)$$

where q^Q is a Poisson jump-arrival process with intensity $\lambda(X) = \lambda_0 + \lambda'_X X$, Z is an $N \times 1$ vector process with fixed probability distribution ν^Q , and $E^Q [Z(t)J(t)'] = V_{ZZ}$. The corresponding jump vector process, $J(t) \equiv \Delta X(t) = Z(t)dq^Q(t)$, is non-zero only if a jump actually occurs. Both q^Q and ν^Q are independent of W^Q . Under these assumptions, instead of equation (7) we have

$$dy_\tau(t) = \mu_{y_\tau}(X(t), t) dt + \frac{B(\tau)'}{\tau} \left[\Sigma\sqrt{S(t)}dW^Q(t) + Zdq^Q(t) \right] \quad (26)$$

and we further obtain,

$$QV_{y_\tau}(t, h) = \int_t^{t+h} \frac{B(T-s)'}{T-s} \Sigma S(s) \Sigma' \frac{B(T-s)}{T-s} ds + \sum_{t-h \leq s \leq t} \left[\frac{B(T-s)'}{T-s} J(s) J(s)' \frac{B(T-s)}{T-s} \right]. \quad (27)$$

³There is a caveat, akin to the invertibility condition of standard affine models. It is possible that the quadratic models naturally embed a form of the USV restriction. This is not readily determined by analytic means, but we have confirmed that the model estimated by Ahn et al. (2003) induces a near perfect linkage between the concurrent yield cross-section and yield volatility as should be the case in an affine model. We accomplish this by simulating the model estimated by Ahn et al. and then ex-post regressing the associated quadratic yield variation on the concurrent yield curve shifts. These results are available upon request from the authors.

⁴This extensive literature includes, e.g., Andersen et al. (2006b), Balduzzi et al. (2001), Bollerslev et al. (2000), Fleming and Remolona (1999), Johannes (2004), and Piazzesi (2005).

Notice that jump realizations induce a quadratic form dependency into the relation between the quadratic yield variation and the state variables so the basic affine spanning restriction is no longer valid. However, as the label ‘affine’ jump-diffusion indicates, it is still the case that the state variables span the first two yield moments. Indeed, upon taking conditional expectations we find,

$$E_t^Q \left[\sum_{t \leq s \leq t+h} J(s)J(s)' \right] = E_t^Q \left[\int_t^{t+h} (\lambda_0 + \lambda'_X X(s)) ds \right] V_{ZZ}. \quad (28)$$

Hence, the *expected* jump contribution to the quadratic variation process is an affine function of the state variables and, since the state variables still may be written as a linear combination of the yields under the Q measure, this property carries over to the full expected quadratic yield variation process. Then, following the line of reasoning in Section 2.2, we obtain the following spanning condition for the (short) future time interval $[t, t + h]$,

$$E_t^Q [QV_{y_\tau}(t + h, h)] = b_{\tau,0} + \sum_{j=1}^J b_{\tau,j} y_{\tau_j}(t). \quad (29)$$

This spanning condition is identical to equation (18) as the latter also holds under the Q measure, and it is straightforward to derive the corresponding version of equation (19) as well. Consequently, the extension to include jumps has no impact on the conditional yield variance forecast spanning conditions under the Q measure. The identical spanning restriction for conditional yield variance forecasts under the actual probability measure, P , applies in the affine jump-diffusion setting as well, if the expected jump distribution and jump intensity continue to be affine functions of the state variables under P , as described in Section 2.2.

The conclusion is that the general formulation in Section 2.2, providing spanning conditions for the conditional yield variance, remains unaltered in the affine jump-diffusion case, while the much stricter realization-by-realization spanning constraint in equation (13) no longer applies. In terms of the regression framework developed in Section 2.3, equation (20) will no longer hold but equation (22) remains valid. Hence, the jump-diffusive case is subsumed within testing procedures based on the latter relation without the need for auxiliary assumptions.

2.6 Realized Yield Volatility Measurement

The concept of realized volatility has been advocated in the recent volatility measurement and forecasting literature as a mean of approximating the daily realizations from the asset return quadratic variation process. Some early applications of this idea may be found in Andersen and Bollerslev (1997, 1998). More formal theoretical justification and assessments of the associated (continuous record) asymptotic theory is provided in Barndorff-Nielsen (2002a, 2002b, 2004) and Andersen et al. (2003a, 2004). The approach is fully nonparametric, and hence model-free, and utilizes the cumulative

squared high-frequency intraday returns to obtain feasible return variation measures. It is straightforward to apply for direct measurement of the yield variation of a bond with maturity τ . Specifically, we compute the realized volatility of the yield y_τ over $[t, t + h]$ by

$$v_{y_\tau}(t + h, h; n)^2 = \sum_{i=1, \dots, n} \left(y_\tau \left(t + \frac{ih}{n} \right) - y_\tau \left(t + \frac{(i-1)h}{n} \right) \right)^2. \quad (30)$$

The realized yield volatility converges, for ever more frequent sampling, to the underlying realization of the quadratic yield variation process. Equation (13) links the quadratic variation of a zero-coupon yield with maturity τ to the cross-section of bond prices. In our application, we rely on the realized yield volatility measure in equation (30) to approximate the contemporaneous quadratic yield variation process $QV_{y_\tau}(t + h, h)$ on the left-hand-side of (13).

Our analysis focuses on the volatility of bond *yields*. There are only few realized volatility studies of U.S. Treasury securities such as, e.g., Andersen et al. (2006b), and these invariably rely on the realized volatility of bond *returns*. As such, it is useful to clarify the link between equation (30) and the realized volatility of the bond return. To this end, we denote the continuously compounded return over the time interval $[t, t + h]$ on a zero-coupon bond with time-to-maturity $\tau = T - t - h$ by

$$r(t + h, h, \tau) = p(t + h, \tau) - p(t, \tau), \quad 0 \leq t \leq t + h \leq T, \quad (31)$$

where $p(t, \tau) \equiv \log(P(t, \tau))$ is the time- t zero-coupon log bond price. Since $P(t, \tau) = e^{-\tau y_\tau(t)}$ this implies the following expression for the intra-day return during a given trading day t , where τ is, by industry convention, constant:

$$r_\tau \left(t + \frac{ih}{n}, \frac{h}{n} \right) = -\tau \left(y_\tau \left(t + \frac{ih}{n} \right) - y_\tau \left(t + \frac{(i-1)h}{n} \right) \right). \quad (32)$$

Thus, the sum of the squared intra-day changes in yields is proportional to the sum of the intra-day squared returns, with a constant of proportionality of τ^2 , i.e., the square of the time-to-maturity. It follows that the realized volatility of the *return* on a bond with maturity τ during $[t, t + h]$ is

$$v_{r_\tau}(t + h, h; n)^2 = \sum_{i=1, \dots, n} \tau^2 \left(y_\tau \left(t + \frac{ih}{n} \right) - y_\tau \left(t + \frac{(i-1)h}{n} \right) \right)^2. \quad (33)$$

Hence, the qualitative features of the realized yield volatility and realized return volatility series are identical for a given maturity zero-coupon bond, and one may be readily derived from the other. Nonetheless, in order to match the yield volatility implications from the affine model class to the cross-section of bond yields, it is necessary to express the estimated quadratic variation in units of yield volatility. Equation (33) then renders comparisons to corresponding findings in the literature for a given maturity bond expressed in terms of realized return volatility straightforward.

3 U.S. Treasury Data

3.1 Intra-Day Yield Data

We rely on the GovPX database to construct intra-day series of bond yields. GovPX consolidates and posts real-time quote and trade data from most of the major interdealer Treasury securities brokers with a notable exception being Cantor Fitzgerald Inc. Taken together, these brokers account for about two-thirds of the interdealer broker market, a fraction that declined to 42% in the first quarter of 2000. In turn, the interdealer market is approximately one half of the total market (see Fleming (1997, 2003)). Hence, while the estimated bills coverage exceeds 90% in every year of the GovPX sample, the availability of thirty-year bond data is limited because of the prominence of Cantor Fitzgerald at the long-maturity segment of the market. Therefore, we use only data on the three-month, six-month, and one-year bills, as well as the two-, five-, and ten-year notes in our analysis. We rely exclusively on quotes for the on-the-run contracts, which are significantly more liquid than off-the-run Treasuries.⁵

Our sample period starts at the inception date of GovPX, June 17, 1991, and ends on June, 15, 2001. More recent data are also available, but we purposely avoid using them for several reasons. First, the 1-year Treasury Bill was no longer auctioned beginning March 2001. Second, after the end of our sample period the GovPX coverage of the U.S. Treasury market started to decline (see, e.g., Fleming (2003)). Third, the period following September 11, 2001, terrorist attacks has been tumultuous for bond markets (see, e.g., Fleming and Garbade (2002)).

The U.S. Treasury market is most active during business days from early morning through the late afternoon. Thus, we start the intra-day transaction record at 7:30AM ET and we end it at 5:00PM ET. This window includes the regular macroeconomic and monetary policy announcements,⁶ which are among the most important determinants of yield changes (see, e.g., Andersen et al. (2003b), Balduzzi et al. (2001), Fleming and Remolona (1997, 1999), Green (2004), and Li and Engle (1998)). Moreover, since the vast majority of the trading takes place during these hours, we capture the associated price discovery process, often seen to reflect the aggregation of heterogeneous private information as well as heterogeneous interpretation of public information, see, e.g., Brandt and Kavajecz (2004) and Pasquariello and Vega (2005).

The GovPX quote frequency for specific maturities turn out to be quite a bit lower than for, e.g., stocks in the Dow Jones 30 index. As such, the recent literature on selecting an optimal intra-day

⁵Fleming (2003) points out that the GovPX raw data files need to be cleaned due to some interdealer brokers' posting errors that are not filtered out by GovPX. Hence, prior to our analysis we implement the error correction procedures recommended by Fleming in the appendix of his paper.

⁶Most scheduled macroeconomic announcements take place at 8:30AM or 10AM. The statements from the Federal Open Market Committee meetings are typically released around 2:15PM. Further, the window includes the Federal Reserve's customary intervention times (11:30AM before 1997, 10:30AM from 1997 to 1999, and 9:30AM from 1999) and the Treasury auctions announcement times (1:30-2PM)

sampling frequency for computing the quadratic variance process in the presence of market microstructure noise, e.g., Aït-Sahalia et al. (2005a,b), Barndorff-Nielsen et al. (2004), Hansen and Lunde (2005), and Bandi and Russell (2005a,b) is not directly applicable. We instead follow the practice of the earlier literature of using a fairly sparse and fixed sampling frequency. A sensible compromise between adding information regarding the strength of the yield movements and incorporating high-frequency microstructure noise seems to be attained around the 10-minute sampling interval where the induced serial correlation in the yield change series are minor. Hence, at the end of each 10-minute interval, we use the immediately preceding on-the-run quote to construct the relevant bid and ask prices. We define the log-price, $\log(P(t))$, as the mid-point of the logarithmic bid and ask. We convert bond prices into zero-coupon yields by using the procedure described in Appendix A. Finally, we compute the series of intra-day yield changes for each Treasury in our sample.

In the sample period, we find a small number of days during which the trading activity is very subdued. We discard those days for which we could not find any trading activity for a period longer than three hours.⁷ This approach delivers a series of 56 intra-day 10-minute yield changes over 2,322 business days, for a total of 130,032 observations for each of the six Treasuries in our sample, as explained in more detail in Appendix B. Further, we compute an average of the daily trading period yield from the 57 intra-day observations so that any i.i.d.-type measurement error is negligible.

Features such as price discreteness and bid-ask spread positioning due to dealer inventory control are among the market microstructure frictions that may induce negative autocorrelation in the recorded series. In order to mitigate the impact of such institutionally driven short-term bouncing in the prices we finally apply an MA(1) filter to the yield change series.⁸

3.2 Daily Constant-Maturity Yield Data

As previously mentioned, the GovPX coverage of the thirty-year bond is limited and we excluded this security from our analysis. Data on the thirty-year bond at the *daily* frequency is, however, available from other sources.⁹ Although such information is not useful for the computation of intraday-based realized volatility series, it can still be used to construct a proxy for a zero-coupon yield and serve as a regressor in the volatility spanning condition (13). Consequently, such auxiliary daily yields may be

⁷Although there are no fixed trading hours for Treasuries, the Bond Market Association (BMA) makes recommendations regarding holiday closures and early closing times. Specifically, the BMA often recommends that the market close early (usually at 2 PM) before a holiday, which typically results in low trading activity for those days. As a robustness check, we also considered eliminating days during which we could not find any trading activity for a period longer than either two or four hours. Neither approach changes the conclusions discussed below.

⁸The MA coefficient estimates are as follows. For the three-month series: -0.064; six-month series: -0.062; one-year series: -0.037; two-year series: -0.007; five-year series: -0.003; ten-year series: -0.048. As a robustness check, we also applied the realized volatility estimator proposed in Hansen and Lunde (2005), which is designed to accommodate the effects of market-microstructure noise. This alternative approach produces results similar to those reported below.

⁹The Treasury last auctioned a nominal thirty-year bond in August 2001, i.e., after the end of our sample period.

used to provide an additional robustness check for our results based on the GovPX quotes.

We therefore consider a panel of daily yields from a constant-maturity series released by the Federal Reserve Board of Governors. In this case, we focus on maturities of three and six months, one, two, three, five, seven, ten, and thirty years.¹⁰ These constant maturity series contain theoretical yields for coupon bonds sold at par. Hence, prior to analysis we convert these series into zero-coupon yields, as discussed in Appendix B.

3.3 Yield Volatility Measures

The quality of the empirical analysis of the affine yield spanning conditions depends critically on the reliability of the realized volatility measures. Since the true (realized) volatility is latent, it is not a trivial matter to confirm the reliability of volatility measures. Thus, before embarking on an extensive analysis of the fundamental spanning restrictions, we document the salient features of our volatility estimates in some detail. One objective is to assess the general properties of the series relative to prior findings in the literature. A second objective is to clarify the difference between our ex-post realized yield volatility measures and the more common (ex-ante) volatility forecasts. This should alleviate any concern that our results are driven by some idiosyncratic features of our realized volatility series. We pay particular attention to the degree of predictability of yield volatility as this is of particular interest for the subsequent empirical evaluation of the affine model class.

3.3.1 Realized Yield Volatility

We construct the realized yield volatility series in equation (30) from intra-day quote data on the three-month, six-month, and one-year bill, as well as the two-, five-, and ten-year note. These realized volatility estimates constitute measures of the zero-coupon yield quadratic variation during business hours (7:30AM to 5PM E.S.T.) alone. In order to relate them to comparable studies based on daily or lower frequency data we express them in units reflecting a yearly percentage. However, we must also convert them from trading-day (business hours) to actual calendar-time yield volatility measures.

The main issue is how to account for the yield volatility outside of business hours. For each maturity series, we compute the inter-daily yield change from the close at 5PM on each trading day to the open at 7:30AM on the next trading day. We then measure ‘total’ daily realized volatility measure by adding the squared inter-daily yield change to the sum of squared intra-day yield changes. This total daily realized volatility measure is then used to compute the sample mean of the annual realized yield variation. Finally, we rescale each intraday-based realized volatility measure proportionally so that, across the entire sample, the daily annualized realized volatility measure based only on intraday yield changes matches that estimated from the total realized volatility series. We emphasize that any

¹⁰Specifically, we use the tcm3m, tcm6m, tcm1y, tcm2y, tcm3y, tcm5y, tcm7y, tcm10y, and tcm30y series from the web site <http://federalreserve.gov/releases/h15/data.htm>.

arbitrariness in the choice of this scaling constant is inconsequential for our empirical analysis of the fundamental spanning restriction.

Figure 1 depicts the square root of the rescaled daily realized volatility series, $v_{y\tau}$, for the zero-coupon yield series we analyze. Casual inspection shows a great deal of covariation in the yield volatilities across the maturity spectrum. However, there are some striking differences, particularly among the more extreme outliers. Although these are manifest in all the series, they often differ distinctly across the maturities. For example, there is a pronounced spike in the realized volatility of short-maturity yields on October 8, 1998, which is much attenuated for the longer yields. On that date, investors appear to react negatively to the concerns of a slowdown in the world economy, a weakening dollar (it lost 12% against the Japanese Yen in less than two days), and the initiation of an impeachment inquiry against President Bill Clinton. Investors were increasingly attracted to short-term bonds, pushing their prices up in a volatile trading session.¹¹ In contrast, the longer maturity yields' volatility spikes up on June 2, 1995, when a dramatic drop in the payroll employment number seems to have raised fears of a recession, sparking a powerful rally that sent prices of longer-term bonds to record highs.¹² Similarly, on March 8, 1996, the employment report revealed that over 700,000 new jobs were added to the payroll, lowering the unemployment rate from 5.8% to 5.5%. This event was contrary to the general perception that the economy was bordering on recession and likely reversed expectations that the Federal Reserve might plan to cut interest rates.¹³

These informal accounts are consistent with the widespread finding that macroeconomic announcement effects are prevalent in Treasury securities, often inducing a jump in the yields and an associated burst of volatility, see, e.g., Johannes (2004) and Andersen et al. (2006b). It also suggests that the reaction across the maturity spectrum is a function of the news content and the prevalent economic conditions. In particular, the volatility response is highly correlated for nearby maturities as one would expect if the economic effects were deemed stronger either at the shorter, medium, or longer-term maturities. The variance of the measurement error for daily realized volatility increases with the level of the underlying volatility, see, e.g., Barndorff-Nielsen and Shephard (2002b). Hence, the coherent response across nearby maturities during extreme events suggests that the realized volatilities capture the relative size of the effects adequately in spite of potential measurement problems.

¹¹“There has been a dramatic steepening of the yield curve,” said Michael Boss, a bond trader with IBJ Lanston Futures in Chicago to CNN, referring to growing differential between yields on long-term and short-term Treasury issues. “Fear is the overriding factor here—it has just been really ugly.”

¹²“Economists fear that downward momentum could feed on itself,” wrote Christopher Georges in the June 5, 1995, issue of the Wall Street Journal. “Once it gets going, the downward spiral is hard to stop,” said Sung Won Sohn, chief economist at NorWest Corp. in Minneapolis. “The correction could go on for longer than anticipated.”

¹³“A paradigm shift on Wall Street today. It all started with a report showing that a stunning number of people got jobs last month, the fastest improvement in the employment picture in 12 years. That started the “stronger economy” neon sign flashing, and gone went hopes of lower interest rates. The bond market fell off a cliff... with the 30 year treasury falling a heart-stopping 3 points,” reported David Brancaccio at CNN.

In Figure 2 we plot the average (sample) daily realized volatility for the zero-coupon yields for the different maturities. For comparison, we also include an alternative measure of yield volatility, obtained by computing the standard deviation of the daily changes (from 3pm to 3pm) in the zero-coupon yields. These graphs provide informal estimates of the ‘unconditional’ term structure of volatility. Both plots exhibit the characteristic ‘snake’ shape documented in, e.g., Piazzesi (2003, 2005). Since our analysis is limited to yields with maturities of at least three months our volatility measures are largely unaffected by short-lived deviations of the short rate from the target zone which can push the ‘head’ of the snake further up. Moreover, the hump in the ‘back’ of the snake appears less pronounced than what has been reported in some studies, including Dai and Singleton (2000, 2003). Such discrepancies may arise from varying degrees of policy inertia across the sample periods covered by the studies, as suggested by Piazzesi (2003). Overall, our realized volatility series replicates the qualitative features of prior studies along this dimension nicely. Moreover, the correspondence between the graph constructed from the intraday-yield-based measures and the daily-yield-based measures adds further credence to the reliability of the realized volatility series.

A final critical feature of the volatility yield series is the degree of temporal dependency. This feature may be gauged from the sample correlogram for the daily logarithmic realized volatility in Figure 3. The displays reveal a distinct hyperbolically declining pattern that is readily, albeit informally, assessed in quantitative terms by reference to the superimposed fitted hyperbolic curves. The findings are again remarkably similar across the maturities. In terms of the standard coefficient for fractional integration, d ($0 < d < 1$), the fitted curves (going from the shorter to longer maturities) imply values of 0.34, 0.33, 0.36, 0.33, 0.31 and 0.30. This evidence is suggestive of the presence of long-memory-type persistence in the volatility process for each of these series, consistent with previous evidence from the analysis of realized volatilities on equities (e.g., Andersen and Bollerslev (1997a) and Andersen, Bollerslev, Diebold, and Ebens (2001)) and currencies (e.g., Andersen (2000), Andersen and Bollerslev (1997a,b, 1998b)).

In conclusion, we find that the realized yield volatility series are consistent with prior evidence in terms of the overall (unconditional) yield volatility level across the term structure, that they appear to capture the volatility bursts associated with the release of macroeconomic announcements in a credible fashion, and that they display the type of temporal dependencies that may be anticipated from existing evidence for equity and foreign exchange markets as well as fixed income markets. These findings, combined with the consistent features observed across nearby maturities, suggest that the series are highly informative regarding the underlying true yield volatility realizations. There is no indication that the measurement errors associated with the use of ten-minute yield changes or the impact of microstructure noise has seriously impaired the quality of these volatility proxies.

3.3.2 Daily EGARCH Volatility Forecasts

Although the realized volatility series generally appear credible and consistent with evidence from daily yield series along various dimensions, the time series behavior is critical for our subsequent volatility spanning regressions. Thus, we now also directly explore the coherence between standard yield volatility forecasts obtained from daily data and the realized yield volatility measures based on intraday data. A thorough exploration of this relation across the entire yield curve would require a full-length independent study so we focus on the daily three-month Treasury bill yields only. Since this series has been analyzed in many prior studies the findings serve both as a robustness check for the (ex-post) realized yield volatility measures and as a reference point for comparison to the existing literature. A number of subtle points are involved in the specification and selection of the appropriate model. Since these are largely unrelated to the main issues of this paper we suppress the details to Appendix C. Later, in Section 4, when we need some candidate yield volatility forecasts we turn instead to simpler models that exploit the superior information in the history of observed intraday-based realized yield volatility series directly.

We obtain the daily yield volatility estimates from an ARMA-EGARCH-type model for the short-term zero-coupon yields. The ARMA structure allows for a flexible conditional mean process while the literature has documented that an EGARCH representation for the conditional yield volatility provides a convenient and successful parsimonious model for the conditional heteroskedasticity in the series. In addition, we introduce an additional source of dynamic interaction via an interest level effect, i.e., the EGARCH conditional variance term is scaled by the power of the yield level, as is common in the preceding time series studies. Finally, we stipulate that the conditional innovations belong to the family of SNP densities introduced by Gallant and Nychka (1987), so that the conditional density may accommodate remaining non-normality and time-series dependence in the residuals. The former consideration is particularly important given the evidence of jump-like outliers in the daily yield series. Detail regarding the interpretation and implementation of these flexible parametric models and their fit to the daily short maturity yield series may be found in Andersen and Lund (1997) and Andersen, Benzoni and Lund (2002, 2004).

The model is applied to the three-month Treasury bill yield series.¹⁴ A long sample is required to identify the dynamics of a highly persistent series so we use data covering July 1, 1983, to June 30, 2005, for a total of 5,498 observations. Earlier data are available, but we exclude the period involving the FED's monetary experiment as this arguably represents a regime shift. Table 1 provides summary statistics for the daily yield series. The first two columns correspond to the full sample used in estimation of the SNP density outlined in Appendix C, while the last two columns represent the

¹⁴We use the H.15 series of daily three-month T-Bill bank discount rates from the St. Louis FED web site, <http://research.stlouisfed.org/fred2/>. Prior to analysis, we convert the H.15 bank-discount-rate data into continuously compounded yields.

June 17, 1991—June 15, 2001, sample period of our realized volatility series. The basic summary statistics are comparable across the two sample periods. In particular, the interest rate levels are close and the kurtosis of the yield changes is exceedingly high for both samples. However, it is apparent that the yield volatility is relatively lower over 1991-2001.

We estimate the model by (quasi-)maximum likelihood (QML). The Bayesian (BIC) and Hannan-Quinn (H-Q) information criteria are used to guide model selection while additional information regarding the choice of the ARMA and EGARCH terms is gleaned from Ljung-Box tests for the autocorrelation of the raw and squared residuals. This analysis leads us to a so-called ARMA(6,1)-Level-EGARCH(2,1)-Kz(8)-Kx(0) specification.¹⁵ The representation implies that the conditional heteroskedasticity is captured by the EGARCH-Level representation while the Kz(8) term accommodates strong departures from conditional normality of the standardized daily yield innovations.

The high orders of the ARMA and EGARCH terms imply a persistent and complex dynamics for the conditional mean and volatility of the yield series. Further, the parameter estimates point towards extremely persistent first- and second-order conditional moments, even if the basic stationarity conditions are satisfied. Specifically, the inverse of the dominant root for the conditional mean polynomial is 0.9998, while the inverse of the roots for the conditional variance polynomial are 0.9970 and 0.9385. Obviously, the mean dynamics is hard to distinguish from the unit root case. The volatility roots fall within the range that produces high-order autocorrelations consistent with the long-memory-type persistence documented in Figure 3 for the realized volatility series.

Figure 4 depicts the EGARCH-Level-SNP model implied one-day-ahead volatility forecasts along with the corresponding daily realized volatility series, $v_y(t)$. This is literally a plot of daily volatility forecasts versus subsequent realizations, in which we use $v_y(t)$ as a proxy for the ex-post volatility realizations in the spirit of equation (16). The smoothing associated with the formation of ex-ante expectations within the EGARCH-Level model is readily apparent in the contrast to the jagged nature of the realized volatility series. Moreover, it is evident that the extreme positive outliers in the realized volatility series, almost by definition, are not a priori predictable. Nonetheless, there is a good qualitative coherence between the two series constructed from distinct data sources, as the long-run movements in the yield volatility forecasts are clearly related to corresponding shifts in the overall intensity of the realized volatility measures. The correlation between the two series is about 44%. Moreover, the overall degree of explanatory power of the forecasts for the variability of the future realized volatility is only slightly lower than reported for futures on the 30-year U.S. Treasury yield in Andersen et al. (2006b). It should be kept in mind, as documented in the latter study, that the yield volatility predictability is considerably lower for the fixed-income markets than for return volatility in, e.g., equity and currency markets. Finally, we also note that the sample means of the two series are close at 0.0448% and 0.0471% per day for the SNP forecast and the realized volatility, respectively.

¹⁵Estimation results are available from the authors upon request.

In sum, our short-term realized yield volatility series is coherent with the volatility forecasts obtained from the extended EGARCH model estimated at the daily frequency. The relation between the two series is analogous to what has been found in prior empirical work and it is consistent with the theoretical results in Andersen, Bollerslev, and Meddahi (2004, 2005). Hence, the EGARCH-Level forecasts appear well calibrated to the realized yield volatility series constructed from an entirely distinct intraday yield series, in turn suggesting that our daily realized yield volatility series provide relatively accurate and useful measures of the underlying quadratic yield variation realizations.

4 Evidence on the Affine Spanning Conditions

This section presents our main empirical findings. First, in Section 4.1 we focus on the empirical implications of multi-factor affine diffusion models previously summarized in Sections 2.1-2.4. Then, in Section 4.2 we extend our empirical analysis to a jump-diffusion setting, as discussed in Section 2.5.2.

4.1 Affine Diffusion Models

4.1.1 Do Bonds Span Volatility Risk?

The fundamental yield spanning condition in equation (13) should be satisfied by any affine (and quadratic) diffusion model that does not embed the unspanned stochastic volatility restriction. In order to assess this prediction empirically, we consider a regression model in which the dependent variable is the realized yield volatility. To construct a proxy for the independent variables, we compute the daily average of the intra-day zero-coupon yields series with three-month, six-month, one-, two-, five-, and ten-year maturity. This approach has the advantage of reducing possible problems due to the presence of non-systematic measurement errors in the zero-coupon yields extracted at different points across the trading day.

In order to allow for more ready interpretation and alleviate the multi-collinearity problem, we extract orthogonal principal components from the panel of average daily yields. As usual, the first three components correlate very highly with empirical measures of the level, the slope, and the curvature of the yield curve. Although these three components capture almost all of the yields' variation over time, we include all six in our regressions to make sure that our right-hand-side variables capture the entire variation that is contained in the original yield series. This approach also provides some robustness by including yields that may be relevant for the quadratic model class which can involve a large number of affine yield factors. We denote the principal components, or factors, by PC_j , $j = 1, \dots, 6$. We use ordinary least squares (OLS) to estimate the empirical version of equation (20) via the regression

model

$$v_{y_\tau}^2(t, h) = \beta_0 + \sum_{j=1}^6 \beta_j PC_j(t, h) + \varepsilon(t). \quad (34)$$

In this section, we report results based on daily realized volatility measures, i.e., $h = 1$. As a robustness check, we also discuss results for volatility measures at the weekly ($h = 5$) or monthly ($h = 21$) frequency in Section 4.1.2.

Panel A of Table 2 reports results based on the full sample. For concreteness, we initially focus on the regression with the shortest maturity (three-month) realized yield volatility, i.e., $v_{y_\tau}^2$, $\tau = 3M$, as the dependent variable. Remarkably, the R^2 coefficient of this regression is less than 4%. In theory, and ignoring measurement error, this regression should virtually have one-hundred per cent explanatory power. The result therefore suggests that a portfolio of bonds has very limited power to span (hedge) volatility risk, in stark contrast to the predictions from the entire affine diffusive class of models. Furthermore, it is noteworthy that the coefficients on the first three principal components are all insignificant, while a few of the higher-order components appear to enter significantly in the regression.¹⁶ Hence, whatever explanatory power the yield cross-section possesses, it is not represented by the usual factors identified in affine term structure models. In other words, the findings are even more problematic for three- or four-factor affine diffusive models which are common in the recent literature. The evidence is also at odds with, e.g., Litterman, Scheinkman, and Weiss (1991), who conclude that interest rate volatility is linked to the curvature factor.

The results from estimating equation (34) with the realized volatility of longer-maturity yields as the dependent variable are reported in rows 2-6 of Table 2, Panel A. The explanatory power is even weaker for the longer yields. In particular, for maturities beyond two years, the R^2 coefficients are essentially zero. Hence, the failings of the affine diffusive models are not linked to idiosyncratic features at the short end of the maturity spectrum. Since the spanning condition applies to the entire affine (non-USV) diffusive model class, we conclude that there is no potential for *any* model of this type to fit the observed yield volatility structure.

We also estimate the covariance matrix of the residuals from the six regressions reported in Panel A and perform a principal component analysis. This is in the spirit of Collin-Dufresne and Goldstein (2002) who perform a similar analysis on the residuals from regressions of returns on at-the-money straddles (i.e., portfolios mainly exposed to volatility risk) against changes in swap rates. As emphasized previously, however, our test procedure should be more powerful. Nonetheless, consistent with their results, we find that the first principal component explains more than 80% of the variation in the model residuals (Panel B). This evidence confirms that the limited explanatory power reported

¹⁶We report Newey-West standard errors that are robust to the presence of heteroskedasticity and serial correlation in the regression residuals. We use 20 lags in the computation of standard errors in all regressions that rely on daily data, 30 lags for the overlapping weekly data, and 40 lags for the overlapping monthly data. We have confirmed that the point-estimates of the t -ratios are robust to the choice of a higher number of lags.

in Panel A is not due to noisy data. In contrast, there is a dominant common factor that drives the volatility of zero-coupon yields at different maturities and is unrelated to the level of the yields. However, the volatility yield innovations are clearly not perfectly correlated across the maturity spectrum so multiple stochastic volatility factors appear necessary to accommodate these features.

4.1.2 Robustness of the Results

As a robustness check, we estimate the spanning regression for many different sub-samples and over longer horizons. For each sub-sample, we fit six separate regressions corresponding to each of the different maturity zero-coupon yield volatility series that are available. However, to conserve space, Table 3 reports only findings for the three-month, two-year, and ten-year realized volatility series which are representative of the full set of results.

Panel A contains estimation results for the period considered in Bikbov and Chernov (2004), May 1, 1994–June 27, 2001. Starting with the first Federal Open Market Committee (FOMC) meeting of 1994, the Fed has announced the new target overnight rate in the federal funds market at the end of each meeting. The Fed also changed the size and timing of target moves. Therefore, this period may be interpreted as a single monetary policy regime. Again, there is no indication that interest rate volatility can be extracted from a panel of bond yields. Further, we again perform a principal component analysis of the regression residuals. Panel B confirms that a single dominant factor explains most of the variation in our realized volatility series that is not accounted for by changes in the yield curve. Virtually identical results were obtained for the period, 02/01/95 to 12/29/2000, which was analyzed by Collin-Dufresne and Goldstein (2002).

Panels C and D report results separately for the first and second parts of our sample. The explanatory power of the regressions with the short-term realized volatility as a dependent variable improves slightly compared to the results for the full sample. However, bond yields still have virtually no explanatory power towards the realized volatility of longer-maturity rates. Further, while the loading on the first principal component now generally is significant, it is positive during the first sample and negative during the second. Such abrupt sign changes suggest that the improvement in the regressions R^2 reflects in-sample over-fitting.¹⁷ An alternative interpretation is that there is a change in regime during the sample period but additional analysis does not support this conjecture.¹⁸

¹⁷We confirmed that such sign changes are not due to a rotation of the yields' principal components. That is, in each sub-sample the first three yields' principal components correlate positively with the level, slope, and curvature of the term structure.

¹⁸Generally, in regime-switching affine models, shorter sample periods should correspond more closely to the presence of one dominant (persistent) regime, which would result in a better in-sample fit. In order to informally test for this possibility, we split the full sample into five distinct two-year periods and ran the regressions for each sub-sample. Although the explanatory power was considerably higher for one of these samples it was still extremely low in most cases. Moreover, the signs of the coefficients associated with different principal components were highly unstable across the sub-samples, indicating problems with overfitting. In summary, the spanning condition seems to fail almost uniformly.

The analysis in Section 3.3 suggests that our realized volatility series provide quite accurate measurements of the true underlying yield variation. Nevertheless, there may be residual concern regarding the impact of measurement errors in the variables entering the spanning regressions. A simple robustness check involves aggregation to a lower sampling frequency, since measurement errors tend to cancel and thus be less important relative to the accumulated signals in the yield changes and variation measures over longer horizons. Consequently, we explore the spanning condition (13) at the monthly frequency. For that purpose, the daily realized volatility measures are aggregated into overlapping monthly series. Similarly, we construct average monthly yields and extract the corresponding principal components. These series are then used in estimating equation (34) by OLS. The results, in Table 4, Panel A, are in line with our previous findings. The increase in overall explanatory power is consistent with the substantially lower idiosyncratic yield variability at the monthly compared to the daily frequency. Nonetheless, the explanatory power remains very low for the longer maturity series. In addition, for the shortest maturity yield volatility series, which is best explained by the yield cross-section, none of the first four principal components, typically associated with the factors in standard affine diffusive models, are significant. Finally, perhaps most tellingly, the residual variation of the yield volatilities continues to have a very strong common component as documented in Panel B of Table 4. Hence, the monthly yield variation also displays a strong covariation across the maturity spectrum that is unrelated to the yield levels.¹⁹

We conduct yet another check to verify that our results are robust to the set of proxies used for the bond yields on the right hand side of the regression model. Specifically, we consider a second panel of daily yields consisting of the constant-maturity series released by the Board of Governors, discussed in detail in Section 3.2. In this case, we have a larger number of maturities available, namely three and six month, and one, two, three, five, seven, ten, and thirty year treasuries. After converting the par coupon-yields into zero-coupon yields, we extract principal components from the series and we employ them in the regression (34), with the dependent variables being the usual six realized volatility measures constructed from intra-daily observations. The results, in Table 5, are consistent with those in Table 2. Most importantly, although the R^2 coefficients improve marginally, we still find little or no evidence that a portfolio of bonds can span (hedge) volatility risk. This reinforces our earlier conclusions. First, our previous findings are seemingly not hampered by the fact that we were not using the yields on the thirty-year bond. Second, the ineffectiveness of increasing the number of yields on the right-hand-side of our regressions suggests that the spanning condition is violated even in higher-order affine (or quadratic) models.

¹⁹All these findings apply for the weekly frequency as well. The detailed results are available from the authors upon request.

4.1.3 Cross-Maturity Correlations in Yield Volatility

We now empirically examine the restrictions that the popular $A_1(N)$ affine term structure model imposes on the quadratic variation of bond yields. These conditions are spelled out explicitly in equations (23) and (24). In order to be consistent with the preceding analysis which is cast in terms of the spanning properties of the yield *levels*, we focus on the first of these two restrictions but we have confirmed that the corresponding analysis based on the yield innovations produces similar results. Table 6 reports the findings. In contrast to the predictions of the $A_1(N)$ model, the correlations are considerably below one and display a systematic pattern across the maturities. For instance, the correlation between the three-month and ten-year yield volatility series is below 40%, and obviously very significantly different from unity. The correlations for nearby maturities are larger, but remain much lower than predicted by an affine model with a single state variable driving the volatility dynamics. These systematic deviations from the theoretical benchmark suggest the need for additional factors driving the term structure of yield volatility. These findings are robust to the analysis of different sub-samples (results available upon request).

These findings shed additional light on the reported tension between the time-series and cross-sectional properties of affine term structure models. This aspect of the problem has been identified by, e.g., Dai and Singleton (2000), who note that the presence of multiple square-root stochastic volatility factors improves the model ability to fit interest rate volatility. However, in the affine diffusive framework such an extension requires additional ‘admissibility’ conditions that force the unconditional correlations among the factors to be non-negative, a restriction that makes it much harder to accommodate the cross-sectional properties of bond yields.²⁰ Of course, the evidence in the preceding sections implies that a multi-factor extension with any potential of fitting the observed volatility yield structure must be cast in a framework that transcends the standard affine diffusion models. One recent approach is to formulate models which embed the USV restriction. These may in principle be extended to have multiple stochastic volatility factors—including a non-affine factor like that advocated by Ahn et al. (2003)—while retaining analytic tractability and without constraining the sign of the correlations among the latent variables. Irrespective of whether such models ultimately will be successful, the joint availability of realized yield volatility measures and fixed-income derivative prices seem to provide a promising basis for future research into these issues.

4.2 Yield Volatility Spanning in Affine Jump-Diffusive Models

As discussed in Section 2.5.2, equation (13) no longer applies in the presence of jumps. Instead, the appropriate yield spanning condition for affine jump-diffusion models is given by equation (29). This

²⁰Ahn et al. (2002, 2003) and Brandt and Chapman (2002) argue that three-factor quadratic-Gaussian models provide a better fit to interest rate volatility than previously considered specifications in the affine class (including the ‘essentially affine’ model of Duffee (2002)).

restriction is also valid, under weak auxiliary assumptions, for the conditional expectation evaluated under the P measure, and this is the version that we focus on here.

4.2.1 Empirical Volatility Spanning in Affine Jump Diffusions

We first consider the regression (22) with only the time t yield cross section (principal components) included as explanatory variables. Since the true conditional expectation of the future quadratic yield variation under the null hypothesis is an affine combination of these variables, the explanatory power of the regression should serve as an upper bound on the degree of predictability that is inherent in the yield volatility series. Table 7, Panel A, reports OLS results for

$$\frac{v_{y_\tau}(t+h, h)^2}{h} = \beta_0 + \sum_{j=1}^6 \beta_j PC_j(t) + \varepsilon(t+h), \quad (35)$$

where the left-hand-side variable is the daily ($h = 1$), overlapping weekly ($h = 5$), and overlapping monthly ($h = 21$) realized volatility for zero-coupon yields with maturity τ , and $PC_j(t)$, $j = 1, \dots, 6$, are the six principal components extracted from the sample of average past daily yields.²¹

It is evident that the yield cross section provides very little information regarding the future yield volatility. The explanatory power of the regressions is very limited at the one-day-ahead forecasting horizon ($h = 1$). At the weekly and monthly horizons the regression's R^2 improves, however the predictive power for the volatility of long-term yields remains negligible (at the 10-year maturity, the R^2 for the $h = 5$ and $h = 21$ cases is lower than 2% and 3%, respectively).

Moreover, the results are remarkably similar to those reported for the corresponding daily contemporaneous spanning regressions in equation (34), intended to mimic the underlying affine diffusive model restriction in equation (20), in Table 2. This finding suggests that the daily innovations to the yield curve are entirely unrelated to the daily yield volatility innovations.²² This is only consistent with the affine jump-diffusive setting if the large component of future yield volatility left unexplained by these regressions is, in a statistical sense, unpredictable. In order to assess the plausibility of this scenario, we turn to an alternative ad hoc approach to forecasts realized volatility.

4.2.2 Realized Volatility Component Model Forecasts

Unlike the spanning condition (13), equation (29) entails forward looking expectations rather than a realization-by-realization linkage. As a result, tests of the jump-diffusive variant of the spanning

²¹These tests may use either the zero-coupon yields extracted at the end-of-the-day (with the potential advantage of providing more current information regarding future yield changes) or the past daily average of the ten-minute yields (likely measured with less error). We report results based on the latter approach, but we have confirmed that near identical results obtain when using only yield observations sampled towards the end of the trading day (2-5PM). This finding provides some further assurance that idiosyncratic yield measurement errors are not contaminating our results.

²²Similar findings apply for the weekly and monthly forecast horizons and for brevity they are suppressed.

condition must concentrate on properties of the conditional expectation as an optimal predictor of future realizations, as emphasized in Section 2.3. Thus, it is convenient to have a simple alternative conditional yield volatility forecasts available.

To this end, we consider a forecasting method that relies on standard time-series models built directly on the history of ex-post volatilities. This approach provides a simple and powerful alternative to the more highly parametric EGARCH-style forecasts considered earlier. A convenient candidate model along these lines is the HAR-RV form studied in Andersen et al. (2006a), Corsi (2003), and Müller et al. (1997):

$$\frac{v_{y_\tau}(t+h, h)^2}{h} = \beta_0 + \beta_D v_{y_\tau}(t, 1)^2 + \beta_W \frac{v_{y_\tau}(t, 5)^2}{5} + \beta_M \frac{v_{y_\tau}(t, 21)^2}{21} + \varepsilon(t+h), \quad (36)$$

where the left-hand-side variable is the daily ($h = 1$), overlapping weekly ($h = 5$), and overlapping monthly ($h = 21$) realized volatility for zero-coupon yields with maturity τ , identical to Section 4.2.1. An important feature of the HAR-RV model is that the mixing of three volatility components allows for a slow volatility autocorrelation decay that is nearly indistinguishable from that of a hyperbolic pattern. This is consistent with the properties of our realized volatility series, as documented in Section 3.3.1. In addition, the forecast procedure is straightforward and fits well within the regression oriented set of diagnostic specification tests we have developed in the preceding sections.

In Table 7, Panel B, we report OLS estimation results for the HAR-RV model (36). We note that the estimates for β_D , β_W , and β_M confirm the existence of highly persistent volatility dependence. Moreover, the R^2 coefficients improve considerably with the forecasting horizon, consistent with the theory in Andersen, Bollerslev, and Meddahi (2004) and the evidence documented in, e.g., Andersen et al. (2003a). Intuitively, the extreme right skew in the realized volatility series induced by occasional volatility bursts renders it difficult to predict in the metric of explained ex-post sample variation, but this effect is less prominent at an intermediate horizon.

All in all, there is good coherence between the HAR-RV forecasts and the ex-post volatility realizations. This is illustrated in Figure 5, which depict the HAR-RV forecasts, $\hat{v}_{y_\tau}(t, h)$, $h = 1$, along with the corresponding realized volatility series. The long-run movements in the forecasts are clearly related to corresponding shifts in the overall intensity of the realized yield volatility measures. Further diagnostics (available from the authors upon request) lend additional support to the quality of fit and show, e.g., that the HAR-RV forecasts are more accurate than the EGARCH-type forecasts introduced in Section 3.3.2.

The HAR-RV regressions in Table 7, Panel B, provide a natural benchmark for the volatility forecasts in Section 4.2.1, which are constructed from the yields alone. Although the predictive power of the HAR-RV volatility components is limited, it is nonetheless significantly higher than that implied by the zero-coupon yields. For example, the R^2 for the one-day-ahead forecast at the extremes of the three-month and ten-year maturities are 4.6% and 0.5% in Panel A versus 8.1% and 2.2% in Panel

B, respectively. The improvement in predictive performance indicates that the ad hoc time series forecasts provide information over-and-above that embedded in the yields' cross-section.

A more direct test of this property is given by the encompassing regression in Panel C. Here, the set of explanatory variables includes the conditional HAR-RV yield volatility forecasts formed at time t as well as the principal components of the time- t zero-coupon yields,

$$\frac{v_{y_\tau}(t+h, h)^2}{h} = \beta_0 + \sum_{j=1}^6 \beta_j PC_j(t) + \beta_D v_{y_\tau}(t, 1)^2 + \beta_W \frac{v_{y_\tau}(t, 5)^2}{5} + \beta_M \frac{v_{y_\tau}(t, 21)^2}{21} + \varepsilon(t+h, h). \quad (37)$$

The three HAR-RV related regressors enter almost uniformly with positive and mostly statistically significant coefficients. In contrast, the coefficients on the yields (not reported) are generally insignificant. Hence, there is evidence of systematic predictability in yield volatility for all the maturities and across daily, weekly, and monthly forecast horizons which is captured by the time series forecasts but not by the yield cross section. In addition, it is noteworthy that the regression coefficients on the past realized volatility components invariably change very little with the addition of the yield principal components as regressors and the increase in the overall explanatory power of the regressions is very small. Overall, the results starkly contradict the null hypothesis that the conditional yield volatility may be expressed as an affine combination of the yields. As such, we conclude that there is compelling evidence against the yield volatility spanning conditions implied by any standard model within the affine class.

5 Conclusions

In recent years, the study of the term structure of interest rates has relied predominantly on continuous-time multi-factor models and, in particular, on affine specifications. A general implication of these models is that the (expected) quadratic yield variation for any fixed-maturity zero-coupon bond is spanned by the contemporaneous yield cross-section. That is, these models predict that interest rate volatility can be extracted from current bond prices.

We rely on model-free, yet efficient, yield volatility measures constructed from high-frequency intraday data to directly test the yield volatility implications for a very broad class of term structure models. Contrary to the affine model predictions we find that neither *realized* nor *expected future* quadratic yield variation is spanned by the term structure of zero-coupon yields. In fact, we find a pronounced and systematic covariation in yield volatility across the maturity spectrum that appears essentially unrelated to the state of the term structure.

Our results shed new light on related empirical findings in the literature. First, there is extensive evidence that short term interest rates display pronounced stochastic volatility features that are largely unrelated to the level of the short term interest rate itself (see, e.g., Brenner, Harjes, and Kroner (1996) and Andersen and Lund (1997)). It does seem natural, in theory, to relate the short rate volatility to the 'curvature' factor in the yield curve, as originally proposed by Litterman, Scheinkman, and Weiss

(1991). This specific linkage has subsequently been questioned by various studies, and we confirm that there is little empirical support for this specific connection. Indeed, our results imply more generally that there is, at best, a very weak relation between the short-rate volatility factor and the term structure. Moreover, any such link becomes virtually non-existent when we relate long maturity yield volatility to the yield curve. Hence, the model failings are even more glaring at the long end of the term structure.

Second, we also gain a fresh perspective on the recent conflicting evidence regarding the affine USV restriction. If we accept that standard affine models are seriously deficient in terms of their implications for the yield volatility dynamics, it is inherently difficult to interpret hypothesis tests concerning the USV restriction within a broader affine model setting. In fact, rejecting the USV hypothesis in favor of an encompassing affine model leads directly back to the empirical conundrum highlighted in the current paper. Moreover, given the additional evidence we have presented against the specification of affine models with a single factor driving the yield volatility, it is important to exert caution when interpreting the inference conducted in studies based on parametric representations that fall within this particular, but commonly adopted, affine framework (e.g., Bikbov and Chernov (2004), CDGJ, and Thompson (2004)).

Taken together, our findings suggest that further extensions to the affine term structure modeling framework are warranted. The literature provides a variety of interesting directions to pursue. One natural and popular approach involves the linkage of the underlying yield curve factors to macroeconomic variables. This may set the stage for a better understanding of the interaction between the term structure dynamics on one hand and monetary policy, the general economic environment, and inflationary expectations on the other.²³ A related strand of literature explores the yield curve reaction to the release of regularly scheduled macroeconomic announcements. These news releases contribute significantly to the observed total quadratic yield variation through both jumps and the associated short-lived volatility bursts in the fixed-income markets. Our realized volatility measures confirm this evidence. Further, they suggest that the reaction across the maturity spectrum is a function of the news content as well as the prevalent economic conditions. In particular, the volatility response is highly correlated for nearby maturities as one would expect if the economic effects vary in terms of their impact for the shorter, medium, or longer-term maturities. Hence, models with general linkages of the term structure of yields and volatilities to macroeconomic and monetary policy variables and associated time-varying reactions to scheduled announcements may be particularly promising.

At the more methodological level, we expect that the use of high-frequency intraday bond data for improved measurement of the real-time evolution in yield volatility will be informative regarding the specification of such future candidate term structure models.

²³Previous studies along these lines include Ang et al. (2005), Bikbov and Chernov (2005), Dai and Philippon (2004), Diebold, Piazzesi, and Rudebusch (2005), and Piazzesi (2005).

Appendix A: Computing Zero-Coupon Yields from Treasury Bond Prices

There is an array of methods to extract zero-coupon yields from a cross section of bond prices and a vast literature has studied the relative advantages of different approaches. The key feature underlying different methods is determined by the choice of a specific interpolation to fit the discount rate (or forward rate) curve. As noted by Dai, Singleton, and Yang (2006), on one extreme is the unsmoothed Fama-Bliss approach, which iteratively extracts forward rates from coupon-bond prices by building a piece-wise linear discount rate function. As such, the implied discount rates exhibit kinks at the maturities of the coupon bonds used. At the opposite extreme the Nelson-Siegel-Bliss and smoothed Fama-Bliss methods approximate the discount rates with exponential functions of time to maturity, which yields a smooth term-structure interpolation. In between these two extremes are methods that rely on a spline curve interpolation for the discount or forward rates.

Here, we use a spline-based method to extract the zero-coupon yields. More specifically, we fit the curve of zero-coupon yields with a cubic spline interpolation $\hat{y}(\tau, \Psi)$, where Ψ is the vector of spline coefficients, τ is time-to-maturity, $\tau_{min} \leq \tau \leq \tau_{max}$, with τ_{min} and τ_{max} the nearest and farthest maturities, respectively. We define P_{τ_k} , $k = 1, \dots, N$, to be the time- t market price of a bond with maturity τ_k ²⁴ and let \hat{P}_{τ_k} be the price of the same bond computed by discounting its coupon and principal payments at the discount rate \hat{y} . Then, we choose Ψ to solve the problem

$$\min_{\Psi} \left(\sum_{k=1}^N (P_{\tau_k} - \hat{P}_{\tau_k})^2 + \int_{\tau_{min}}^{\tau_{max}} \lambda(\tau) \hat{y}''(\tau, \Psi)^2 d\tau \right). \quad (38)$$

This is the approach preferred by Waggoner (1997) and used in, e.g., Dai, Singleton, and Yan (2006), except that we fit the smoothed cubic spline directly on the zero-coupon yields curve (similar to McCulloch (1975)), while Waggoner (1997) and Dai, Singleton, and Yan (2006) fit the smoothed spline on the forward rates curve. We follow Fisher, Nychka, and Zervos (1995) and we impose a penalty for the ‘roughness’ of the spline interpolation, i.e., we fit a ‘smoothed’ spline. As in Waggoner (1997) and Dai, Singleton, and Yan (2006) we calibrate the penalty function $\lambda(\tau)$ to be more severe at long maturities, which stiffens the spline interpolation and reduces its oscillations at the long end of the term structure. In contrast, we impose a small penalty for roughness at short maturities, where more flexibility is necessary to fit the prices of short-term securities, which are available in larger number than long-term bonds. Waggoner (1997) recommends to use a number of knots approximately equal to one third the number of bonds used in the construction of the spline. Dai, Singleton, and Yan (2006) use instead a larger number of knots, as many as 50 to 60. In our application, we take an approach that falls in between these two by using a number of knots approximately equal to the number of

²⁴For the GovPX data set, $N = 6$ and $\tau_k = 3$ and 6 months, 1, 2, 5, and 10 years. For the data set of daily constant maturity Treasuries $N = 9$ and $\tau_k = 3$ and 6 months, 1, 2, 3, 5, 7, 10, and 30 years.

bonds used in the estimation.

It is useful to illustrate the properties of the zero-coupon yields that we extract from bond prices through the spline interpolation. To this end, we construct three series of zero-coupon yields using on-the-run GovPX quotes at the ten-minute frequency for three-, six-month, and one-year Treasury Bills. These series do not rely on the spline term structure interpolation method. As such, depending on the Treasury’s issuing calendar there is a small mismatch in the yields’ maturities compared to the zero-coupon yields series that we use in our analysis. Further, we construct three series of *coupon* yields using on-the-run GovPX quotes at the ten-minute frequency for two-, five-, and ten-year Treasury Notes. Also in this case, these series do not rely on the the term structure interpolation procedure. Thus, in addition to being subject to a small maturity mismatch, these series differ from those used in our analysis because they contain coupon rather than zero-coupon yields.

We compute the correlation between each one of these six series and the corresponding series of zero-coupon yields constructed by solving problem (38). The results are in Table 8, Panel A. The sample correlations are nearly perfect at all maturities, which suggests that our spline interpolation method does not change dramatically the nature of the yield series that we investigate in the paper. We repeat this analysis for the data set of daily constant-maturity (TCM) par yields released by the Federal Reserve Board. Again, we find that the sample correlations between the series of TCM par yields and the series of zero-coupon yields extracted through our term structure interpolation are nearly perfect (Table 8, Panel B).

Next, we measure the sample correlations between realized volatility series constructed using the GovPX zero-coupon yield series and the series of ‘raw’ yields. Also in this case, for Treasury Bills the ‘raw’ yields are *zero-coupon* yields with maturities closest to three, six months, and one year. For Treasury Notes, the ‘raw’ yields are *coupon* yields with maturities closest to two, five, and ten years. These sample correlations are nearly perfect for maturities up to five years, and are approximately 97% at the ten-year horizon (Table 9).

Finally, we use the realized volatility series constructed from ‘raw’ GovPX yields to carry over the entire analysis that we perform in the paper with the zero-coupon yields. Consistent with the evidence in Table 9, the results are very close to those reported in the paper for zero-coupon yield series. For short-maturity yields, they are nearly identical.

Overall, this evidence suggests that the conclusions of our analysis are not biased by the noise due to the procedure that we use to extract the series of zero-coupon yields.

Appendix B: Quotes Updates and Sampling Scheme

In our analysis, at the end of each 10-minute interval we use the immediately preceding on-the-run quote to construct the relevant bid and ask prices. Due to infrequent trading, a quote may not be updated for more than 10 consecutive minutes. In these cases, our approach is to pad the last observed

quote until a new quote becomes available.

Such sampling scheme is economically meaningful. If there is no arrival of new information, brokers do not feel the need to update their quotes. As such, the previous quote still reflects the broker's assessment of the security's value. This is the price that we use in our analysis. In some cases quotes could be 'stale' and therefore misaligned with the securities' values. Eventually, however, such deviation will be corrected when the broker adjusts the quote. Such adjustment will be reflected in our realized volatility measures. When activity is very subdued this delay could create some bias in our volatility estimates. Thus, we discard those days for which we could not find any trading activity for a period longer than three hours.

In Table 10 below, we report the percentage number of unique yield observations in the intra-day GovPX dataset. The figures in Panel A are computed using the full sample period, from June 17, 1991 to June 15, 2001. In Panel B we report the same percentage ratios computed after discarding trading days with three or more hours of inactivity. In all cases, more than 80% of yields' observations are 'unique' quotes. This ratio increases to nearly 100% for the two-year notes (which is among the most liquid Treasury contracts) and is well in excess of 90% for most other maturities. As expected, discarding days with sparse trading activity results in an even higher percentage of unique quotes.

Appendix C: The SNP Conditional Density for the Daily Three-Month Yields

In summary form, the ARMA-EGARCH-Level-SNP model takes the general form,

$$y_t = \mu_t + y_{t-1}^\delta \sqrt{h_t} z_t \quad (39)$$

$$\mu_t = \phi_0 + \sum_{i=1}^s \phi_i y_{t-i} + \sum_{i=1}^u \zeta_i (y_{t-i} - \mu_{t-i}), \quad (40)$$

$$\ln h_t = \omega \left(1 - \sum_{i=1}^p \beta_i \right) + \sum_{i=1}^p \beta_i \ln h_{t-i} + (1 + \alpha_1 L + \dots + \alpha_q L^q) [\theta_1 z_{t-1} + \theta_2 (b(z_{t-1}) - \sqrt{2/\pi})]. \quad (41)$$

Finally, the conditional density for the yield innovation, z_t , is given by a flexible SNP representation. The full density takes the form,

$$f_K(y_t | x_t; \xi) = \left(\nu + (1 - \nu) \times \frac{[P_K(z_t, x_t)]^2}{\int_{\mathbb{R}} [P_K(z_t, x_t)]^2 \phi(u) du} \right) \frac{\phi(z_t)}{y_{t-1}^\delta \sqrt{h_t}}, \quad \nu = 0.01, \quad (42)$$

$$\begin{aligned}
z_t &= \frac{y_t - \mu_t}{y_{t-1}^\delta \sqrt{h_t}}, \\
\mu_t &= \phi_0 + \sum_{i=1}^s \phi_i y_{t-i} + \sum_{i=1}^u \zeta_i (y_{t-i} - \mu_{t-i}), \\
\ln h_t &= \omega \left(1 - \sum_{i=1}^p \beta_i\right) + \sum_{i=1}^p \beta_i \ln h_{t-i} + (1 + \alpha_1 L + \dots + \alpha_q L^q) [\theta_1 z_{t-1} + \theta_2 (b(z_{t-1}) - \sqrt{2/\pi})], \\
b(z) &= |z| \text{ for } |z| \geq \pi/2K, \quad b(z) = (\pi/2 - \cos(Kz))/K \text{ for } |z| < \pi/2K, \\
P_K(z, x) &= \sum_{i=0}^{K_z} a_i(x) z^i = \sum_{i=0}^{K_z} \left(\sum_{|j|=0}^{K_x} a_{ij} x^j \right) z^i, \quad a_{00} = 1,
\end{aligned}$$

where j is a multi-index vector, $x^j \equiv (x_1^{j_1}, \dots, x_M^{j_M})$, and $|j| \equiv \sum_{m=1}^M j_m$.

References

- Ahn, Dong-Hyun, Robert F. Dittmar, A. Ronald Gallant, and Bin Gao, 2003, Purebred or Hybrid?: Reproducing the Volatility in Term Structure Dynamics, *Journal of Econometrics* 116, 147-180.
- Ahn, Dong-Hyun, Robert F. Dittmar, and A. Ronald Gallant, 2002, Quadratic Term Structure Models: Theory and Evidence, *Review of Financial Studies* 15, 243-288.
- Aït-Sahalia, Yacine, Per A. Mykland, and Lan Zhang, 2005a, How Often to Sample a Continuous-Time Process in the Presence of Market Microstructure Noise, *Review of Financial Studies* 18, 351-416.
- Aït-Sahalia, Yacine, Per A. Mykland, and Lan Zhang, 2005b, Ultra High-Frequency Volatility Estimation with Dependent Microstructure Noise, NBER Working Paper #11380.
- Andersen, Torben G., 2000, Some Reflections on Analysis of High-Frequency Data, *Journal of Business and Economic Statistics* 18, 146-153.
- Andersen, Torben G., Luca Benzoni, and Jesper Lund, 2002, An Empirical Investigation of Continuous-Time Equity Return Models, *Journal of Finance* 57, 1239-1284.
- Andersen, Torben G., Luca Benzoni, and Jesper Lund, 2004, Stochastic Volatility, Mean Drift and Jumps in the Short Term Interest Rate, Working Paper, Northwestern University, University of Minnesota, and Nykredit Bank.
- Andersen, Torben G., and Tim Bollerslev, 1997a, Intraday Periodicity and Volatility Persistence in Financial Markets, *Journal of Empirical Finance* 4, 115-158.

- Andersen, Torben G., and Tim Bollerslev, 1997b, Heterogeneous Information Arrivals and Return Volatility Dynamics: Uncovering the Long-Run in High Frequency Returns, *Journal of Finance* 52, 975-1005.
- Andersen, Torben G., and Tim Bollerslev, 1998a, Answering the Skeptics: Yes, Standard Volatility Models do Provide Accurate Forecasts, *International Economic Review* 39, 885-905.
- Andersen, Torben G., and Tim Bollerslev, 1998b, Deutsche MarkDollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements, and Longer Run Dependencies, *Journal of Finance* 53, 219-265.
- Andersen, Torben G., Tim Bollerslev, and Francis X. Diebold, 2004, Parametric and Nonparametric Volatility Measurement, in Lars P. Hansen and Yacine Ait-Sahalia, eds., *Handbook of Financial Econometrics*, North-Holland, Amsterdam, forthcoming.
- Andersen, Torben G., Tim Bollerslev, and Francis X. Diebold, 2006a, Roughing It Up: Including Jump Components in the Measurement, Modeling and Forecasting of Return Volatility, *Review of Economics and Statistics*, forthcoming.
- Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Heiko Ebens, 2001, The Distribution of Realized Stock Return Volatility, *Journal of Financial Economics* 61, 43-76.
- Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Paul Labys, 2003a, Modeling and Forecasting Realized Volatility, *Econometrica* 71, 579-625.
- Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Clara Vega, 2003b, Micro Effects of Macro Announcements: Real-Time Price Discovery in Foreign Exchange, *American Economic Review* 93, 38-62.
- Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Clara Vega, 2006b, Real-Time Price Discovery in Stock, Bond and Foreign Exchange markets, Working Paper, Northwestern University, Duke University, University of Pennsylvania, and the Federal Reserve Board.
- Andersen, Torben G., Tim Bollerslev, and Nour Meddahi, 2004, Analytical Evaluation of Volatility Forecasts, *International Economic Review* 45, 1079-1110.
- Andersen, Torben G., Tim Bollerslev, and Nour Meddahi, 2005, Correcting the Errors: Volatility Forecast Evaluation Using High-Frequency Data and Realized Volatilities, *Econometrica* 73, 279-296.
- Andersen, Torben G., and Jesper Lund, 1997, Estimating Continuous-Time Stochastic Volatility Models of the Short-Term Interest Rate, *Journal of Econometrics* 77, 343-377.

- Ang, Andrew, Monika Piazzesi, and Min Wei, 2005, What Does the Yield Curve Tell us about GDP Growth?, *Journal of Econometrics*, forthcoming.
- Balduzzi, Pierluigi, Edwin J. Elton, and T. Clifton Green, 2001, Economic News and Bond Prices: Evidence from the U.S. Treasury Market, *Journal of Financial and Quantitative Analysis* 36, 523-543.
- Bandi, Federico M., and Jeffrey R. Russell, 2005a, Microstructure Noise, Realized Variance, and Optimal Sampling, Working paper, Chicago GSB.
- Bandi, Federico M., and Jeffrey R. Russell, 2005b, Separating Microstructure Noise from Volatility, *Journal of Financial Economics*, forthcoming.
- Barndorff-Nielsen, Ole E., and Neil Shephard, 2001, Non-Gaussian Ornstein-Uhlenbeck-Based Models and some of their Uses in Financial Economics (with discussion), *Journal of the Royal Statistical Society, Series B* 63, 167-241.
- Barndorff-Nielsen, Ole E., and Neil Shephard, 2002a, Econometric Analysis of Realized Volatility and its Use in Estimating Stochastic Volatility Models, *Journal of the Royal Statistical Society, Series B* 64, 253-280.
- Barndorff-Nielsen, Ole E., and Neil Shephard, 2002b, Estimating Quadratic Variation Using Realised Variance, *Journal of Applied Econometrics* 17, 457-477.
- Barndorff-Nielsen, Ole E., and Neil Shephard, 2004, Econometric Analysis of Realised Covariation: High Frequency Based Covariance, Regression and Correlation in Financial Economics, *Econometrica* 72, 885-925.
- Barndorff-Nielsen, Ole E., Peter R. Hansen, Asger Lunde, and Neil Shephard, 2004, Regular and Modified Kernel-Based Estimators of Integrated Variance: The Case with Independent Noise, Working Paper, University of Aarhus, Stanford University, and University of Oxford.
- Beber, Alessandro, and Michael Brandt, 2005, The Effect of Macroeconomic News on Beliefs and Preferences: Evidence from the Options Market, Working Paper, Duke University.
- Bikbov, Ruslan, and Mikhail Chernov, 2004, Term Structure and Volatility: Lessons from the Eurodollar Markets, Working Paper, Columbia GSB.
- Bikbov, Ruslan, and Mikhail Chernov, 2005, No-Arbitrage Macroeconomic Determinants of the Yield Curve, Working Paper, Columbia GSB.

- Bollerslev, Tim, Jun Cai, and F. M. Song, 2000, Intraday Periodicity, Long-Memory Volatility, and Macroeconomic Announcement Effects in the U.S. Treasury Bond Market, *Journal of Empirical Finance* 7, 37-55.
- Brandt, Michael W., and David A. Chapman, 2002, Comparing Multifactor Models of the Term Structure, Working Paper, Duke University and Boston College.
- Brandt, Michael W., and Kenneth A. Kavajecz, 2004, Price Discovery in the U.S. Treasury Market: The Impact of Order Flow and Liquidity on the Yield Curve, *Journal of Finance* 59, 2623-2654.
- Brenner, Robin J., Richard H. Harjes, and Kenneth F. Kroner, 1996, Another Look at Models of the Short-Term Interest Rate, *Journal of Financial and Quantitative Analysis* 31, 85-107.
- Casassus, Jamie, Pierre Collin-Dufresne, and Robert S. Goldstein, 2004, Stochastic (Unspanned) Volatility and Fixed Income Derivative Pricing, *Journal of Banking and Finance* forthcoming.
- Chapman, David A., John B. Long, Jr., and Neil D. Pearson, 1999, Using Proxies for the Short Rate: When Are Three Months Like an Instant? *Review of Financial Studies* 12, 763-806.
- Cheng, Peng, and Olivier Scaillet, 2005, Linear-Quadratic Jump-Diffusion Modeling with Application to Stochastic Volatility, Working Paper, HEC Geneve and FAME.
- Cheridito, Patrick, Damir Filipović, and Robert L. Kimmel, 2005, Market Price of Risk Specifications for Affine Models: Theory and Evidence, *Journal of Financial Economics*, forthcoming.
- Collin-Dufresne, Pierre, and Robert S. Goldstein, 2002, Do Bonds Span the Fixed Income Markets? Theory and Evidence for Unspanned Stochastic Volatility, *Journal of Finance* 57, 1685-1730.
- Collin-Dufresne, Pierre, Robert S. Goldstein, and Christopher S. Jones, 2004, Can Interest Rate Volatility be Extracted from the Cross Section of Bond Yields? An Investigation of Unspanned Stochastic Volatility, Working Paper, UC Berkeley, University of Minnesota, and University of Southern California.
- Comte, Fabienne, and Eric Renault, 1998, Long Memory in Continuous-Time Stochastic Volatility Models, *Mathematical Finance* , 291-323.
- Corsi, Fulvio, 2003, A Simple Long Memory Model of Realized Volatility, Working Paper, University of Southern Switzerland.
- Dai, Qiang, and Thomas Philippon, 2004, Fiscal Policy and the Term Structure of Interest Rates, Working Paper, University of North Carolina at Chapel Hill and NYU.
- Dai, Qiang, and Kenneth J. Singleton, 2000, Specification Analysis of Affine Term Structure Models, *Journal of Finance* 55, 1943-1978.

- Dai, Qiang, Kenneth J. Singleton, and Wei Yang, 2006, Regime Shifts in a Dynamic Term Structure Model of U.S. Treasury Bond Yields, *Review of Financial Studies*, forthcoming.
- Das, Sanjiv, 2002, The Surprise Element: Jumps in Interest Rates, *Journal of Econometrics* 106, 27-65.
- Diebold, Francis X., and Canlin Li, 2005, Forecasting the Term Structure of Government Bond Yields, *Journal of Econometrics*, forthcoming.
- Diebold, Francis X., Monika Piazzesi, and Glenn Rudebusch, 2005, Modeling Bond Yields in Finance and Macroeconomics, *American Economic Review P&P*, forthcoming.
- Duarte, Jefferson, 2004, Evaluating An Alternative Risk Preference in Affine Term Structure Models, *Review of Financial Studies* 17, 370-404.
- Duffee, Greg R., 1996, Idiosyncratic Variation of Treasury Bill Yields, *Journal of Finance* 51, 527-551.
- Duffee, Greg R., 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.
- Duffie, Darrell, Jun Pan, and Kenneth J. Singleton, 2000, Transform Analysis and Asset Pricing for Affine Jump Diffusions, *Econometrica* 68, 1343-1376.
- Fan, Rong, Anurag Gupta, and Peter Ritchken, 2003, Hedging in the Possible Presence of Unspanned Stochastic Volatility: Evidence from Swaption Markets, *Journal of Finance* 58, 2219-2248.
- Fisher, Mark, Douglas Nychka, and Daniel Zervos, 1995, Fitting the Term Structure of Interest Rates with Smoothing Splines, Working Paper, Federal Reserve Board.
- Fleming, Michael J., 1997, The Round-the-Clock Market of U.S. Treasury Securities, *Federal Reserve Bank of New York Economic Policy Review* 3, 9-32.
- Fleming, Michael J., 2003, Measuring Treasury Market Liquidity, *Federal Reserve Bank of New York Economic Policy Review* 9, 83-108.
- Fleming, Michael J., and Kenneth D. Garbade, 2002, When the Back Office Moved to the Front Burner: Settlement Fails in the Treasury Market after 9/11, *Economic Policy Review* 8, 35-57.
- Fleming, Michael J., and Eli M. Remolona, 1997, What Moves the Bond Market?, *Federal Reserve Bank of New York Economic Policy Review* 3, 31-50.

- Fleming, Michael J., and Eli M. Remolona, 1999, Price Formation and Liquidity in the U.S. Treasury Market: The Response to Public Information, *Journal of Finance* 54, 1901-1915.
- French, Kenneth R., G. William Schwert, and Robert F. Stambaugh, 1987, Expected Stock Returns and Volatility, *Journal of Financial Economics* 19, 3-29.
- French, Kenneth R., and Richard Roll, 1986, Stock Return Variances: The Arrival of Information and the Reaction of Traders, *Journal of Financial Economics* 17, 5-26.
- Gallant, A. Ronald, and Douglas W. Nychka, 1987, Semi-Nonparametric Maximum Likelihood Estimation, *Econometrica* 55, 363-390.
- Green, T. Clifton, 2004, Economic News and the Impact of Trading on Bond Prices, *Journal of Finance* 59, 1201-1234.
- Hansen, Peter R., and Asger Lunde, 2005, Realized Variance and Market Microstructure Noise, The 2005 Invited Address, *Journal of Business and Economic Statistics*, forthcoming.
- Hsieh, David A., 1991, Chaos and Nonlinear Dynamics: Application to Financial Markets, *Journal of Finance* 46, 1839-1877.
- Johannes, Michael, 2004, The Statistical and Economic Role of Jumps in Continuous-Time Interest Rate Models, *Journal of Finance* 59, 227-260.
- Joslin, Scott, 2005, Completeness in Fixed Income Markets, Working Paper, Stanford University.
- Kimmel, Robert L., 2004, Modeling the Term Structure of Interest Rates: A New Approach, *Journal of Financial Economics*, 72, 143183.
- Li, Haitao, and Feng Zhao, 2003, Unspanned Stochastic Volatility, Is It There After All? Evidence From Hedging Interest Rate Caps, Working Paper, Cornell University.
- Li, Li, and Robert F. Engle, 1998, Macroeconomic Announcements and Volatility of Treasury Futures, Working Paper, UC San Diego.
- Litterman, Robert and Jose A. Scheinkman, 1991, Common Factors Affecting Bond Returns, *Journal of Fixed Income* 1, June, 54-61.
- Litterman, Robert, Jose A. Scheinkman, and Laurence Weiss, 1991, Volatility and the Yield Curve, *Journal of Fixed Income* 1, June, 49-53.
- McCulloch, J. Huston, 1975, The Tax-Adjusted Yield Curve, *Journal of Finance* 30, 811-830.
- Meddahi, Nour, 2002, A Theoretical Comparison Between Integrated and Realized Volatility, *Journal of Applied Econometrics* 17, 475-508.

- Müller, Ulrich A., Michel M. Dacorogna, Rakhil D. Davé, Richard B. Olsen, Olivier V. Pictet, and Jacob E. von Weizsäcker, 1997, Volatilities of different time resolutions Analyzing the dynamics of market components, *Journal of Empirical Finance* 4, 213-239.
- Pasquariello, Paolo, and Clara Vega, 2004, Informed and Strategic Order Flow in the Bond Markets, Working Paper, University of Michigan and University of Rochester.
- Piazzesi, Monika, 2003, Affine Term Structure Models, *Handbook of Financial Econometrics*, forthcoming.
- Piazzesi, Monika, 2005, Bond Yields and the Federal Reserve, *Journal of Political Economy*, 113, 311-344.
- Poterba, James M., and Lawrence H. Summers, 1986, The Persistence of Volatility and Stock Market Fluctuations, *American Economic Review* 76, 1142-1151.
- Schwert, G. William, 1989, Why Does Stock Market Volatility Change Over Time?, *Journal of Finance* 44, 1115-1153.
- Schwert, G. William, 1990, Stock Volatility and the Crash of '87, *Review Financial Studies* 3, 77-102.
- Thompson, Samuel, 2004, Identifying Term Structure Volatility from the LIBOR-Swap Curve, Working Paper, Harvard University.
- Umantsev, Leonid A., 2001, Econometric Analysis of European LIBOR-Based Options within Affine Term-Structure Models, Ph.D. dissertation, Stanford University.
- Waggoner, Daniel F., 1997, Spline Methods for Extracting Interest Rate Curves from Coupon Bond Prices, Working Paper, Atlanta FED.

Tables and Figures

Table 1: Summary Statistics for the U.S. Three-Month Treasury Bill Yield Data. All figures are computed using daily yield data and expressed in percent per annum.

	07/01/1983–06/30/2005		06/17/1991–06/15/2001	
	y	Δy	y	Δy
Mean	5.0876	-0.0010	4.6441	-0.0009
Std. Dev.	2.3041	0.0568	0.9198	0.0467
Skewness	0.0475	-0.5392	-0.5512	-0.7656
Kurtosis	2.5134	15.6448	2.2427	19.8250

Table 2: Evidence on the Affine Diffusive Spanning Condition. OLS estimates for the regression $v_{y_\tau}^2(t, 1) = \beta_0 + \sum_{j=1}^6 \beta_j PC_j(t, 1) + \varepsilon(t)$. The dependent variable is the daily realized volatility (in percent per annum) for zero-coupon yields with maturity $\tau = 3M, 6M, 1Y, 2Y, 5Y,$ and $10Y$. The explanatory variables, $PC_j, j = 1, \dots, 6$, are the six principal components extracted from the panel of the six zero-coupon yields. The coefficient t -ratios in square brackets are based on (Newey-West) heteroskedasticity and autocorrelation robust standard errors.

Panel A: Daily observations from 06/17/1991 to 06/15/2001.								
Dep. variable	β_0	β_1	β_2	β_3	β_4	β_5	β_6	R_{Adj}^2
	[t -ratio]	[t -ratio]	[t -ratio]	[t -ratio]	[t -ratio]	[t -ratio]	[t -ratio]	
$v_{y_\tau}^2, \tau = 3M$	0.5579 [12.66]	0.0261 [0.96]	-0.0834 [-1.40]	-0.0256 [-0.14]	0.4313 [1.61]	-2.2271 [-2.41]	-4.6142 [-3.48]	3.88%
$v_{y_\tau}^2, \tau = 6M$	0.5403 [13.76]	0.0177 [0.81]	-0.0041 [-0.08]	0.3458 [1.86]	0.2092 [0.78]	-1.3257 [-1.95]	-1.9765 [-2.13]	1.37%
$v_{y_\tau}^2, \tau = 1Y$	0.6973 [15.32]	0.0156 [0.70]	0.0432 [0.87]	0.5210 [2.28]	0.0936 [0.37]	-1.1546 [-1.89]	-1.2352 [-1.29]	0.89%
$v_{y_\tau}^2, \tau = 2Y$	1.5974 [15.72]	0.0095 [0.21]	0.0084 [0.09]	0.9198 [1.77]	-0.3487 [-0.52]	-1.8517 [-1.61]	0.0179 [0.01]	0.30%
$v_{y_\tau}^2, \tau = 5Y$	1.4251 [17.86]	0.0121 [0.34]	-0.0108 [-0.15]	0.1982 [0.49]	-0.3527 [-0.64]	-2.0917 [-2.42]	0.9495 [0.57]	0.23%
$v_{y_\tau}^2, \tau = 10Y$	1.4296 [20.60]	-0.0037 [-0.11]	0.0977 [1.24]	0.3932 [1.11]	-0.2367 [-0.60]	-1.2139 [-1.35]	-0.6718 [-0.44]	0.19%

Panel B: Percent of residual variation explained by principal components extracted from the regressions error terms.						
1st	2nd	3rd	4th	5th	6th	
82.38%	9.67%	4.94%	1.79%	0.85%	0.38%	

Table 3: Evidence on the Affine Diffusive Spanning Condition for Sub-Samples. OLS estimates for the regression $v_{y_\tau}^2(t, 1) = \beta_0 + \sum_{j=1}^6 \beta_j PC_j(t, 1) + \varepsilon(t)$. The dependent variable is the daily realized volatility (in percent per annum) for zero-coupon yields with maturity $\tau = 3M, 2Y$, and $10Y$. The explanatory variables, $PC_j, j = 1, \dots, 6$, are the six principal components extracted from the panel of the six zero-coupon yields. The coefficient t -ratios in square brackets are based on (Newey-West) heteroskedasticity and autocorrelation robust standard errors.

Panel A: Daily observations from 01/05/1994 to 06/27/2001.								
Dep. variable	β_0	β_1	β_2	β_3	β_4	β_5	β_6	R_{Adj}^2
	[t -ratio]	[t -ratio]	[t -ratio]	[t -ratio]	[t -ratio]	[t -ratio]	[t -ratio]	
$v_{y_\tau}^2, \tau = 3M$	0.6159 [12.02]	-0.0630 [-1.04]	0.0191 [0.35]	0.1537 [0.74]	0.2630 [0.82]	-3.8623 [-2.66]	-4.0718 [-3.75]	5.19%
$v_{y_\tau}^2, \tau = 2Y$	1.6898 [12.76]	-0.0442 [-0.58]	-0.0414 [-0.28]	1.0882 [1.23]	-1.4754 [-1.97]	-2.5200 [-1.46]	0.5283 [0.28]	0.56%
$v_{y_\tau}^2, \tau = 10Y$	1.3789 [18.62]	-0.0334 [-0.65]	0.1719 [1.75]	0.4928 [1.15]	-0.9925 [-2.27]	-1.3089 [-1.14]	-0.5815 [-0.40]	1.10%
Panel B: Percent of residual variation explained by principal components extracted from the regressions error terms (01/05/1994—06/27/2001).								
	1st	2nd	3rd	4th	5th	6th		
	90.76%	5.10%	2.17%	0.98%	0.66%	0.33%		
Panel C: Daily observations from 06/17/1991 to 06/14/1996.								
$v_{y_\tau}^2, \tau = 3M$	0.4985 [14.65]	0.0882 [6.35]	0.0677 [1.47]	-0.3413 [-2.06]	1.0819 [2.48]	-1.0299 [-1.94]	-4.2817 [-3.04]	6.52%
$v_{y_\tau}^2, \tau = 2Y$	1.5363 [11.57]	0.0961 [2.24]	-0.0949 [-0.44]	0.6233 [1.17]	-0.3226 [-0.34]	-0.7125 [-0.38]	2.0541 [0.56]	0.13%
$v_{y_\tau}^2, \tau = 10Y$	1.5013 [15.01]	0.0487 [1.36]	0.0298 [0.21]	0.1786 [0.39]	-0.7365 [-0.91]	-1.0992 [-0.76]	-0.8120 [-0.25]	-0.23%
Panel D: Daily observations from 06/17/1996 to 06/15/2001.								
$v_{y_\tau}^2, \tau = 3M$	0.6181 [9.37]	-0.1976 [-2.13]	-0.0974 [-1.03]	0.3224 [1.51]	-0.4162 [-1.11]	-2.8122 [-1.32]	-6.2825 [-3.31]	7.15%
$v_{y_\tau}^2, \tau = 2Y$	1.6610 [12.62]	-0.2571 [-2.03]	-0.2087 [-0.85]	-0.3130 [-0.37]	-1.7625 [-2.61]	-2.7281 [-1.01]	-3.7712 [-1.67]	2.73%
$v_{y_\tau}^2, \tau = 10Y$	1.3600 [15.91]	-0.0936 [-1.06]	0.0375 [0.25]	-0.0489 [-0.11]	-1.0511 [-2.04]	-1.5319 [-0.82]	-2.8608 [-1.54]	1.42%

Table 4: Evidence on the Diffusive Volatility Spanning Condition for Monthly Realized Volatility Series. OLS estimates for the regression $v_{y_\tau}^2(t, 21) = \beta_0 + \sum_{j=1}^6 \beta_j PC_j(t, 21) + \varepsilon(t)$. The dependent variable is the overlapping *monthly* realized volatility (in percent per annum) for zero-coupon yields with maturity $\tau = 3M, 6M, 1Y, 2Y, 5Y, \text{ and } 10Y$. The explanatory variables, $PC_j, j = 1, \dots, 6$, are the principal components extracted from overlapping monthly averages of the zero-coupon yields. The coefficient t -ratios in square brackets are based on (Newey-West) heteroskedasticity and autocorrelation robust standard errors.

Panel A: Overlapping <i>monthly</i> observations from 06/17/1991 to 06/15/2001.								
Dep. variable	β_0	β_1	β_2	β_3	β_4	β_5	β_6	R_{Adj}^2
	[t -ratio]	[t -ratio]	[t -ratio]	[t -ratio]	[t -ratio]	[t -ratio]	[t -ratio]	
$v_{y_\tau}^2, \tau = 3M$	0.5593 [11.87]	0.0366 [1.37]	-0.0781 [-1.17]	-0.0405 [-0.23]	0.5807 [1.61]	-2.8117 [-2.07]	-5.2707 [-3.14]	22.27%
$v_{y_\tau}^2, \tau = 6M$	0.5413 [12.80]	0.0251 [1.10]	0.0048 [0.09]	0.3314 [1.77]	0.3090 [0.87]	-2.0244 [-1.85]	-2.1106 [-1.49]	12.09%
$v_{y_\tau}^2, \tau = 1Y$	0.6983 [14.05]	0.0236 [0.97]	0.0532 [0.91]	0.4883 [2.14]	0.1691 [0.49]	-1.4962 [-1.42]	-1.1981 [-0.79]	9.05%
$v_{y_\tau}^2, \tau = 2Y$	1.6012 [14.59]	0.0256 [0.50]	0.0307 [0.28]	0.8158 [1.50]	-0.4100 [-0.46]	-1.0284 [-0.55]	0.0834 [0.03]	3.07%
$v_{y_\tau}^2, \tau = 5Y$	1.4299 [16.29]	0.0234 [0.57]	0.0087 [0.11]	0.0357 [0.08]	-0.4585 [-0.63]	-1.2199 [-0.87]	0.8729 [0.38]	1.24%
$v_{y_\tau}^2, \tau = 10Y$	1.4196 [19.63]	-0.0093 [-0.27]	0.0802 [1.02]	0.2045 [0.56]	-0.2719 [-0.53]	-0.6915 [-0.49]	-1.2108 [-0.58]	2.14%

Panel B: Percent of residual variation explained by principal components extracted from the regressions error terms.						
1st	2nd	3rd	4th	5th	6th	
86.33%	7.33%	4.78%	0.94%	0.46%	0.15%	

Table 5: Evidence on the Diffusive Volatility Spanning Condition for Daily Constant-Maturity Yields. OLS estimates for the regression $v_{y_\tau}^2(t, 1) = \beta_0 + \sum_{j=1}^9 \beta_j PC_j(t, 1) + \varepsilon(t)$. The dependent variable is the daily realized volatility (in percent per annum) for zero-coupon yields with maturity $\tau = 3M, 6M, 1Y, 2Y, 5Y, \text{ and } 10Y$. The explanatory variables, PC_j , $j = 1, \dots, 9$, are the principal components extracted from the panel of constant-maturity daily zero-coupon yields, with maturities of 3M, 6M, 1Y, 2Y, 5Y, 7Y, 10Y, and 30Y. The coefficient t -ratios in square brackets are based on (Newey-West) heteroskedasticity and autocorrelation robust standard errors.

Panel A: Daily observations from 06/17/1991 to 06/15/2001.											
Dep. variable	β_0	β_1	β_2	β_3	β_4	β_5	β_6	β_7	β_8	β_9	$R^2_{Adj.}$
	[t -ratio]	[t -ratio]	[t -ratio]	[t -ratio]	[t -ratio]	[t -ratio]	[t -ratio]	[t -ratio]	[t -ratio]	[t -ratio]	
$v_{y_\tau}^2, \tau = 3M$	0.5579 [13.07]	0.0137 [0.52]	-0.0689 [-1.91]	-0.1033 [-0.76]	0.2248 [0.79]	-2.6019 [-2.78]	-3.5753 [-3.16]	-0.2076 [-0.27]	-1.8598 [-1.74]	6.7725 [1.50]	4.96%
$v_{y_\tau}^2, \tau = 6M$	0.5403 [14.46]	0.0150 [0.70]	-0.0048 [-0.16]	0.2095 [1.50]	0.0176 [0.06]	-2.0371 [-3.03]	-1.8609 [-2.24]	1.4178 [2.07]	-1.6242 [-1.49]	5.5111 [1.74]	2.95%
$v_{y_\tau}^2, \tau = 1Y$	0.6973 [15.97]	0.0184 [0.90]	0.0303 [0.99]	0.3207 [1.81]	-0.0599 [-0.20]	-1.9149 [-2.82]	-0.8696 [-0.93]	1.2308 [1.42]	-2.6999 [-2.02]	6.0932 [2.13]	2.00%
$v_{y_\tau}^2, \tau = 2Y$	1.5974 [16.11]	0.0118 [0.30]	0.0205 [0.34]	0.6472 [1.48]	-0.0812 [-0.14]	-2.1859 [-1.61]	-0.3312 [-0.16]	2.5463 [1.18]	-4.9705 [-1.65]	11.7131 [2.24]	0.91%
$v_{y_\tau}^2, \tau = 5Y$	1.4251 [18.20]	0.0128 [0.41]	-0.0010 [-0.02]	0.1511 [0.43]	0.2500 [0.54]	-1.7633 [-1.73]	0.0941 [0.06]	1.9884 [1.19]	-3.3096 [-1.40]	8.3585 [2.61]	0.64%
$v_{y_\tau}^2, \tau = 10Y$	1.4296 [22.03]	0.0110 [0.38]	0.0719 [1.51]	0.1045 [0.40]	-0.6781 [-1.32]	-2.3210 [-2.38]	-0.4861 [-0.34]	2.6225 [1.79]	-3.4607 [-1.61]	11.5905 [3.12]	1.60%

Panel B: Percent of the residual variation explained by principal components extracted from the regressions' error terms.

	1st	2nd	3rd	4th	5th	6th
	82.29%	9.71%	4.95%	1.81%	0.85%	0.39%

Table 6: Sample Correlations Between Daily Realized Volatility Series. The pairwise percentage sample correlations between daily realized yield volatility series for maturities of three and six months, one, two, five, and ten years. Newey-West standard errors, robust for the presence of heteroskedasticity and autocorrelation, are in round brackets.

	3M	6M	1Y	2Y	5Y
6M	75.34 (11.46)				
1Y	65.35 (17.18)	89.38 (7.55)			
2Y	53.93 (18.91)	80.07 (11.53)	93.10 (6.07)		
5Y	48.23 (20.03)	71.16 (13.35)	84.62 (8.76)	93.37 (5.28)	
10Y	39.25 (14.45)	55.38 (7.47)	61.48 (5.67)	67.44 (5.03)	70.78 (3.70)

Table 7: Daily, Weekly, and Monthly Volatility Forecasts. We report OLS estimates for $\frac{v_{y_\tau}(t+h,h)^2}{h} = \beta_0 + \beta_D v_{y_\tau}(t,1)^2 + \beta_W \frac{v_{y_\tau}(t,5)^2}{5} + \beta_M \frac{v_{y_\tau}(t,21)^2}{21} + \sum_{j=1}^6 \beta_j PC_j(t) + \varepsilon(t+h)$, where the dependent variable is the daily ($h = 1$), overlapping weekly ($h = 5$), and overlapping monthly ($h = 21$) realized volatility (in percent per annum) for zero-coupon yields with maturity $\tau = 3M, 6M, 1Y, 2Y, 5Y$, and $10Y$. $PC_j, j = 1, \dots, 6$, are the principal components extracted from the panel of zero-coupon yields. The sample period is 06/17/1991–06/15/2001. The coefficient t -ratios in square brackets are based on (Newey-West) heteroskedasticity and autocorrelation robust standard errors.

Panel A: Estimation results for the constrained model. The set of explanatory variables contains the principal components of the past yields, i.e., the coefficients β_D, β_W , and β_M are fixed at zero.

	$v_{y_\tau}^2, \tau = 3M$			$v_{y_\tau}^2, \tau = 6M$			$v_{y_\tau}^2, \tau = 1Y$		
h	1	5	21	1	5	21	1	5	21
β_0	0.5586 [12.72]	0.5597 [11.93]	0.5621 [11.58]	0.5395 [13.90]	0.5409 [13.03]	0.5428 [12.96]	0.6981 [15.31]	0.6989 [14.35]	0.6998 [14.30]
β_1	0.0371 [1.46]	0.0416 [1.56]	0.0505 [2.23]	0.0256 [1.23]	0.0287 [1.31]	0.0374 [1.93]	0.0228 [1.05]	0.0265 [1.15]	0.0342 [1.56]
β_2	-0.0782 [-1.19]	-0.0776 [-1.07]	-0.0701 [-0.98]	0.0036 [0.07]	0.0040 [0.07]	0.0105 [0.19]	0.0513 [1.00]	0.0537 [0.94]	0.0612 [1.08]
β_3	0.0287 [0.16]	0.0096 [0.06]	0.0308 [0.19]	0.4172 [2.24]	0.4177 [2.32]	0.3997 [2.32]	0.6097 [2.60]	0.6022 [2.57]	0.5388 [2.59]
β_4	0.5462 [2.19]	0.6019 [2.24]	0.5369 [1.71]	0.2987 [1.21]	0.3167 [1.18]	0.3649 [1.33]	0.1385 [0.54]	0.1786 [0.68]	0.2305 [0.83]
β_5	-2.3164 [-2.54]	-2.0695 [-2.27]	-1.4686 [-2.41]	-1.7423 [-2.43]	-1.7134 [-2.40]	-1.1334 [-2.28]	-1.6523 [-2.46]	-1.5736 [-2.33]	-0.9788 [-1.96]
β_6	-5.2141 [-3.71]	-4.8909 [-3.26]	-3.1640 [-2.44]	-2.2739 [-2.34]	-1.9423 [-1.70]	-1.5063 [-1.39]	-1.2986 [-1.32]	-1.0459 [-0.88]	-1.0685 [-0.95]
R_{Adj}^2	4.63%	12.71%	16.19%	2.29%	8.03%	13.15%	1.51%	6.14%	11.83%
	$v_{y_\tau}^2, \tau = 2Y$			$v_{y_\tau}^2, \tau = 5Y$			$v_{y_\tau}^2, \tau = 10Y$		
h	1	5	21	1	5	21	1	5	21
β_0	1.6039 [15.58]	1.6052 [14.69]	1.6072 [14.80]	1.4320 [17.76]	1.4343 [16.62]	1.4376 [16.52]	1.4233 [20.52]	1.4246 [19.61]	1.4269 [19.83]
β_1	0.0233 [0.49]	0.0291 [0.57]	0.0438 [0.86]	0.0225 [0.60]	0.0250 [0.62]	0.0293 [0.70]	-0.0116 [-0.34]	-0.0097 [-0.27]	-0.0052 [-0.15]
β_2	0.0248 [0.26]	0.0324 [0.30]	0.0549 [0.53]	0.0060 [0.08]	0.0115 [0.15]	0.0215 [0.29]	0.0780 [1.04]	0.0793 [1.00]	0.0787 [1.06]
β_3	1.0842 [2.03]	1.0824 [1.92]	0.9694 [1.95]	0.3117 [0.75]	0.3197 [0.72]	0.2397 [0.58]	0.4144 [1.19]	0.4017 [1.12]	0.3328 [1.02]
β_4	-0.3677 [-0.50]	-0.2930 [-0.40]	-0.1168 [-0.16]	-0.3217 [-0.54]	-0.3083 [-0.51]	-0.1848 [-0.32]	-0.1933 [-0.46]	-0.1265 [-0.29]	-0.1489 [-0.33]
β_5	-2.9856 [-2.27]	-2.5130 [-1.99]	-1.1685 [-1.36]	-2.9442 [-3.08]	-2.5099 [-2.74]	-1.3953 [-2.12]	-2.1235 [-2.22]	-1.6833 [-1.78]	-0.9174 [-1.28]
β_6	0.0543 [0.03]	0.7580 [0.31]	-0.0883 [-0.04]	0.9018 [0.51]	1.5460 [0.84]	0.9717 [0.66]	-0.8127 [-0.46]	-0.4707 [-0.26]	-0.1519 [-0.11]
R_{Adj}^2	0.69%	3.08%	5.27%	0.64%	2.42%	2.36%	0.46%	1.86%	2.99%

Table 7, continued:

Panel B: Estimation results for the constrained model. The set of explanatory variables contains volatility components only, i.e., the coefficients $\beta_j, j = 1, \dots, 6$, are fixed at zero.									
	$v_{y\tau}^2, \tau = 3M$			$v_{y\tau}^2, \tau = 6M$			$v_{y\tau}^2, \tau = 1Y$		
h	1	5	21	1	5	21	1	5	21
β_0	0.2042 [4.47]	0.2534 [5.34]	0.3481 [7.31]	0.2146 [4.91]	0.2438 [5.66]	0.3350 [7.30]	0.3178 [5.94]	0.3424 [6.14]	0.4252 [6.81]
β_D	0.0603 [0.92]	0.0479 [2.67]	0.0252 [3.20]	-0.0094 [-0.35]	0.0236 [1.83]	0.0121 [2.22]	0.0253 [1.11]	0.0122 [1.16]	0.0067 [1.55]
β_W	0.2339 [3.10]	0.1728 [2.96]	0.1245 [3.03]	0.2428 [2.81]	0.1617 [2.99]	0.1332 [3.18]	0.0937 [1.78]	0.0833 [1.98]	0.0806 [2.93]
β_M	0.3449 [2.67]	0.3321 [3.34]	0.2364 [3.09]	0.3738 [4.54]	0.3703 [4.54]	0.2431 [3.23]	0.4290 [5.69]	0.4186 [5.00]	0.3086 [4.14]
R_{Adj}^2	8.07%	16.85%	16.95%	5.59%	15.58%	16.77%	3.40%	11.22%	16.40%
	$v_{y\tau}^2, \tau = 2Y$			$v_{y\tau}^2, \tau = 5Y$			$v_{y\tau}^2, \tau = 10Y$		
h	1	5	21	1	5	21	1	5	21
β_0	0.8113 [6.75]	0.8666 [6.81]	1.1050 [7.80]	0.6707 [7.92]	0.7386 [8.01]	0.9559 [8.95]	0.8178 [8.06]	0.8649 [8.41]	1.0629 [9.06]
β_D	0.0290 [1.52]	0.0052 [0.61]	0.0058 [1.56]	0.0473 [1.45]	0.0158 [1.66]	0.0092 [2.18]	0.0155 [0.46]	0.0170 [2.07]	0.0044 [1.11]
β_W	0.0644 [1.37]	0.0939 [2.38]	0.0926 [3.48]	0.1332 [2.78]	0.1407 [4.29]	0.0989 [3.86]	0.1780 [2.22]	0.1226 [2.15]	0.0756 [3.52]
β_M	0.4033 [4.45]	0.3638 [3.97]	0.2162 [3.13]	0.3542 [4.82]	0.3322 [4.89]	0.2302 [4.14]	0.2334 [3.27]	0.2551 [3.37]	0.1767 [2.65]
R_{Adj}^2	2.53%	8.90%	11.18%	3.79%	11.12%	12.81%	2.18%	6.81%	7.51%

Table 7, continued:

Panel C: Estimation results for the unconstrained model. The set of explanatory variables contains both volatility components and past yields. The coefficients $\beta_j, j = 1, \dots, 6$, are not reported.

	$v_{y\tau}^2, \tau = 3M$			$v_{y\tau}^2, \tau = 6M$			$v_{y\tau}^2, \tau = 1Y$		
h	1	5	21	1	5	21	1	5	21
β_0	0.2685 [5.94]	0.3221 [6.79]	0.4055 [8.49]	0.2474 [5.31]	0.2771 [6.06]	0.3692 [7.97]	0.3488 [6.23]	0.3746 [6.60]	0.4610 [7.53]
β_D	0.0522 [0.79]	0.0399 [2.16]	0.0194 [2.62]	-0.0145 [-0.53]	0.0185 [1.53]	0.0082 [1.76]	0.0210 [0.93]	0.0079 [0.80]	0.0037 [0.95]
β_W	0.2034 [2.65]	0.1433 [2.49]	0.1031 [2.74]	0.2202 [2.67]	0.1417 [2.80]	0.1156 [3.01]	0.0748 [1.53]	0.0647 [1.67]	0.0643 [2.60]
β_M	0.2611 [2.23]	0.2408 [2.62]	0.1567 [2.17]	0.3331 [3.92]	0.3285 [3.97]	0.1969 [2.55]	0.4053 [5.91]	0.3928 [5.11]	0.2743 [3.73]
$R^2_{Adj.}$	9.50%	21.82%	24.42%	6.56%	19.57%	24.31%	4.18%	14.83%	23.48%
	$v_{y\tau}^2, \tau = 2Y$			$v_{y\tau}^2, \tau = 5Y$			$v_{y\tau}^2, \tau = 10Y$		
h	1	5	21	1	5	21	1	5	21
β_0	0.8205 [6.87]	0.8790 [6.99]	1.1297 [8.13]	0.6635 [7.85]	0.7327 [7.96]	0.9554 [9.08]	0.8221 [7.88]	0.8703 [8.47]	1.0763 [9.46]
β_D	0.0247 [1.28]	0.0013 [0.16]	0.0035 [0.97]	0.0419 [1.28]	0.0108 [1.19]	0.0063 [1.55]	0.0122 [0.36]	0.0141 [1.80]	0.0026 [0.69]
β_W	0.0503 [1.11]	0.0796 [2.00]	0.0820 [3.31]	0.1171 [2.55]	0.1254 [3.74]	0.0894 [3.69]	0.1645 [2.15]	0.1118 [2.07]	0.0686 [3.41]
β_M	0.4145 [4.91]	0.3718 [4.42]	0.2111 [3.17]	0.3771 [5.02]	0.3527 [5.14]	0.2402 [4.41]	0.2420 [3.30]	0.2615 [3.48]	0.1738 [2.69]
$R^2_{Adj.}$	3.06%	11.35%	15.02%	4.35%	13.34%	14.82%	2.50%	8.25%	9.71%

Table 8: Sample Correlations Between Zero-Coupon and ‘Raw’ Bond Yields.

Panel A: Correlations between intra-daily zero-coupon and ‘raw’ yields.					
For T-Bills, the ‘raw’ yields are zero-coupon yields with maturities closest to 3M, 6M, and 1Y.					
For T-Notes, the ‘raw’ yields are coupon yields with maturities closest to 2Y, 5Y, and 10Y.					
3M	6M	1Y	2Y	5Y	10Y
99.99%	100.00%	99.94%	99.97%	99.84%	99.31%

Panel B: Correlations between daily constant-maturity zero-coupon and par yields.								
3M	6M	1Y	2Y	3Y	5Y	7Y	10Y	30Y
100.00%	100.00%	100.00%	100.00%	99.98%	99.86%	99.74%	99.52%	96.97%

Table 9: Sample Correlations Between Realized Volatility Measures Constructed from Intra-Daily Zero-Coupon Yields and from ‘Raw’ Bond Yields. For Treasury Bills, the ‘raw’ yields are zero-coupon yields with maturities closest to three and six months, and one year. For Treasury Notes, the ‘raw’ yields are coupon yields with maturities closest to two, five, and ten years.

3M	6M	1Y	2Y	5Y	10Y
99.49%	99.93%	99.79%	99.99%	99.87%	96.86%

Table 10: Percentage Number of Unique Yield Observations in the Intra-Day GovPX Sample. Panel A reports the percentage ratio between unique and total yield observations during the full June 17, 1991—June 15, 2001 sample period. Panel B reports the same percentage ratios computed after discarding trading days with three or more hours of inactivity.

Panel A: Full sample (June 17, 1991—June 15, 2001)					
3M	6M	1Y	2Y	5Y	10Y
80.74%	81.32%	91.04%	97.65%	94.21%	87.62%
Panel B: Trading days with 3 or more hours of inactivity are discarded					
3M	6M	1Y	2Y	5Y	10Y
81.77%	82.54%	92.69%	99.10%	97.05%	91.20%




Figure 1: Realized Volatility Series. The plots depict the square root of the daily percentage realized volatility measures (in percent per day), $v_{y\tau}$, for the three- and six-month, one-, two-, five-, and ten-year maturity yields (sample period: 06/17/1991–06/15/2001).




Figure 2: The Term Structure of Yield Volatility. The continuous line depicts the sample standard deviation of daily changes in zero-coupon yields (in percent per day). The dashed line depicts the average daily realized volatility v_y (in percent per day). In both cases, the plots are constructed using yields with maturities of three and six months, one, two, five, and ten years (sample period: 06/17/1991–06/15/2001).




Figure 3: Sample Autocorrelations for the Realized Volatility Series. The continuous line plots the sample autocorrelations for daily logarithmic realized volatility. The dotted line depicts the minimum-distance estimates of the hyperbolic decay rate, $cLag^{2d-1}$. Sample period: 06/17/1991–06/15/2001.




Figure 4: U.S. Three-Month Treasury Bill Yield Volatility. The plots depict daily realized volatility (i.e., $v_y(t)$, in percent per day) versus one-day-ahead volatility forecasts based on the SNP model (i.e., $\sqrt{V(y_t|x_t;\xi)}$, in percent per day). Sample period: 06/17/1991–06/15/2001.

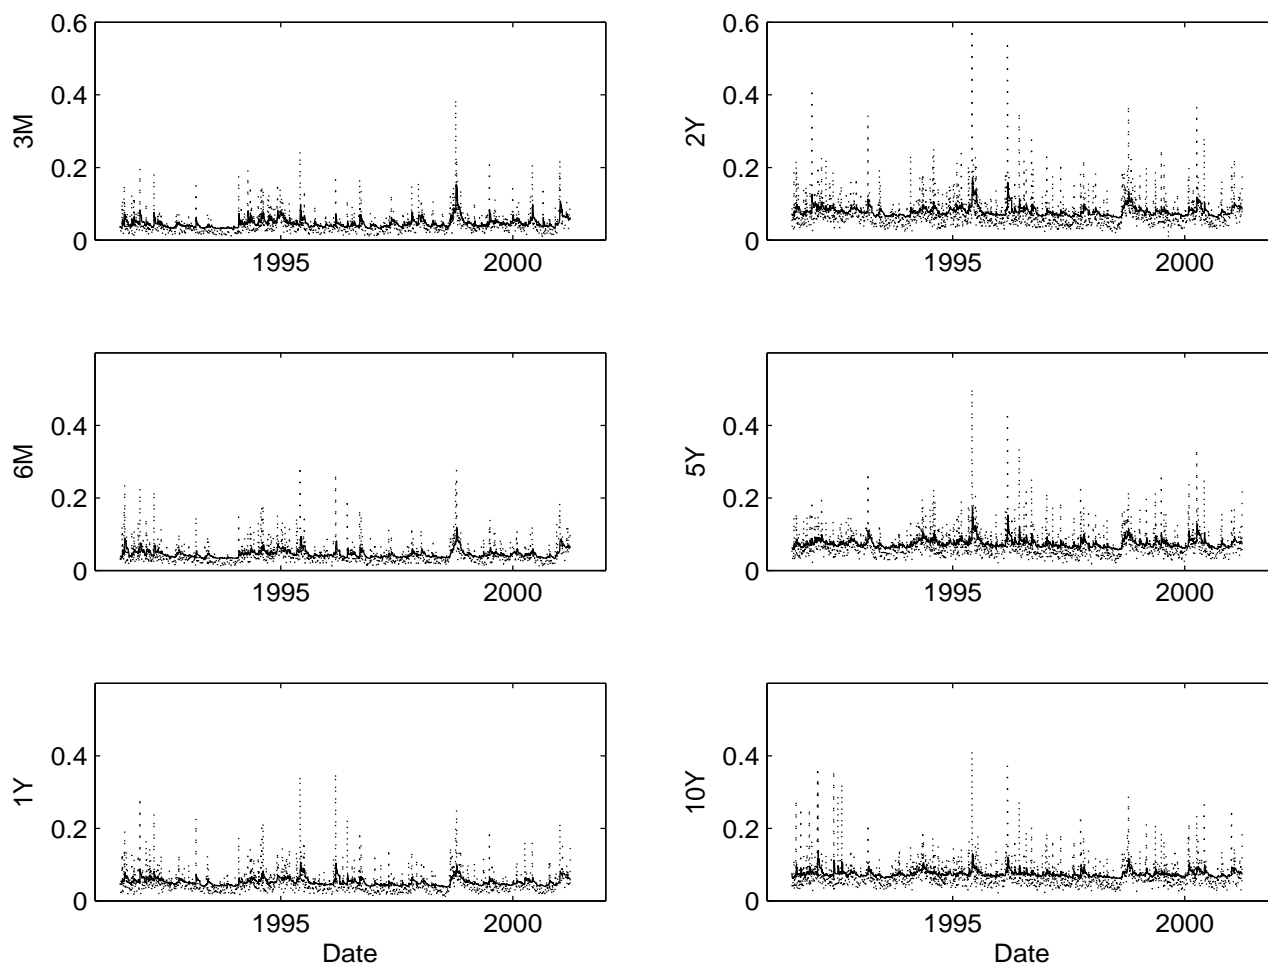


Figure 5: Volatility Forecasts and Realized Volatility Series. The plots depict one-day-ahead HAR-RV forecasts, $\hat{v}_{y_\tau}(t, 1)$, and realized volatility series, $v_{y_\tau}(t, 1)$ (in percent per day). Sample period: 07/17/1991—03/27/2001.