

Can Bonds Hedge Volatility Risk in the U.S. Treasury Market? A Specification Test for Affine Term Structure Models

Torben G. Andersen and Luca Benzoni*

First Draft: January 15, 2005

This Draft: March 28, 2006

*Torben G. Andersen is at the Kellogg School of Management, Northwestern University, 2001 Sheridan Road, Evanston, IL 60208, 847-467-1285, t-andersen@northwestern.edu and the NBER. Luca Benzoni is at the Carlson School of Management, University of Minnesota, 321 19th Ave S, Minneapolis, MN 55455, 612-624-1075, lbenzoni@umn.edu. We are grateful to Darrell Duffie, Michael Fleming, Bob Goldstein, Mike Johannes, 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, 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 most recent version of the paper can be downloaded from <http://www.umn.edu/~lbenzoni>.

Can Bonds Hedge Volatility Risk in the U.S. Treasury Market? A Specification Test for Affine Term Structure Models

Abstract

We investigate whether bonds can hedge 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 appear largely unrelated to the cross-section of yields. We conclude that a broad class of affine diffusive, quadratic diffusive and affine jump-diffusive models is incapable of accommodating the observed yield volatility dynamics. Hence, yield volatility risk per se cannot be hedged by taking positions in the Treasury bond market. We also advocate using these empirical yield volatility measures more broadly as a basis for specification testing and (parametric) model selection within the term structure literature.

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 has averaged approximately \$539 billion, about tenfold the daily volume at, e.g., 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 times. The Federal Reserve System uses this 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. Related, these securities serve as an input and a 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 and the associated derivatives. In particular, Litterman, Scheinkman, and Weiss (1991) note that investors have long realized that the relative attractiveness of bonds with different maturities and coupons depends not only on 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 yields on bonds with any maturity is a linear combination of the term structure of bond yields. As such, these models predict that interest rates volatility risk can be hedged by trading in a portfolio of bonds. In this paper, we empirically examine this prediction.

Previous studies have investigated this issue by using data on the London Interbank Offered Rate (LIBOR), swap rates, and 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. They refer to this feature as the ‘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, which supports the notion that bond markets are complete.¹

More recently, other studies have investigated the properties of term structure models that embed the USV restriction. Also in this case, 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 rates curve can be negatively correlated with the time-series of volatility estimated with a standard GARCH approach. In response, they argue that a 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 (a result that he attributes to the 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 in the model when they confront it with data on caps and swaptions. They conclude that derivatives should be used for 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. A potential and 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 data on derivatives prices are required to obtain efficient inference along this dimension since measurement errors may render the theoretical link between yield levels and volatility elusive from observed bond data 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

¹The results in Fan et al. (2003) are consistent with the early study by Litterman, Scheinkman, and Weiss (1991), who argue that the yield spreads of certain ‘butterfly’ combinations (which are highly sensitive to volatility) correlate highly with the curvature factor.

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)).² 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 volatility estimates of the yields of these securities. 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 yields' quadratic variation by summing the squared intra-daily changes in these yields. We may consequently test directly whether bonds can hedge 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, we emphasize that the affine model restrictions we test are based exclusively on the affine structure under the so-called 'equivalent martingale' or 'pricing' measure, so it is independent of whether the representation under the 'actual' measure is non-affine as proposed by, e.g., 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 in the construction of the zero-coupon bond yields. Third, we avoid using data for any market other than the specific fixed-income market we analyze. This approach sidesteps potentially serious concerns regarding the reliability of derivatives prices obtained from secondary over-the-counter markets due to liquidity and market microstructure issues. Fourth, we also avoid having to specify an ad hoc time series model for the conditional yield variance (expected quadratic variation) process. This is one approach previously adopted to gauge the coherence between the volatility dynamics implied by the model and the data (see, e.g., CDGJ and Dai and Singleton (2003)). On the other hand, the availability of the high-frequency based realized volatility series also allows us, if necessary or convenient, to construct simple, yet efficient, forecasts of future quadratic yield variation. In fact, such forecasts typically outperform those obtained from standard time series volatility models based on daily or lower frequency data (see, 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

²Direct measures of the low-frequency return variance have been obtained from cumulative higher-frequency squared returns previously, although the theoretical basis for such procedures were not articulated. Early studies include French, Schwert, and Stambaugh (1987), Hsieh (1991), Poterba and Summers (1986), Schwert (1989, 1990), Taylor and Xu (1997).

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 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 $N = 4$ and/or $N = 3$ 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 also expand our specification analysis beyond the affine diffusive model class. It turns out that the ‘quadratic term structure model’ studied by Ahn et al. (2002) have 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 changes the spanning restriction qualitatively. For the affine jump-diffusion class, we document that the direct spanning condition fails while the conditionally expected future quadratic variation, as before, will be spanned by the yield cross-section. We therefore conduct a second set of specification tests exploiting efficient quadratic yield variation forecasts generated directly from the realized yield volatility series. These volatility forecasts can still be constructed for an arbitrary horizon so we retain many of the advantages discussed above. However, they are no longer entirely model-free, so we also compare these predictions to a more familiar parametric volatility estimate, or proxy, obtained from daily data.

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 as well as the associated quadratic variation forecasts generated from the realized yield volatilities. We confirm that the properties of these series are qualitatively consistent with the anticipated 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 not subject to any idiosyncratic variation or systematic measurement errors which 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) have insignificant coefficients in the majority of these regressions. Higher-order principal components often enter significantly, although with limited explanatory power. This finding is at odds with the notion that curvature is related to interest rate volatility. We have also confirmed that these results carry over when the spanning condition is tested by using weekly and monthly realized volatility measures. Moreover, an analysis of sub-samples shows that the volatility spanning condition is violated consistently across different sample periods. It is consequently not surprising that we also reject the auxiliary implications of the affine multi-factor term structure models in which a single factor determines the conditional variance of the state variables. Finally, we provide evidence against a conditional version of the volatility spanning condition that holds in 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 the USV restriction—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, e.g., the hedging of interest rate volatility risk. It remains a topic for future research to determine which type of extension to such models can offer a framework that is both tractable and empirically successful.

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-diffusion term structure models. We clarify how this link 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 while in Section 4 we document the salient features of our volatility estimates relative to prior findings in the literature. Section 5 contains our main empirical findings. In Section 5.1, we focus on affine-diffusion models, while in Section 5.2 we extend our analysis for the presence of jumps. Concluding remarks are in Section 6.

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

This section reviews known aspects of affine diffusion term structure models with an emphasis on those features that we examine in our subsequent empirical inquiry. Following Duffie 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 Q -measure, \mathcal{K} and Θ are $N \times N$ matrices, and $S(t)$ is a diagonal matrix with the i th diagonal element given by $[S(t)]_{ii} = \alpha_i + \beta'_i X(t)$.

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 functions $A(\tau)$ and $B(\tau)$ solve a system of ordinary differential equations.

This result provides a direct 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 given by

$$P(t, \tau) = e^{-\tau y_\tau(t)}. \quad (4)$$

Thus, we have

$$y_\tau(t) = -\frac{A(\tau)}{\tau} + \frac{B(\tau)'}{\tau} X(t). \quad (5)$$

An application of Itô's Lemma to equation (5) shows that the yield y_τ follows a diffusion process:

$$dy_\tau(t) = \mu_{y_\tau}(X(t), t) dt + \frac{B(\tau)'}{\tau} \Sigma \sqrt{S(t)} dW^Q(t). \quad (6)$$

Consequently, the (instantaneous) quadratic variation of the yield given as the squared yield volatility coefficient for y_τ is

$$V_{y_\tau}(t) = \frac{B(\tau)'}{\tau} \Sigma S(t) \Sigma' \frac{B(\tau)}{\tau}. \quad (7)$$

Now, the elements of the $S(t)$ matrix are affine in the state vector $X(t)$, i.e., $[S(t)]_{ii} = \alpha_i + \beta_i' X(t)$. Further, equation (5) implies that each state variable in the vector $X(t)$ is an affine function of the bond yields $Y(t) = \{y_{\tau_j}(t), j = 1, \dots, J\}$, where we assume that we observe a larger set of yields than there are state variables, i.e., $J \geq N$. Thus, for any τ we can find a set of constants $a_{\tau,j}$, $j = 0, \dots, J$, so that

$$V_{y_\tau}(t) = a_{\tau,0} + \sum_{j=1}^J a_{\tau,j} y_{\tau_j}(t). \quad (8)$$

This derivation underscores the fact that the quadratic 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 (8). Since the quadratic variation, almost by definition, is also 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 cross section of bond yields and the short rate variation. Of course, the same relationship remains valid for any fixed maturity yield, as indicated above, implying a range of simultaneous constraints across the yield volatility spectrum.³ One crucial implication is that an investor can use a portfolio of zero-coupon bonds to hedge volatility risk in the Treasury market.

There are some subtleties involved in constructing the appropriate empirical yield variation measures to represent the yield quadratic variation which appears prominently in the affine model restrictions above. Hence, we spell out the role of the quadratic yield variation and its relationship to the cross-section of yields in the affine model setting in detail. First, we recall the definition of the quadratic variation process for the constant maturity yield y_τ initiated at time $t_0 = 0$,

$$QV_{y_\tau}(t) \equiv \int_0^t V_{y_\tau}(s) ds. \quad (9)$$

³A notable exceptions is the USV class of models of Casassus et al. (2004), Collin-Dufresne and Goldstein (2002), CDGJ, and the related model class explored in Kimmel (2004).

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

$$QV_{y_\tau}(t, h) \equiv QV_{y_\tau}(t) - QV_{y_\tau}(t - h) = \int_{t-h}^t V_{y_\tau}(s) ds. \quad (10)$$

Next, observe that equation (8) implies,

$$\int_{t-h}^t V_{y_\tau}(s) ds = a_{\tau,0} + \sum_{j=1}^J a_{\tau,j} \int_{t-h}^t y_{\tau_j}(s) ds. \quad (11)$$

We may rewrite equation (11) in a more readily interpretable manner by defining $\bar{y}_{\tau_j}(t, h)$ as the average yield of y_{τ_j} over $[t - h, t]$. This term corresponds directly to the integral on the extreme right of equation (11). Then, also exploiting equation (10), we obtain the following restriction,

$$QV_{y_\tau}(t, h) = a_{\tau,0} + \sum_{j=1}^J a_{\tau,j} \bar{y}_{\tau_j}(t, h). \quad (12)$$

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 open and/or 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 data. Perhaps as a consequence of this fact, the quadratic variation of the yields has not been the focus of direct measurement or testing within the term structure literature. Instead, most studies exploring the affine model restrictions rely on parametric conditional yield variance estimates or implied volatility measures backed out from derivatives prices. Although this approach quite generally can be rigorously justified as arising from the relevant theory, the replacement of the quadratic variation with an alternative volatility proxy is not innocuous. It inevitably entails a loss of power in terms of testing the affine spanning condition. We discuss these issues at length below.

2.2 Spanning Restrictions for the Conditional Yield Variance

The pure diffusive no-arbitrage setting and the associated semi-martingale representation of bond prices imply that the *predictable* yield variation over short daily or intraday periods are negligible (of order dt^2) relative to the variation of the yield innovations (of order $dW^Q(t)^2 = dt$). In practical terms this means that we safely may ignore the conditional mean of yield changes in computing the conditional yield variance over short horizons. It follows that the conditional yield variance over $[t - h, t]$ is simply,

$$\text{Var}_{t-h}^P [y_\tau(t)] = E_{t-h}^P [(y_\tau(t) - y_\tau(t - h))^2]. \quad (13)$$

where the subscript $t-h$ indicates that the variance and expectation are evaluated conditional on the time $t-h$ information set, and the superscript P indicates the so-called actual or ‘physical’ probability measure as opposed to the equivalent martingale pricing measure, Q . Using this observation and letting the integer $n \geq 1$ denote the number of equidistant intraday yield changes sampled over the (short) interval $[t-h, t]$, we have

$$E_{t-h}^P [(y_\tau(t) - y_\tau(t-h))^2] = E_{t-h}^P \left[\sum_{i=1, \dots, n} \left(y_\tau \left(t-h + \frac{ih}{n} \right) - y_\tau \left(t-h + \frac{(i-1)h}{n} \right) \right)^2 \right]. \quad (14)$$

Equation (14) holds for an arbitrary n , so by letting n increase towards infinity we have, by basic properties of the quadratic variation process, that

$$\text{Var}_{t-h}^P [y_\tau(t)] = E_{t-h}^P [QV_{y_\tau}(t, h)]. \quad (15)$$

This relation highlights important differences between these two concepts of yield volatility. The conditional variance is a forward looking expectation of the future sample path variation, and it is thus fundamentally an ex-ante concept. In contrast, the quadratic variation denotes the actual realized variation in the sample path, so it is an ex-post (realization) measure. If the volatility is (conditionally) deterministic as when volatility is constant, then the two notions of yield variation coincides. In general, however, the yield variation has a sizeable, genuinely unpredictable innovation component which renders volatility stochastic. As such, the sample variability of the quadratic yield variation process will inevitably be substantially larger than for the conditional yield variance process because sample realizations, by construction, fluctuate more than their a priori expectations. The spanning condition in equation (12) ties the contemporaneous yield level and yield variation together directly in terms of realizations. Since the (realized) quadratic variation is inherently more variable than the ex-ante expectation, there is much more sample variation for the cross-section of yields to rationalize than there is for 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 (12) into equation (15), to obtain,

$$\text{Var}_{t-h}^P [y_\tau(t)] = a_{\tau,0} + \sum_{j=1}^J a_{\tau,j} E_{t-h}^P [\bar{y}_{\tau_j}(t, h)]. \quad (16)$$

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 drift specification, 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

$$\text{Var}_{t-h}^P [y_\tau(t)] = b_{\tau,0} + \sum_{j=1}^J b_{\tau,j} y_{\tau_j}(t-h). \quad (17)$$

This affine spanning condition for the (true) conditional variance process has been tested in a number of prior empirical studies. As already alluded to above, it is inherently less powerful in terms of testing the underlying affine model than equation (12) and it requires the model to be affine under both the P and Q measures. A final caveat is that the condition is only valid if it is the true conditional variance process that appears on the left hand side of (17). In other words, in an affine model the true conditional variance is given by *some* linear combination of the current yields. Hence, if the conditional variance is specified as an ad hoc time series model such as, e.g., a GARCH style model, which inevitably is subject to some degree of misspecification, then the class of forecasts spanned by the yield cross-section on the right hand side of (17) should forecast future realizations of the quadratic yield variation better, or at least no worse, than the GARCH model. Of course, such comparisons of relative predictive ability require a fairly long sample in order to achieve sufficient statistical power. In contrast, the fundamental affine spanning condition (12) should in principle apply to day-by-day realizations which can be tested straightforwardly from small samples.

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 prices are available. If this is the case, the implied conditional variance forecasts (under Q) should also be spanned by the cross-section of yields. As for the fundamental spanning condition (12) this formulation relies only on the model being affine under Q , so it applies also for the Duarte style extensions of the basic affine model. On the other hand, the forecast horizon must necessarily equal the maturity of the derivatives contracts, which typically will entail 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 $\Delta y_\tau(t) = y_\tau(t) - y_\tau(t-h)$ and $\Delta \text{Var}_t^P [y_\tau(t+h)] = \text{Var}_t^P [y_\tau(t+h)] - \text{Var}_{t-h}^P [y_\tau(t)]$ denote period-by-period yield changes and conditional variance changes respectively, it follows from equation (17) that

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

Of course, we can derive an equivalent expression for changes in the implied volatility forecasts under the Q measure. Several studies employ these specifications of the affine spanning conditions as the basis for their tests. This approach is obviously closely related to the specification in equation

(17) so we only report results for 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 (18).

2.3 Correlation in Yield Volatility Innovations

The volatility spanning condition in equation (12) applies to any affine model of the form (1)-(2). Naturally, additional restrictions may apply in more specific model 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 is a positive integer equal to the number of latent factors. For example, Bikbov and Chernov (2004), CDGJ, and Thompson (2004) use this model with $N = 4$ and/or $N = 3$ in their studies.

Because of the prominence of the affine $A_1(N)$ specification, we further develop the implied volatility restrictions within this model class. Itô’s Lemma applied to the general affine condition stated in equation (7) shows that the yield’s quadratic variation is a diffusion process:

$$dV_{y_\tau}(t) = \mu_{y_\tau}^V(X(t), t) dt + \sigma_{y_\tau}^V(X(t), t) dW^Q(t). \quad (19)$$

In the special case of the $A_1(N)$ model, simple algebra yields that the *stochastic* component of the dV_{y_τ} diffusion in equation (19) simplifies to

$$dV_{y_\tau}^{stochastic}(t) = \sigma_\tau(t) \sqrt{x_1(t)} dW_1^Q(t), \quad (20)$$

where $x_1(t)$ denotes the state variable that drives the volatility dynamics and $\sigma_\tau(t)$ is a deterministic term.

Consequently, for all $A_1(N)$ model the volatility innovations for any pair of yields y_{τ_1} and y_{τ_2} are perfectly correlated:

$$\text{corr}(dV_{y_{\tau_1}}^{stochastic}, dV_{y_{\tau_2}}^{stochastic}) = 1. \quad (21)$$

This is, of course, also not a novel insight. However, as for the volatility spanning condition, this relation has not previously been subjected to direct empirical scrutiny based on nonparametric or model-free measures of quadratic yield variation. Below, we use our realized volatility measures to examine whether this model specific prediction is consistent with empirical evidence.

2.4 Extensions to Quadratic and Affine Jump-Diffusion 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, e.g., Ahn et al. (2003), the QTSM is isomorphic to the ATSM in its mechanism for generating volatility as volatility 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.

A modification of the yield spanning condition is required, however, to accommodate the possibility of jumps in the state variables and yields. This is a relevant extension since the empirical evidence strongly suggests that macroeconomic announcements may induce instantaneous jumps in the yields upon release.⁴ Following Duffie et al. (2000), the state vector X in an affine jump-diffusion model has Q -dynamics

$$dX(t) = \mathcal{K}(\Theta - X(t))dt + \Sigma\sqrt{S(t)}dW^Q(t) + Zdq^Q(t), \quad (22)$$

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 a fixed probability distribution ν^Q , and $J(t) \equiv \Delta X(t) = Z(t)dq^Q(t)$ is the corresponding vector jump process which 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 (6) we have

$$dy_\tau(t) = \mu_{y_\tau}(X(t), t) dt + \frac{B(\tau)'}{\tau} \left[\Sigma \sqrt{S(t)} dW^Q(t) + Z dq^Q(t) \right] \quad (23)$$

and we further obtain,

$$QV_{y_\tau}(t, h) = \int_t^{t+h} \frac{B(T-s)'}{T-s} \Sigma S(s) \Sigma' \frac{B(T-s)}{T-s} ds + \sum_{t-h \leq s \leq t} \left[\frac{B(T-s)'}{T-s} J(s) J(s)' \frac{B(T-s)}{T-s} \right]. \quad (24)$$

Notice that jump realizations induce a quadratic form dependency into the relation between the quadratic yield variation process and the state variables so the basic affine spanning restriction no longer applies. 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,

$$E_{t-h}^Q \left[\sum_{t-h \leq s \leq t} J(s) J(s)' \right] = E_{t-h}^Q \left[Z Z' \int_{t-h}^t (\lambda_0 + \lambda'_X X(s)) ds \right]. \quad (25)$$

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

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

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 variant of the spanning condition for the pure diffusive case (12) over the (short) future time interval $[t, t + h]$,

$$E_t^Q [QV_{y_\tau}(t + h, h)] = b_{\tau,0} + \sum_{j=1}^J b_{\tau,j} y_{\tau_j}(t). \quad (26)$$

This spanning condition is entirely equivalent to equation (17) and it is straightforward to derive also the corresponding version of equation (18). Consequently, the extension to include jumps has no impact on the conditional yield variance forecast spanning conditions under the Q measure—it should continue to hold in this more general setting. Thus, the current yields (yield changes) should span a correctly specified and measured implied volatility forecast (changes in implied volatility forecasts) derived from derivatives contracts. The identical spanning restriction for conditional yield variance forecasts under the actual probability measure, P , will apply 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 —which is in line with existing formulations of empirical models in the literature—and the diffusive dynamics remain affine, excluding only, as in Section 2.2, models such the one in Duarte (2004).

The conclusion is that the general formulation in Section 2.2, providing spanning conditions for the conditional yield variance, remains largely unaltered in the affine jump-diffusion case, while the much stricter realization-by-realization spanning constraint in equation (12) no longer applies. On the other hand, all formulations of the yield spanning conditions are invariant to the quadratic Gaussian model assumption as these literally can be encompassed within the affine diffusive setting.

2.5 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 actual 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 the concept for direct measurement of the quadratic variation of the yield on a bond with maturity τ . Specifically, we compute the realized volatility of the yield y_τ over the interval $[t - h, t]$ by

$$v_{y_\tau}(t, h; n)^2 = \sum_{i=1, \dots, n} \left(y_\tau \left(t - h + \frac{ih}{n} \right) - y_\tau \left(t - h + \frac{(i-1)h}{n} \right) \right)^2. \quad (27)$$

In line with the logic outlined above equation (15), the realized yield volatility converges, for ever more frequent sampling, towards the underlying realization of the quadratic yield variation process. Equation (12) 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 (27) to approximate the contemporaneous quadratic yield variation process $QV_{y_\tau}(t, h)$ on the left-hand-side of (12).

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. (2005), and these invariably rely on the realized volatility of bond *returns*. As such, it is useful to clarify the link between equation (27) and the realized volatility of the bond return. To this end, we denote the continuously compounded return over the time interval $[t - h, t]$ on a zero-coupon bond with time-to-maturity $\tau = T - t$ by

$$r(t, h, \tau) = p(t, \tau) - p(t - h, \tau), \quad 0 \leq h \leq t \leq T, \quad (28)$$

where $p(t, \tau) \equiv \log(P(t, \tau))$ is the time- t Treasury Bill log-price. Equation (4) yields an expression for the intra-day return during a given date t :

$$r_\tau \left(t - h + \frac{ih}{n}, \frac{h}{n} \right) = -\tau \left(y_\tau \left(t - h + \frac{ih}{n} \right) - y_\tau \left(t - h + \frac{(i-1)h}{n} \right) \right). \quad (29)$$

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

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

It is evident that the qualitative features of the realized yield volatility and realized return volatility series are identical for a given maturity zero-coupon bond and that one may be 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 (29) 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

(a notable exception is 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)). We note, however, that 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 in the long-maturity-bonds 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 close it at 5:00PM ET. This window includes the time of 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 in U.S. Treasuries occurs during these hours, we also capture the activity associated with the price discovery process driven by the aggregation of heterogeneous private information and heterogeneous interpretation of public information through trading in the market, see, e.g., Brandt and Kavajecz (2004) and Pasquariello and Vega (2005).

The GovPX quote frequency for the specific maturities turn out not to be as high as for, e.g., the quotes on the individual stocks in the Dow Jones 30 index. As such, the recent literature on selecting an optimal intra-day 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 realized volatility literature of using a fairly sparse and fixed sampling frequency. A sensible compromise between obtaining improved information regarding the

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 regularly scheduled macroeconomic announcements take place at 8:30AM or 10AM. The statements from the regular Federal Open Market Committee meetings are typically released around 2:15PM. Further, this 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)

strength of the underlying yield movements and adding high-frequency microstructure noise seems to be achieved around the 10-minute sampling interval where the induced serial correlation in the yield change series are relatively 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 so-called ‘bootstrapping’ method (see, e.g., Tuckman (2002)). 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. Hence, we discard those days for which we could not find any trading activity for a period longer than three hours from the sample.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. Further, we compute an average of the daily trading period yield from the 57 intra-day yield observations, so that any i.i.d.-type measurement error becomes immaterial. This approach should remove much concern about the measurement errors for the yield level in our tests of the spanning condition.

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, partly because the database does not include quotes from Cantor Fitzgerald Inc.. Thus, we exclude this security from our sample of intra-day Treasury quotes. 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 construction of intraday-based realized volatility series, it may nonetheless serve as a proxy for a zero-coupon yield that may be used as a regressor in the volatility spanning condition (12). Consequently, such auxiliary daily yields may be used to provide an additional robustness check for our results based on the intraday

⁷Although there are no fixed trading hours for Treasuries, the Bond Market Association (BMA) makes recommendations regarding holiday closes and early closes. Specifically, the BMA often recommends that the market close early (usually at 2 PM) before a holiday, which typically results in low trading activity in 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. Either approach did not change the conclusions discussed below.

⁸The MA coefficient estimates are as follows. For the three-month series: -0.068; six-month series: -0.066; one-year series: -0.050; two-year series: -0.007; five-year series: -0.002; ten-year series: -0.043. 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.

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 coupon-bond yields for bonds sold at par. Hence, prior to analysis we convert these series into zero-coupon yields via the so-called bootstrapping method.

4 Volatility Measures and Forecasts

The quality of the empirical analysis of the affine yield spanning conditions is critically dependent upon the reliability of both the realized volatility measures and the ex-ante (model dependent) volatility forecasts. Since the true (realized) volatility is a latent variable, it is not a trivial matter to confirm the quality of the volatility measures. Hence, 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 clarify the difference between our ex-post realized volatility measures and the more common (ex-ante) volatility forecasts, where the latter, in spirit, are similar to the volatility measures considered in previous studies examining the spanning condition. A second objective is to assess the general properties of the series relative to prior findings in the literature. This should alleviate any concern that our results are driven by some idiosyncratic features of our volatility series. We pay particular attention to the degree of predictability in the yield volatility series as this is of paramount importance for the subsequent empirical evaluation of the affine model class.

4.1 Realized Yield Volatility

We construct the realized yield volatility series in equation (27) 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) volatility measures 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 in the sample.¹¹ Next, we construct an extended measure

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>.

¹¹On the days in which a new on-the-run bond is introduced in the market, we use the 7:30AM on-the-run quotes to fit a cubic spline on the term structure of three- and six-month, one-, two-, five- and ten-year yields. We use such

of daily realized volatility by adding the squared inter-daily yield change to the sum of squared intraday yield changes. This extended measure is then used to estimate the longer term (average) realized yield volatility by simply adding the daily total realized volatility across the full sample. This forms the basis for a direct measure of the average annual realized yield variation. Finally, we rescale each intraday-based realized volatility measure proportionally so that, across the entire sample, the average daily annualized (standard deviation) volatility measure matches that estimated from the total realized volatility series. We emphasize that any 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 realized volatility series, v_{y_t} , for all of the zero-coupon yields that we analyze. Casual inspection reveals a great deal of covariation in the yield volatilities across the maturity spectrum. However, there are also some striking differences. Particularly eye-catching are the extreme outliers. Although these typically are manifest in all the yields, it is noteworthy that they appear distinctly different 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 in the longer-maturity series. On that date, investors appear to have reacted negatively to the concerns of a slowdown in the world economy, a weakening dollar (which had lost 12% against the Japanese Yen in less than two days), and political turmoil associated with the House of Representatives favorable vote to begin an impeachment inquiry into President Bill Clinton. Investors were increasingly attracted to the safety of short-term bonds, pushing their prices up in a volatile trading session.¹²

In contrast, the longer maturity yields' volatility spikes up on, e.g., June 2, 1995, when a dramatic drop in the payroll employment number seems to have raised fears of a recession, sparking a powerful bond market rally that sent prices of longer-term bonds to record highs.¹³ 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.¹⁴

interpolation to measure the yield on the bond that just went off-the-run. This approach compensates for possible liquidity effects that might impact the bond price when it goes off-the-run.

¹²“There has been a dramatic steepening of the yield curve,” said Michael Boss, a bond trader with IBJ Lanston Futures in Chicago, on October 8, 1998, 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.”

¹³“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.”

¹⁴“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.

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. (2005a). It also suggests 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 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 market events indicates that the realized volatilities capture the relative size of the effects adequately in spite of potential measurement error problems.

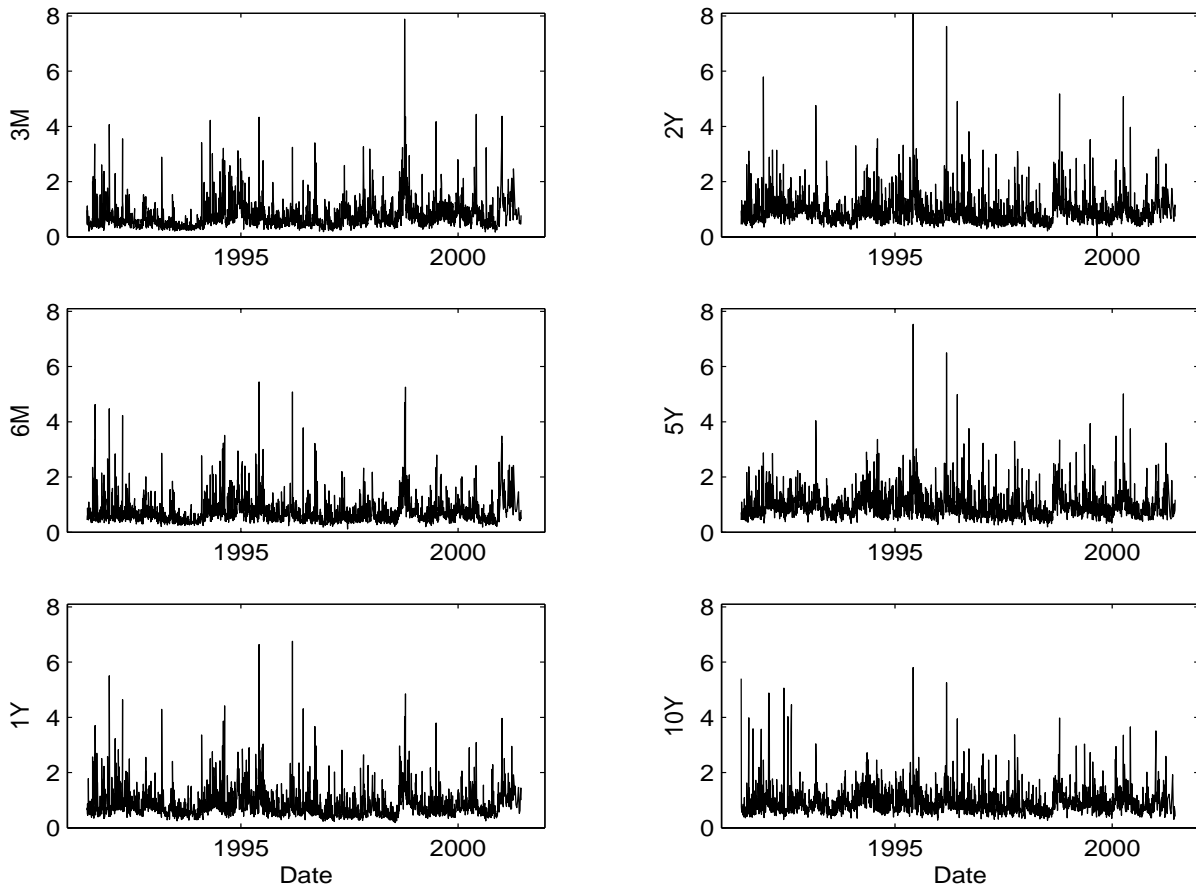


Figure 1: Realized Volatility Series. The plots depict the square root of the rescaled (yearly percentage) realized volatility measures, v_{y_r} , for the three- and six-month, one-, two-, five-, and ten-year maturity yields (sample period: 06/17/1991–06/15/2001).

In Figure 2 we plot the average (sample) yearly 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 annualized standard deviation of the daily changes (from 4pm to 4pm) 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 measures of volatility 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 quite 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.




Figure 2: The Term Structure of Yield Volatility. The continuous line depicts the annualized percentage sample standard deviation of daily changes in zero-coupon yields. The dashed line depicts the yearly percentage average realized volatility v_y . 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).

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 0.35, 0.34, 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)).




Figure 3: Sample Autocorrelations. 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.

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 in combination with the consistent features observed across the 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.

4.2 Volatility Forecasts

In line with the approach in the preceding section, we explore yield volatility forecasts based on intraday yield data as well as daily three-month Treasury bill yield observations. This serves both as a robustness check for the realized yield volatility measures and facilitates direct comparison to existing studies. We first study the properties of standard daily time series forecasts of volatility and later explore forecasts generated directly from the historical realized yield volatility series.

4.2.1 EGARCH-Type Volatility Forecasts

This section reports on the estimation of a daily EGARCH-type model for the short-term zero-coupon yield. The model generates daily volatility forecasts which we compare and contrast to the daily realized volatility series of Section 4.1.

In order to retain the nonparametric flavor of our analysis we focus on the family of SNP densities introduced by Gallant and Nychka (1987). Nonetheless, as is commonly done, we rely on a conditionally Gaussian leading term, designed to capture the bulk of the dependency in the conditional mean and variance of the series, in order to maintain parsimony and avoid excessive overfitting. In particular, we use an ARMA specification for the conditional mean and an EGARCH structure for the conditional variance to capture the heteroskedasticity in the short-rate dynamics. Further, we allow for an additional source of interaction via an interest level effect, i.e., the EGARCH conditional variance term is scaled by $y_t^{2\delta}$. In line with the SNP approach, we allow a squared Hermite polynomial expansion to accommodate any remaining non-normality and time-series dependence in the innovation process. The former consideration is particularly important given the evidence of jump-like outliers in the daily yield series. In sum, we obtain a conditional density for the daily yield of the following form:

$$f_K(y_t|x_t;\xi) = c(y_t, x_t) \frac{\phi(z_t)}{y_{t-1}^\delta \sqrt{h_t}}, \quad (31)$$

where $\phi(\cdot)$ is the standard normal density, $x_t = \{y_1, \dots, y_{t-1}\}$ reflects the information set, ξ is the SNP density parameter vector, and $c(y_t, x_t)$ denotes the SNP density expansion designed to allow for nonnormal innovations and to accommodate remaining conditionally heteroskedastic features in the innovation process. The $c(y_t, x_t)$ part of the density representation turns out to be immaterial for the line of argument we pursue, especially since we end up not incorporating any additional time series dependence into the model through this term. Hence, we suppress the explicit form of this component

here, but for completeness it is given in the Appendix. The scaling of $c(y_t, x_t)$ automatically ensures that the SNP density appropriately integrates to unity. Finally,

$$z_t = \frac{y_t - \mu_t}{y_{t-1}^\delta \sqrt{h_t}}, \quad \mu_t = \phi_0 + \sum_{i=1}^s \phi_i y_{t-i} + \sum_{i=1}^u \zeta_i (y_{t-i} - \mu_{t-i}),$$

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

As in Andersen and Lund (1997) and Andersen et al. (2002, 2004), $b(z)$ is used merely as a numerical device to provide a smooth, twice-differentiable approximation to the absolute value operator in the EGARCH variance equation. More in-depth discussion of the implementation of the SNP approach within a similar setting may also be found in those papers.

Because of its prominence in the literature, we focus on the three-month Treasury bill yield series.¹⁵ It is well known that a long sample period is required to pin down the moments of a persistent series like the three-month rate and its volatility. Hence, we use the data series from July 1, 1983, to June 30, 2005, for a total of 5,498 observations. Earlier data are also available, but we purposely avoid 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 period used for estimation of the SNP density (31), while the last two columns represent the June 17, 1991–June 15, 2001, sample period of our realized volatility series. The basic summary statistics are quite similar across the two sample periods. In particular, the interest rate levels are nearly identical and the kurtosis of the yield changes is exceedingly high for both samples. However, it is apparent that the 1991-2001 period is characterized by a comparatively low level of yield volatility.

Table 1: Summary statistics for U.S. three-month T-Bill yield data. All figures are computed using daily yields expressed in percentage form on a yearly basis.

	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

We estimate the SNP density in (31) by (quasi-)maximum likelihood (QML). The Bayesian (BIC)

¹⁵We 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.

and Hannan-Quinn (H-Q) information criteria are used to guide model selection while additional information regarding the proper 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 an ARMA(6,1)-Level-EGARCH(2,1)-Kz(8)-Kx(0) specification.¹⁶ The Kx(0) representation implies that we, as alluded to earlier, found no need for additional conditional heteroskedasticity beyond what is captured by the EGARCH-Level representation of the conditionally Gaussian leading term. On the other hand, the Kz(8) term accommodates strong departures from conditional normality of the standardized innovations.

High-order ARMA and EGARCH terms are needed to capture the conditional mean and volatility dynamics of the three-month yield series. Further, our analysis points towards extremely persistent first- and second-order conditional moments, even if the absolute value of the roots of the ARMA and EGARCH autoregressive polynomials remain outside the unit circle, i.e., the 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, while the volatility roots fall within the range that produces an autocorrelation structure which closely mimics the hyperbolic decay associated with long-memory-type persistence.

The EGARCH-Level-SNP density (31) can be used to construct estimates for the one-day-ahead (conditional) yield variance:

$$E(y_t|x_t;\xi) = \int y f_K(y|x_t;\xi) dy \quad (32)$$

$$V(y_t|x_t;\xi) = \int (y - E(y|x_t;\xi))^2 f_K(y|x_t;\xi) dy, \quad (33)$$

where y_t is the time- t realization of the yield and $x_t = \{y_1, \dots, y_{t-1}\}$.

Figure 4 depicts the one-day-ahead volatility forecasts obtained from the EGARCH-Level-SNP model (i.e., $\sqrt{V(y_t|x_t;\xi)}$) along with the corresponding daily realized volatility series, $v_y(t)$. Hence, this is literally a plot of volatility forecasts versus subsequent realizations (proxied by $v_y(t)$) in the spirit of equation (15). 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 coherence between the two series constructed from distinct data sources, as the long-run movements in the yield volatility forecasts clearly are related to corresponding shifts in the overall intensity of the realized volatility measures. The correlation among the two series amounts to 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.

¹⁶Estimation results are available from the authors upon request.

(2005). It should be kept in mind, as also 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., the equity and currency markets. Finally, we also note that the sample means of the two series are very close at 0.0426% and 0.0486% per day for the SNP forecast and the realized volatility, respectively.




Figure 4: U.S. three-month T-Bill yield volatility: daily realized volatility (i.e., $v_y(t)$) versus one-day-ahead volatility forecasts based on the SNP model (i.e., $\sqrt{V(y_t|x_t;\xi)}$). Sample period: 06/17/1991–06/15/2001.

In sum, we conclude that our short-term yield realized volatility series is coherent with the volatility forecasts obtained from the extended EGARCH model estimated at the daily frequency. The overall relationship between the forecasts and the subsequent volatility realizations is qualitatively analogous to what has been found in prior empirical work and it is consistent with the theoretical inquiry regarding this relationship in Andersen, Bollerslev, and Meddahi (2004, 2005). Hence, the EGARCH-Level forecasts appear well calibrated, which in turn suggests that the daily realized yield volatility series provide informative and useful measures of the underlying quadratic yield variation realizations.

4.2.2 Realized Volatility Component Model Forecasts

The availability of daily model-free realized volatility measures facilitates forecast procedures that rely on standard time-series models built directly on the history of ex-post volatilities. This approach provides a simple and powerful alternative to the EGARCH-type forecasts explored above. A convenient model along these lines is the HAR-RV form studied in Andersen et al. (2005), Corsi (2003),

and Müller et al. (1997):

$$\frac{v_{y_\tau}(t+h, h)^2}{h} = \beta_0 + \beta_D v_{y_\tau}(t, 1)^2 + \beta_W \frac{v_{y_\tau}(t, 5)^2}{5} + \beta_M \frac{v_{y_\tau}(t, 21)^2}{21} + \varepsilon(t+h), \quad (34)$$

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 . 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 previously documented in Section 4.1.

In Table 2, we report ordinary least squares (OLS) estimation results for the HAR-RV model (34). We note that the estimates for β_D , β_W , and β_M confirm the existence of highly persistent volatility dependence. Moreover, the relative importance of the monthly volatility component increases from the daily to the weekly and monthly regressions, and vice-versa for the daily volatility component. Finally, 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 the slightly longer horizons.

Nonetheless, there is good coherence between the HAR-RV forecasts and the ex-post volatility realizations. This is illustrated in Table 3, where we report sample correlations between the square-root of the HAR-RV forecasts, $\hat{v}_{y_\tau}(t, h)$, and the realized volatility series, $v_{y_\tau}(t, h)$, for different forecasting horizons, $h = 1, 5$, and 21. This point is further 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. As seen earlier in Figure 4, the long-run movements in the forecasts are clearly related to corresponding shifts in the overall intensity of the realized volatility measures.

In sum, we conclude that our HAR-RV forecasts are informative about the subsequent volatility realizations. Moreover, the relation between a-priori and ex-post estimates is consistent with the results in Section 4.2.1 and it is in line with the predictions of the theoretical literature (see, e.g., Andersen, Bollerslev, and Meddahi (2004, 2005)). This evidence lends some credibility to the empirical analysis of the affine spanning conditions below.

Table 2: Daily, Weekly, and Monthly HAR-RV Regressions. We report OLS estimates for

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

where the dependent variable is the daily ($h = 1$), overlapping weekly ($h = 5$), and overlapping monthly ($h = 21$) realized volatility for zero-coupon yields with maturity $\tau = 3M, 6M, 1Y, 2Y, 5Y$, and $10Y$. Realized volatilities are constructed by using 10-minute yield-changes sampled from 06/17/1991 to 06/15/2001. Standard errors estimates are robust with respect to both autocorrelation and heteroskedasticity. Coefficient t -ratios are in square brackets.

	$v_{y\tau}^2, \tau = 3M$			$v_{y\tau}^2, \tau = 6M$			$v_{y\tau}^2, \tau = 1Y$		
h	1	5	21	1	5	21	1	5	21
β_0	0.0012 [4.03]	0.0015 [4.44]	0.0021 [6.87]	0.0012 [4.79]	0.0014 [5.35]	0.0019 [7.72]	0.0018 [6.06]	0.0020 [6.67]	0.0025 [8.07]
β_D	0.0576 [0.92]	0.0442 [2.39]	0.0248 [2.07]	-0.0084 [-0.30]	0.0241 [1.68]	0.0125 [1.84]	0.0266 [1.10]	0.0131 [1.07]	0.0072 [1.20]
β_W	0.2263 [3.35]	0.1765 [3.37]	0.1349 [2.36]	0.2460 [2.78]	0.1652 [3.04]	0.1338 [2.88]	0.0979 [1.74]	0.0857 [1.98]	0.0844 [2.41]
β_M	0.3658 [2.50]	0.3456 [2.77]	0.2370 [2.53]	0.3715 [4.29]	0.3676 [4.14]	0.2439 [2.89]	0.4339 [5.50]	0.4247 [5.20]	0.3116 [4.17]
R_{Adj}^2	8.30%	17.83%	17.99%	5.69%	15.73%	16.85%	3.65%	11.83%	17.03%
	$v_{y\tau}^2, \tau = 2Y$			$v_{y\tau}^2, \tau = 5Y$			$v_{y\tau}^2, \tau = 10Y$		
h	1	5	21	1	5	21	1	5	21
β_0	0.0024 [6.91]	0.0026 [7.51]	0.0033 [9.55]	0.0023 [8.10]	0.0025 [8.77]	0.0033 [11.15]	0.0025 [8.17]	0.0026 [9.03]	0.0033 [11.42]
β_D	0.0291 [1.53]	0.0055 [0.57]	0.0059 [1.10]	0.0470 [1.44]	0.0154 [1.43]	0.0091 [1.60]	0.0185 [0.57]	0.0170 [1.77]	0.0047 [0.99]
β_W	0.0657 [1.32]	0.0936 [2.25]	0.0919 [2.53]	0.1313 [2.65]	0.1398 [4.01]	0.0980 [3.01]	0.1713 [2.27]	0.1206 [2.43]	0.0755 [2.99]
β_M	0.4021 [4.16]	0.3638 [3.95]	0.2175 [3.14]	0.3553 [4.48]	0.3328 [4.65]	0.2316 [4.22]	0.2431 [3.37]	0.2622 [3.66]	0.1825 [3.07]
R_{Adj}^2	2.53%	8.88%	11.20%	3.77%	11.09%	12.85%	2.23%	6.96%	7.83%

Table 3: Sample correlation between the square-root of the HAR-RV forecasts, $\hat{v}_{y_\tau}(t, h)$, and the realized volatility series, $v_{y_\tau}(t, h)$, $\tau = 3M, 6M, 1Y, 2Y, 5Y$, and $10Y$. The forecasting horizons are one day ($h = 1$), one week ($h = 5$), and one month ($h = 21$), respectively.

h	3M	6M	1Y	2Y	5Y	10Y
1	44.37%	36.37%	34.01%	30.60%	33.19%	25.92%
5	50.82%	47.26%	44.37%	40.11%	42.30%	33.36%
21	47.94%	45.63%	46.02%	38.74%	40.93%	32.51%

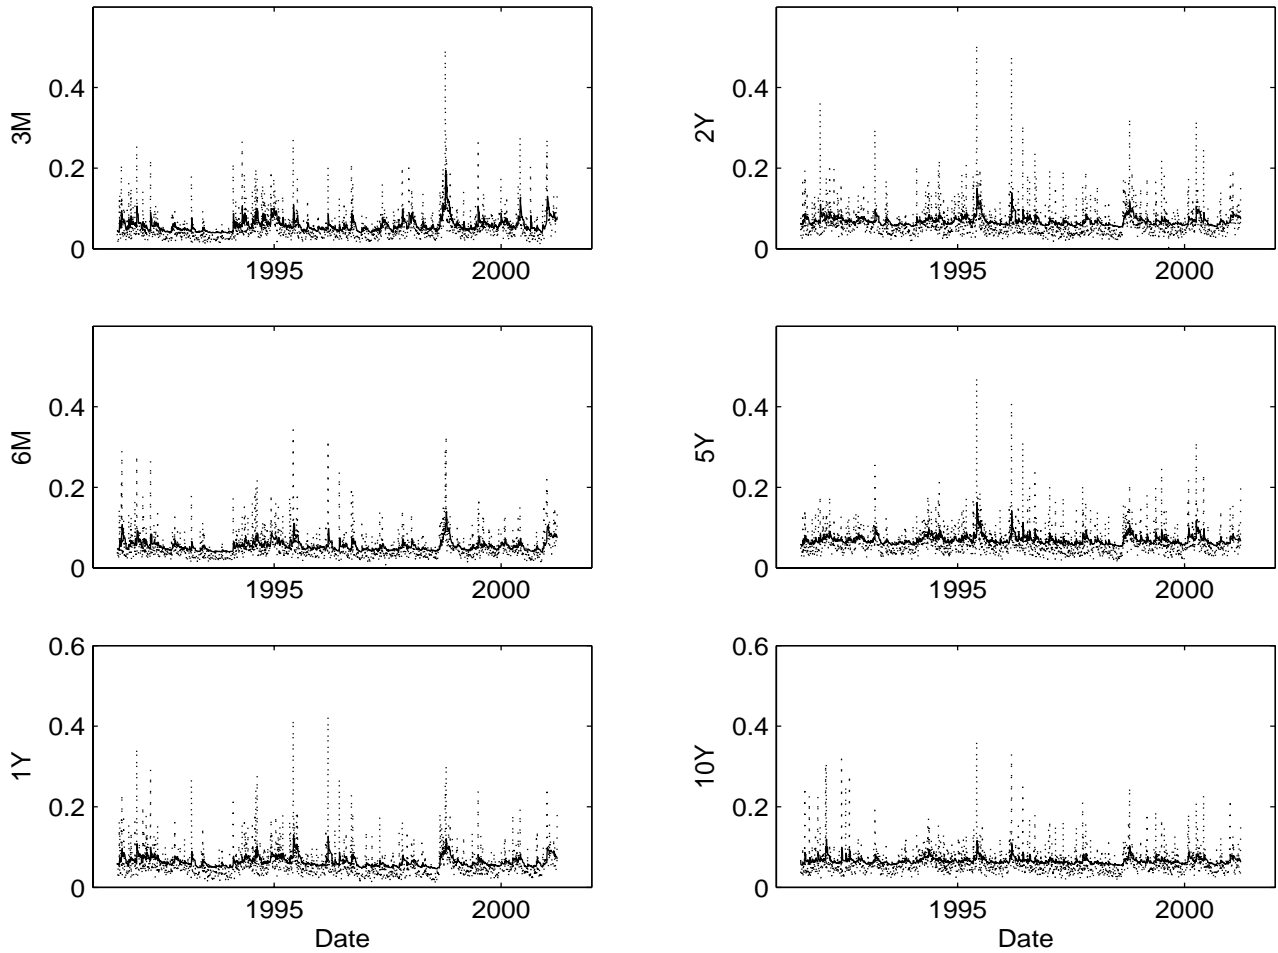


Figure 5: One-day-ahead HAR-RV forecasts, $\hat{v}_{y_\tau}(t, 1)$, and realized volatility series, $v_{y_\tau}(t, 1)$ (daily percentage scaling). Sample period: 07/17/1991—03/27/2001.

5 Evidence on the Affine Spanning Conditions

This section presents our main empirical findings. First, in Section 5.1 we focus on the empirical implications of multi-factor affine diffusion models previously summarized in Sections 2.1-2.3. Then, in Section 5.2 we extend our analysis to a jump-diffusion setting, as discussed in Section 2.4.

5.1 Affine Diffusion Models

5.1.1 Can Bonds Hedge Volatility Risk?

The fundamental yield spanning condition in equation (12) should be satisfied by any affine (and quadratic) diffusion model. 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 important advantage of reducing possible problems due to the presence of any non-systematic measurement error in the zero-coupon yields.

To alleviate the multi-collinearity problem, we extract orthogonal principal components from the panel of average daily yields. As usual, we find that the first three components correlate very highly with empirical measures of the yield level, the slope of the yield curve and the curvature of the yield curve. Although these first three components, as commonly found, capture almost all of the yields' variation over time, we include all six of them in our regressions to make sure that our right-hand-side variables capture the entire yields' variation. This approach provides some robustness towards also including yields that may be relevant for the quadratic model class, which may involve a substantial amount of affine yield factors. We denote the principal components, or factors, by PC_j , $j = 1, \dots, 6$. We use ordinary least squares to estimate the model

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

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 5.1.2 below.

Panel A of Table 4 reports results based on the full sample of daily data from 06/17/1991 to 06/15/2001. 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 only 3%. 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. In further interpreting this result, it is noteworthy that the coefficients on the first three principal components are insignificant, while higher-order factors enter significantly in the regression. 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 quite standard in the recent literature. The evidence is also at odds with the findings of, e.g., Litterman, Scheinkman, and Weiss (1991), who conclude that interest rate volatility is linked to the curvature factor.

The results from estimating the model (35) with the realized volatility of longer-maturity yields as the dependent variable are given in Table 4, Panel A. The explanatory power deteriorates considerably compared to the case involving the shorter rate yield volatility. In particular, for maturities longer than two years, the R^2 coefficients are essentially zero. This evidence has striking implications. In particular, the failings of the affine diffusive models do not appear to be driven by some idiosyncratic features of the yield dynamics at the short end of the maturity spectrum. These findings strengthen the observations in, e.g., Ahn et al. (2003) and Dai and Singleton (2000)) that the standard affine models have difficulty in accommodating the volatility dynamics of yields at different maturities. Since the spanning condition we explore is applicable across the entire model class, we conclude that there is no potential for *any* model of this type to fit the observed yield volatility structure. Instead, there seems to be a need for multiple stochastic volatility factors that are not directly tied to the yield cross-section. Such specifications will either fall outside the affine diffusion model class or incorporate multiple volatility factors subject to the USV restriction.

We also estimate the covariance matrix of the residuals from the six regressions reported in Panel A and we use it to perform a principal component analysis. This exercise is in the spirit of Collin-Dufresne and Goldstein (2002), who perform a principal component analysis of the residuals from the 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 the more powerful one of the two. Nonetheless, consistent with their results, we find that the first principal component explains more than 77% of the variation in the model residuals (Panel B). This evidence confirms that the limited explanatory power of the regressions in Panel A is not due to noisy data. In contrast, there appears to be a dominant common factor that drives the volatility of zero-coupon yields at different maturities and is unrelated to the level of the yields.

5.1.2 Robustness of the Results

As a robustness check, we estimate the model over many different sub-samples and for the yield variation measured over different horizons. For each sub-sample, we fit six distinct regressions corresponding to having the realized zero-coupon yield volatility for a specific maturity as the dependent

variable. However, to conserve space in Table 5 we only report estimates for the regressions based on the three-month, two-year, and ten-year realized volatility 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. As noted in Piazzesi (2005), starting with the first Federal Open Market Committee (FOMC) meeting of 1994 the Fed has been announcing 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 arguably 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 perform a principal component analysis of the residuals from the regressions estimated over this sample period. As reported in Panel B, we confirm that a single dominant factor explains most of the variation in our realized volatility series that is not associated with changes in the yields. Virtually identical results were found for the related sample period, 02/01/95 to 12/29/2000, 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 what we have found for the full sample. However, bond yields have virtually no explanatory power towards the realized volatility of longer-maturity rates. Further, contrary to the findings for the full sample, the loading on the first principal component is generally significant in both sub-samples. However, it is positive during the first sample and negative during the second. Since this factor reflects shifts in yield levels, it suggests that the so-called ‘level’ effect in volatility may not be stable over time. Further, the lack of robustness in coefficient estimates over time suggests that the improvement in the regressions R^2 reflects in-sample over-fitting.¹⁷

The analysis in Section 4 suggests that our realized volatility series are highly informative regarding the true underlying yield variation. In particular, there is no indication that measurement errors associated with the use of ten-minute yield changes or the presence of microstructure noise seriously impair the quality of these volatility proxies. Nevertheless, there may be residual concern regarding the impact of measurement errors in our realized volatility measures. Similarly, we would also like to assess whether observation errors might contaminate the yield proxies on the right hand side of the regressions and cause artificial deterioration in the cohesion among the yield level and the yield

¹⁷An alternative interpretation is that the switch in sign corresponds to a change in regime within the affine model. More generally, in such affine diffusive regime-switching models shorter sample periods would correspond more closely to the presence of one dominant (persistent) regime over the period, which would result in a better fit over those shorter periods. In order to informally test for this possibility, we split the full sample in five two-year periods and ran the regressions for each sub-sample. Although the explanatory power was considerably higher for one of the shorter 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. Hence, we still did not detect any sign of a reliable link between the yield volatility and the yield cross-section. These results are available upon request.

volatility variation. A simple robustness check involves aggregation to a lower sampling frequency, since random measurement errors will tend to diversify and be less important relative to the accumulated signals in the yield changes and yield variation measures. Consequently, we test the spanning condition (12) at the weekly frequency. For that purpose, we aggregate the daily realized volatility measures in non-overlapping weekly series. Similarly, we construct a weekly average of the daily yields and extract principal component from the panel of such series. We use this data to fit model (35) by OLS. The results, in Table 6, Panel A, are consistent with our previous findings. The small increase in overall explanatory power is consistent with the substantially lower yield variability at the weekly compared to the daily frequency. Nonetheless, the explanatory power remains very low and virtually non-existent for the longer maturity series. Moreover, the residual variation of the yields continues to have a strong common component as documented in Panel B of Table 6. Hence, the weekly yield variation also displays a strong covariation across the maturity spectrum which is unrelated to the yield levels. Finally, we have confirmed that these results carry over to the monthly frequency.¹⁸

Finally, we conduct a second check to verify that our results are robust to the proxy 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 Federal Reserve 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, 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 (35). Here, the dependent variables are the usual six realized volatility measures constructed from intra-daily yield observations. The results, in Table 7, are consistent with those previously reported in Table 4. Most importantly, although the R^2 coefficients for these regressions improve marginally, we still find little or no evidence that a portfolio of bonds can span (hedge) volatility risk. This evidence makes our earlier conclusions more powerful. First, it shows that our previous findings were not hampered by the fact that we were not using the yields on the thirty-year bond. Second, by increasing the number of yields in the right-hand-side of our regressions we effectively provide evidence suggesting that the spanning condition is violated even in higher-order affine (or quadratic) models.

¹⁸Since the findings are very similar to those for the weekly horizon we do not report them here but they are available upon request.

Table 4: OLS Regressions Results. We report OLS estimates for the model

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

where $\tau = 3M, 6M, 1Y, 2Y, 5Y,$ and $10Y$, and $PC_j, j = 1, \dots, 6$, are the six principal components extracted from the panel of six zero-coupon yields. Standard errors estimates are robust with respect to both autocorrelation and heteroskedasticity. Coefficient t -ratios are in square brackets.

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]	
$v_{y_\tau}^2, \tau = 3M$	0.8768 [12.70]	0.0430 [0.99]	-0.1450 [-1.46]	-0.4276 [-1.48]	2.6139 [2.10]	-4.6304 [-3.00]	3.2651 [2.36]	3.12%
$v_{y_\tau}^2, \tau = 6M$	0.7818 [14.75]	0.0273 [0.90]	-0.0134 [-0.19]	0.4079 [1.44]	2.2262 [2.57]	-2.3046 [-2.55]	3.4137 [2.62]	2.00%
$v_{y_\tau}^2, \tau = 1Y$	1.0385 [16.54]	0.0285 [0.91]	0.0614 [0.86]	0.6734 [2.04]	1.9632 [2.11]	-1.9673 [-1.99]	4.4724 [2.51]	1.46%
$v_{y_\tau}^2, \tau = 2Y$	1.2134 [16.71]	0.0094 [0.29]	0.0078 [0.11]	0.7082 [1.80]	1.4372 [1.45]	-0.8654 [-0.79]	3.9283 [1.75]	0.56%
$v_{y_\tau}^2, \tau = 5Y$	1.2274 [18.85]	0.0115 [0.40]	0.0019 [0.03]	0.1663 [0.46]	1.2741 [1.54]	-0.7163 [-0.66]	3.2514 [1.75]	0.27%
$v_{y_\tau}^2, \tau = 10Y$	1.1106 [21.88]	0.0016 [0.06]	0.0749 [1.33]	0.3054 [1.10]	0.8373 [1.20]	-0.7222 [-0.83]	3.1961 [2.03]	0.49%

Panel A: Daily observations from 06/17/1991 to 06/15/2001.

1st	2nd	3rd	4th	5th	6th
77.33%	11.83%	6.26%	2.69%	1.11%	0.78%

Panel B: Percentage of the variance explained by the principal components extracted from the OLS regressions' residuals.

Table 5: Robustness Checks. For different sub-samples, we report OLS estimates of the model

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

where $\tau = 3M, 6M, 1Y, 2Y, 5Y,$ and $10Y$, and $PC_j, j = 1, \dots, 6$, are the six principal components extracted from the panel of six zero-coupon yields. Standard errors estimates are robust with respect to both autocorrelation and heteroskedasticity. Coefficient t -ratios are in square brackets.

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]	
$v_{y_\tau}^2, \tau = 3M$	0.9515 [12.41]	-0.1032 [-1.11]	0.0070 [0.09]	-0.8208 [-1.66]	4.7671 [2.43]	-4.6604 [-3.11]	4.1679 [2.35]	4.60%
$v_{y_\tau}^2, \tau = 2Y$	1.2404 [14.10]	-0.0302 [-0.61]	-0.0302 [-0.32]	1.3084 [1.57]	2.1761 [1.94]	-1.2886 [-1.28]	5.7523 [1.75]	0.99%
$v_{y_\tau}^2, \tau = 10Y$	1.0855 [20.26]	-0.0204 [-0.56]	0.1331 [1.95]	0.5950 [1.34]	0.4005 [0.48]	-1.1702 [-1.35]	2.6763 [1.40]	0.92%
Panel A: Daily observations from 01/05/1994 to 06/27/2001.								
	1st	2nd	3rd	4th	5th	6th		
	82.52%	12.30%	2.62%	1.22%	0.78%	0.56%		
Panel B: Percentage of the variance explained by the principal components extracted from the OLS regressions' residuals (01/05/1994—06/27/2001).								
$v_{y_\tau}^2, \tau = 3M$	0.7475 [15.54]	0.1373 [6.97]	0.1451 [2.40]	-0.6135 [-2.53]	2.8796 [3.69]	-0.2054 [-0.14]	5.3885 [2.51]	7.58%
$v_{y_\tau}^2, \tau = 2Y$	1.3019 [12.60]	0.0854 [2.65]	-0.0627 [-0.37]	0.5212 [1.24]	0.5744 [0.50]	3.3419 [1.05]	6.3412 [1.58]	0.55%
$v_{y_\tau}^2, \tau = 10Y$	1.1892 [15.79]	0.0438 [1.56]	0.0299 [0.28]	0.2560 [0.74]	1.4946 [1.63]	1.6768 [0.65]	4.0235 [1.46]	0.24%
Panel C: Daily observations from 06/17/1991 to 06/14/1996.								
$v_{y_\tau}^2, \tau = 3M$	0.9942 [9.37]	-0.3326 [-2.12]	-0.2181 [-1.62]	-0.0822 [-0.11]	1.1435 [0.61]	-6.7605 [-2.31]	1.8958 [0.51]	5.24%
$v_{y_\tau}^2, \tau = 2Y$	1.1194 [13.42]	-0.1888 [-2.27]	-0.1680 [-1.17]	0.5320 [0.72]	0.1181 [0.11]	-1.0952 [-0.77]	2.6106 [0.98]	1.80%
$v_{y_\tau}^2, \tau = 10Y$	1.0265 [17.28]	-0.0747 [-1.14]	0.0045 [0.05]	0.0265 [0.05]	-0.5060 [-0.60]	-0.0770 [-0.06]	2.8336 [1.67]	0.39%
Panel D: Daily observations from 06/17/1996 to 06/15/2001.								

Table 6: Robustness Checks. We report OLS estimates for the model

$$v_{y_\tau}^2(t, 5) = \beta_0 + \sum_{j=1}^6 \beta_j PC_j(t, 5) + \varepsilon(t).$$

The dependent variable is the *annualized* non-overlapping *weekly* realized volatility 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 principal components extracted from non-overlapping weekly averages of the zero-coupon yields. Standard errors estimates are robust with respect to both autocorrelation and heteroskedasticity. Coefficient t -ratios are in square brackets.

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]	
$v_{y_\tau}^2, \tau = 3M$	0.9051 [10.29]	0.0483 [0.90]	-0.1491 [-1.18]	-0.3938 [-1.23]	2.6805 [1.67]	-5.1080 [-2.74]	3.7347 [2.42]	8.84%
$v_{y_\tau}^2, \tau = 6M$	0.8069 [12.08]	0.0289 [0.73]	-0.0140 [-0.16]	0.4492 [1.41]	2.2826 [2.15]	-2.4539 [-2.23]	3.3300 [2.36]	6.25%
$v_{y_\tau}^2, \tau = 1Y$	1.0710 [13.43]	0.0310 [0.75]	0.0623 [0.65]	0.7299 [1.86]	1.9733 [1.64]	-2.4555 [-1.96]	4.0574 [2.21]	4.82%
$v_{y_\tau}^2, \tau = 2Y$	1.2521 [13.88]	0.0110 [0.26]	0.0069 [0.08]	0.7743 [1.58]	1.2623 [1.03]	-1.1761 [-0.93]	3.4252 [1.39]	1.61%
$v_{y_\tau}^2, \tau = 5Y$	1.2665 [15.30]	0.0133 [0.35]	0.0031 [0.04]	0.2084 [0.44]	1.2706 [1.22]	-1.2214 [-0.93]	3.0622 [1.33]	0.49%
$v_{y_\tau}^2, \tau = 10Y$	1.1467 [19.15]	0.0014 [0.05]	0.0778 [1.25]	0.3433 [1.04]	0.9491 [1.15]	-1.0355 [-0.93]	3.1208 [1.84]	1.87%

Panel A: Non-overlapping *weekly* observations from 06/17/1991 to 06/15/2001.

1st	2nd	3rd	4th	5th	6th
77.79%	13.47%	5.09%	2.14%	0.91%	0.60%

Panel B: Percentage of the variance explained by the principal components extracted from the OLS regressions' residuals.

Table 7: Robustness Checks. We report OLS estimates for the model

$$v_{y\tau}^2(t, 1) = \beta_0 + \sum_{j=1}^9 \beta_j PC_j(t, 1) + \varepsilon(t).$$

The dependent variable is the daily realized volatility for zero-coupon yields with maturity $\tau = 3M, 6M, 1Y, 2Y, 5Y, 7Y, 10Y,$ and $30Y$. Standard errors estimates are robust with respect to both 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 $30Y$. Standard errors estimates are robust with respect to both autocorrelation and heteroskedasticity. Coefficient t -ratios are in square brackets.

Dep. variable	β_0 [t -ratio]	β_1 [t -ratio]	β_2 [t -ratio]	β_3 [t -ratio]	β_4 [t -ratio]	β_5 [t -ratio]	β_6 [t -ratio]	β_7 [t -ratio]	β_8 [t -ratio]	β_9 [t -ratio]	R_{Adj}^2
$v_{y\tau}^2, \tau = 3M$	0.8768 [13.37]	0.0233 [0.56]	-0.1128 [-2.04]	-0.2035 [-1.00]	0.3165 [0.73]	-4.3109 [-2.86]	5.4624 [3.30]	0.4657 [0.38]	-2.6926 [-1.59]	10.7043 [1.51]	5.06%
$v_{y\tau}^2, \tau = 6M$	0.7818 [15.08]	0.0226 [0.75]	-0.0093 [-0.22]	0.3127 [1.60]	-0.0467 [-0.11]	-2.9762 [-3.09]	2.2661 [2.09]	2.4131 [2.36]	-2.2262 [-1.44]	7.7202 [1.73]	2.94%
$v_{y\tau}^2, \tau = 1Y$	1.0385 [16.79]	0.0312 [1.05]	0.0437 [1.02]	0.4572 [1.83]	-0.1724 [-0.39]	-2.9171 [-2.92]	1.1062 [0.88]	1.5974 [1.24]	-4.0811 [-2.11]	9.0061 [2.15]	2.09%
$v_{y\tau}^2, \tau = 2Y$	1.2134 [16.78]	0.0093 [0.32]	0.0153 [0.35]	0.5003 [1.56]	-0.0860 [-0.20]	-1.6860 [-1.68]	-0.0377 [-0.03]	1.8720 [1.17]	-3.6421 [-1.61]	8.9498 [2.31]	0.92%
$v_{y\tau}^2, \tau = 5Y$	1.2274 [18.94]	0.0113 [0.44]	0.0041 [0.11]	0.1381 [0.47]	0.2657 [0.71]	-1.4683 [-1.75]	-0.3513 [-0.27]	1.5954 [1.19]	-2.9563 [-1.51]	7.2892 [2.73]	0.66%
$v_{y\tau}^2, \tau = 10Y$	1.1106 [22.61]	0.0117 [0.53]	0.0534 [1.48]	0.0747 [0.38]	-0.4605 [-1.21]	-1.7815 [-2.43]	0.0667 [0.06]	2.1314 [1.98]	-2.5040 [-1.51]	8.8369 [3.20]	1.56%

Panel A: Daily observations from 06/17/1991 to 06/15/2001.

	1st	2nd	3rd	4th	5th	6th
	77.36%	11.78%	6.25%	2.71%	1.11%	0.79%

Panel B: Percentage of the variance explained by the principal components extracted from the OLS regressions' residuals.

5.1.3 Correlation in Realized Volatility Innovations

Here, we examine the restrictions that the $A_1(N)$ affine term structure model imposes on the quadratic variation of bond yields. In order to check whether the condition in equation (21) is consistent with the empirical evidence, we use the HAR-RV model discussed in Section 4.2.2 to estimate the innovations in realized volatility series. Specifically, we construct one-day-ahead HAR-RV forecasts and compute six innovations series by:

$$\hat{\varepsilon}_{y_\tau}(t) = v_{y_\tau}(t, 1)^2 - E[v_{y_\tau}(t, 1)^2 | t - 1], \quad \tau = 3M, \dots, 10Y. \quad (36)$$

After removing the predictable component from the realized volatility series, we compute the pairwise sample correlation between the volatility innovation series. The results are reported in Table 8. 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 innovations is around 40%, a number significantly below unity. The correlations for nearby maturities are larger, but they 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).

Table 8: Sample Correlations in Realized Volatility Innovations. We filter out the predictable component in the realized volatility series for yields with maturities of three and six months, one, two, five, and ten years. Below, we report the pairwise percentage sample correlations between the series of realized volatility innovations. Standard errors are in round brackets.

	3M	6M	1Y	2Y	5Y
6M	70.21 (12.45)				
1Y	62.75 (17.82)	89.21 (8.43)			
2Y	50.67 (19.30)	79.92 (13.29)	92.11 (7.56)		
5Y	44.87 (18.74)	70.84 (15.61)	83.98 (11.13)	93.71 (6.45)	
10Y	39.11 (13.51)	58.25 (9.32)	64.42 (7.24)	71.63 (4.33)	75.18 (3.21)

This evidence sheds 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 appears to be inconsistent with the cross-sectional properties of bond yields.¹⁹ Of course, our evidence in the preceding section 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 derivatives prices seem to provide a promising basis for future research into these issues.

5.2 Extension to Affine Jump-Diffusion Models

As explained in Section 2.4, the spanning condition (12) for the diffusive case no longer applies in the presence of jumps. Instead, the appropriate yield spanning condition for the affine jump-diffusion model is given by equation (26). Of course, under relatively weak auxiliary assumptions, this restriction is also valid for the conditional expectation evaluated under the P measure, as discussed below equation (26), and this is the version that we focus on here.

The main difference between the two spanning conditions is that equation (26) represents forward looking expectations rather than a realization-by-realization linkage. As a result, powerful tests of the jump-diffusive variant of the spanning condition hinge upon the availability of reliable forecasts for the quadratic yield variation. To this end, we rely mainly on the HAR-RV forecasts discussed in Section 4.2.2, but for the three-month yield series we also consider the EGARCH one-day-ahead forecast discussed in Section 4.2.1. Since condition (26) is strictly valid only when the expected yield change is negligible, we concentrate on conditional variability forecasts for the daily horizon, although we report results based on one-week- and one-month-ahead forecasts as a robustness check. For all these horizons, the expected yield changes are small and the approximation error associated with the spanning condition should be trivial.

In addition, the implementation of the empirical tests requires data on a set of contemporaneous yields, representing the right-hand-side regressors of the spanning restriction (26). To this end, we could either use end-of-the-day yields’ measures (with the potential advantage of providing more

¹⁹Ahn 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)).

current informative regarding future yield changes) or a daily average of the past intra-daily yields (likely measured with less error). Below, we report results based on the latter approach, but we have confirmed that we obtain similar results using only yield observations sampled towards the end of the trading day (2-5PM). This provides some further assurance that measurement error issues are not contaminating our results.

In Table 9, we report OLS estimation results for

$$E[v_{y_r}(t, h)^2 | t - h] = \beta_0 + \sum_{j=1}^6 \beta_j PC_j(t - h) + \varepsilon(t), \quad (37)$$

where $PC_j(t - h)$, $j = 1, \dots, 6$, are the six principal components extracted from the sample of average daily yields.

Panel A presents results for one-day-ahead volatility forecasts ($h = 1$). It is evident that the explanatory power is higher than what we found in Section 5.1.1. Specifically, the R^2 coefficient for the three-month series is approximately 10% when the dependent variable is the EGARCH-type forecast, and 11% in the case of the HAR-RV measure, thus clearly exceeding the 3% R^2 found when examining the contemporaneous yield-variation relation (see Table 4, Panel A). This improvement is not surprising as the volatility forecasts are smooth ex-ante measures of (expected) volatility, whereas the ex-post realized volatility measures also reflect the unpredictable yield innovations and, thus, fluctuate much more. As such, condition (26) is considerably less stringent than version (12) from Section 5.1.1, simply because there is less sample variation for the cross-section of yields to rationalize. On the other hand, the 10-11% R^2 is still very small given the fact that the yield cross-section should provide the optimal forecast for future realized volatility. In fact, this finding is only compatible with theory if the yield cross-section outperforms the time-series forecasts of yield volatility, as discussed in detail below equation (17). However, comparing the performance of the HAR-RV forecasts reported in Table 2 with those of the yield cross-section provided in Tables 4 and 6 for the three-month yield series reveals that the former have considerably more explanatory power for future realized yield variation both at the daily and weekly horizon.

Moving to results for the longer maturities in Table 9, we find that the explanatory power of these regressions deteriorates significantly when we consider HAR-RV forecasts of longer-maturity yield series. For instance, the R^2 coefficient for two-, five-, and ten-year yields are approximately 4.4%, 2.2%, and 5.5%, respectively. Moreover, it is again evident from our prior results that the yield volatility forecasts constructed from the yield cross-section underperforms those of the HAR-RV model for these maturities. Finally, it is telling that Panel B documents a strong covariation in the components of the HAR-RV forecasts which is not explained by the yield cross-section. Hence, the univariate time series models capture systematic comovements across the yield volatilities that cannot be extracted from the current yield levels. In other words, there is predictable yield volatility variation which is unrelated to the yield cross-section. Likewise, Section 5.1 documents that the realized yield variation—

which is heavily influenced by unexpected yield innovations—cannot be explained by (unexpected) changes in the yield levels. We have further verified that identical conclusions hold for the weekly and monthly forecast horizons. Overall, we conclude that there is very compelling evidence against the yield volatility spanning conditions implied by any standard model within the affine class.

Table 9: OLS Regressions Results. We report OLS estimates for the model

$$E[v_{y_\tau}(t, h)^2 | t - h] = \beta_0 + \sum_{j=1}^6 \beta_j PC_j(t - h) + \varepsilon(t),$$

where $\tau = 3M, 6M, 1Y, 2Y, 5Y,$ and $10Y$, and $h = 1, 5,$ and 21 days. The variance forecasts $E[v_{y_\tau}(t, h)^2 | t - h]$ are based on the HAR-RV model estimates reported in Section 4.2.2. In the $\tau = 3M$ case, we also report results based on the one-day-ahead forecast from the SNP model estimated in Section 4.2.1. $PC_j, j = 1, \dots, 6$, are the principal components extracted from the panel of zero-coupon yields. The sample period is 06/17/1991–06/15/2001. Standard errors estimates are robust with respect to both autocorrelation and heteroskedasticity. Coefficient t -ratios are in square brackets.

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]	
SNP, $\tau = 3M$	0.5043 [19.44]	0.0158 [1.00]	-0.0523 [-1.38]	0.0237 [0.19]	0.9715 [2.48]	-1.4407 [-3.41]	2.3697 [3.59]	10.25%
HAR-RV, $\tau = 3M$	0.8782 [21.83]	0.0366 [1.45]	-0.0812 [-1.35]	-0.1954 [-1.30]	1.4895 [2.08]	-1.5876 [-2.27]	2.7980 [2.76]	11.01%
HAR-RV, $\tau = 6M$	0.7800 [27.45]	0.0219 [1.32]	0.0078 [0.20]	0.2342 [1.65]	1.0881 [2.44]	-0.6961 [-1.69]	2.0852 [2.69]	9.67%
HAR-RV, $\tau = 1Y$	1.0391 [33.74]	0.0199 [1.29]	0.0473 [1.28]	0.3451 [2.32]	0.7069 [1.62]	-0.5247 [-1.23]	2.0495 [2.41]	9.13%
HAR-RV, $\tau = 2Y$	1.2181 [39.00]	0.0091 [0.63]	0.0142 [0.47]	0.3111 [1.95]	0.1271 [0.31]	-0.2658 [-0.63]	1.6965 [2.00]	4.36%
HAR-RV, $\tau = 5Y$	1.2342 [40.72]	0.0134 [0.96]	0.0155 [0.58]	0.0633 [0.39]	0.2546 [0.68]	-0.3681 [-0.85]	1.6192 [1.99]	2.24%
HAR-RV, $\tau = 10Y$	1.1062 [60.25]	-0.0005 [-0.06]	0.0317 [1.63]	0.0928 [0.99]	0.1915 [0.77]	-0.2390 [-0.85]	1.6106 [3.19]	5.51%

Panel A: The dependent variable is the one-day-ahead variance forecast ($h = 1$ day).

1st	2nd	3rd	4th	5th	6th
75.66%	17.82%	3.49%	1.86%	0.79%	0.38%

Panel B: Percentage of the variance explained by the principal components extracted from the OLS regressions' residuals.

6 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 no 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 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 and both monetary policy, the general economic environment, and inflationary expectations.²⁰ 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 were deemed stronger either at the shorter, medium, or longer-term maturities. Hence, we find the prospects of combining models which have general linkages to macroeconomic and monetary policy variables with models that incorporate time-varying reactions to macroeconomic announcements to be particularly promising. This type of approach has the potential to improve the explanatory power for the term structure dynamics and to render the implied evolution of the yield curve more directly interpretable in economic terms.

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.

Appendix

The full SNP density takes the form,

$$f_K(y_t|x_t;\xi) = \left(\nu + (1 - \nu) \times \frac{[P_K(z_t, x_t)]^2}{\int_{\mathbb{R}} [P_K(z_t, x_t)]^2 \phi(u) du} \right) \frac{\phi(z_t)}{y_{t-1}^\delta \sqrt{h_t}}, \quad \nu = 0.01, \quad (38)$$

$$z_t = \frac{y_t - \mu_t}{y_{t-1}^\delta \sqrt{h_t}},$$

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

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

$$b(z) = |z| \text{ for } |z| \geq \pi/2K, \quad b(z) = (\pi/2 - \cos(Kz))/K \text{ for } |z| < \pi/2K,$$

$$P_K(z, x) = \sum_{i=0}^{K_z} a_i(x) z^i = \sum_{i=0}^{K_z} \left(\sum_{|j|=0}^{K_x} a_{ij} x^j \right) z^i, \quad a_{00} = 1,$$

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

²⁰Previous 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).

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, 2005a, Roughing It Up: Including Jump Components in the Measurement, Modeling and Forecasting of Return Volatility, Working Paper, Northwestern University, Duke University, and University of Pennsylvania.
- 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, 2005b, 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, 2005, 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.
- 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.
- 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, 143-183.
- 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.
- 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.

- Taylor, Stephen J., and Xinzhong Xu, 1997, The Incremental Volatility Information in one Million Foreign Exchange Quotations, *Journal of Empirical Finance* 4, 317-340.
- Thompson, Samuel, 2004, Identifying Term Structure Volatility from the LIBOR-Swap Curve, Working Paper, Harvard University.
- Tuckman, Bruce, 2002, Fixed Income Securities: Tools for Today's Markets Wiley, John & Sons, Incorporated.
- Umantsev, Leonid A., 2001, Econometric Analysis of European LIBOR-Based Options within Affine Term-Structure Models, Ph.D. dissertation, Stanford University.