London School of Economics and Political Science

Essays on Asset Pricing

Hoyong Choi

A thesis submitted to the Department of Finance of the London School of Economics for the degree of Doctor of Philosophy, London, July 2016

Declaration

I certify that the thesis I have presented for examination for the MPhil/PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of 43,760 words.

Statement of conjoint work

I confirm that Chapter 3 was jointly co-authored with Philippe Mueller and Andrea Vedolin, and I contributed 33% of this work.

Abstract

The first chapter studies the impact of variance risk in the Treasury market on both term premia and the shape of the yield curve. Under minimal assumptions shared by standard structural and reduced-form asset pricing models, I show that an observable proxy of variance risk in the Treasury market can be constructed via a portfolio of Treasury options. The observable variance risk has the ability to explain the time variation in term premia, but is largely unrelated to the shape of the yield curve. Using the observable variance risk, I also propose a new representation of no-arbitrage term structure models. All the pricing factors in the model are observable, tradable, and hence economically interpretable. The representation can also accommodate both unspanned macro risks and unspanned stochastic volatility in the term structure literature.

The second chapter shows that it is beneficial to incorporate a particular zero-cost trading strategy into approaches that extract a stochastic discount factor from asset prices in a model-free manner (e.g. the Hansen-Jagannathan minimum variance stochastic discount factor). The strategy mimics the Radon-Nikodym derivative between two pricing measures with alternative investment horizons, and is hence characterized by the term structure of the SDF (or the dynamics of the SDF). Incorporating the strategy into the Euler equation significantly enhances the ability of the extracted stochastic discount factor to explain cross-sectional variation of expected asset returns. Furthermore, the strategy remarkably tightens various lower bounds for the stochastic discount factor, hence setting a more stringent hurdle for equilibrium asset pricing models.

The third chapter studies variance risk premiums in the Treasury market. We first develop a theory to price variance swaps and show that the realized variance can be perfectly replicated by a static position in Treasury futures options and a dynamic position in the underlying. Pricing and hedging is robust even if the underlying jumps. Using a large options panel data set on Treasury futures with different tenors, we report the following

findings: First, the term-structure of implied variances is downward sloping across maturities and increases in tenors. Moreover, the slope of the term structure is strongly linked to economic activity. Second, returns to the Treasury variance swap are negative and economically large. Shorting a variance swap produces an annualized Sharpe ratio of almost two and the associated returns cannot be explained by standard risk factors. Moreover, the returns remain highly statistically significant even when accounting for transaction costs and margin requirements.

Contents

Chapter 1

Information in (and not in) Treasury Options

1.1 Introduction

What is the role of variance risk in the Treasury market? How big is its impact on the risk-return trade-off? How does it affect the shape of the yield curve? What kind of macroeconomic uncertainty drives it? The first step to addressing these questions is to *identify* the variance risk in the Treasury market. In this paper, I suggest a novel approach to identifying variance risk, by utilizing information in Treasury bond options to answer the above questions.

I first show that variance risk can be proxied by implied variance measures from bond option markets and that this is true under a set of mild assumptions which are shared by many well-known structural and reduced form asset pricing models. Specifically, I prove that a bond VIX^2 (a portfolio of Treasury options constructed akin to the VIX^2 in the equity market[1]) represents the variance risk in the Treasury market under the assumptions that (i) the short-term interest rate is a linear function of the state variables and (ii) the state follows an affine diffusion process under the risk-neutral measure. In other words, the bond VIX^{2}s span time-varying variances in Treasury yields under the two assumptions. As a consequence, the impact of variance risk on both term premia and the shape of the yield curve is directly measurable via the *observable* variance risk: the bond VIX^{2}s. Using this theoretical framework, I obtain the following three novel results.

First, I propose a novel return-forecasting factor that *jointly* exploits the bond VIX^{2}s and the implication of leading macro-finance asset pricing models. The bond VIX^{2}s identify

[1]The VIX is a measure of volatility, and hence the VIX^2 is a measure of variance.

economic fundamentals that determine the conditional variances of bond yields in many well-known consumption-based asset pricing models: for example, the time-varying variance of consumption growth in Bansal and Yaron (2004), the probability of a rare disaster in Wachter (2013), and the external habit in Le, Singleton, and Dai (2010). Interestingly, the unobservable fundamentals are both the drivers of the variances of Treasury yields and the sole sources of time variation in term premia under these frameworks (see e.g. Le and Singleton (2013)). Hence, one common implication of the models is that excess returns on bonds should be completely explained by the bond VIX^2s. In particular, in the long-run risk framework of Bansal and Yaron (2004), the risk premium is time-varying solely due to the time variation in the quantity of risk. Moreover, changes in the variance of yields are the manifestation of time-varying macroeconomic uncertainties in the long-run risk framework. Because the short-term interest rate is postulated to be linear in affine diffusion states in the economy, the bond VIX^2s span the time-varying variance in yields. Hence, the time variation in expected excess returns should be captured by the bond VIX^2s. The same implication can also be obtained from the rare disaster framework of Wachter (2013). Time-varying probability of a rare disaster is assumed to follow an affine diffusion process in the framework, and it is also a sole driver of time variation in both risk premia and interest rate variance. Hence, the bond VIX^2s are the manifestation of time-varying disaster probabilities, and should have the ability to predict future excess returns. The affine$^{\mathbb{Q}}$ habit model in Le, Singleton, and Dai (2010) is another class of models in which the bond VIX^2s should be driven by the factor underlying the time variation in risk premia. In this framework, the external habit of the representative agent - the source of time-varying price of risks - is the only factor driving both the time variation in the variance of yields and the time-varying risk premia. Moreover, the drift in the pricing kernel is assumed to be linear in the state variables following an affine-diffusion process under the risk-neutral measure, and hence the bond VIX^2s reflect the time-varying price of risk. In sum, the space of time-varying risk premia and the space of the bond VIX^2s are identical under the standard macro-finance asset pricing models, the long-run risk, rare disaster, and affine$^{\mathbb{Q}}$ habit formation frameworks.

In the above three frameworks, the noise in realized excess returns on bonds can be completely removed by projecting realized excess returns onto the bond VIX^{2}s space. Hence, using implied variance measures from options on Treasury futures with different tenors, I use a projection in line with Cochrane and Piazzesi (2005). Similar to their regressions where they project bond excess returns of different maturity onto forward rates, I find that the linear combination of bond VIX^{2}s also produces a tent-shape factor forecasting excess returns. Interestingly, the predictive ability of this return-forecasting factor mainly stems from the excess returns of relatively short-term bonds, while the linear combination of forward rates in Cochrane and Piazzesi (2005) is superior in predicting excess returns on long-term bonds. Moreover, the single factor from the bond VIX^{2}s and the Cochrane-Piazzesi factor are complementary, and the predictability for bond returns increases significantly in joint regressions.

Second, I analyze the observable variance risk's impact on the shape of the yield curve. Its marginal impact is assessed by projecting yields onto the bond VIX^{2}s as well as the first three yield principal components: level, slope and curvature. After controlling for these factors, I find that variance risk is largely unrelated to the shape of the yield curve. This result corroborates earlier evidence of unspanned stochastic volatility (USV) whereby yield variance can only be very weakly identified from the cross-section of yields (see e.g., Collin-Dufresne and Goldstein (2002) among many others). However, the strict condition for the USV effect is rejected by a newly devised statistical test exploiting the observational variance risk. In sum, it is hard to identify the volatility of interest rates from the yield curve movements, but the knife-edge conditions for the USV effect do not seem to hold in the data.

Third, to assess the variance risk's impact on term premia and the cross-section of yields within a fully-fledged framework, I suggest a new representation of affine no-arbitrage term structure models that incorporate the observable variance risk. The representation follows in the spirit of Joslin, Singleton, and Zhu (2011), and extends their work to affine models with stochastic volatility. The risk factors are represented as a portfolio of yields and options. Hence, all the pricing factors are observable, tradable, and economically interpretable. In addition, due to the observable variance risk, the factor dynamics under the physical measure can easily be estimated by generalized least squares. Furthermore, the observable proxy of variance incorporates information in volatility-sensitive instruments, namely the Treasury options. As a result, the variance risk in interest rates is well identified, in contrast to the conventional latent factor approaches. Finally, the model can be easily extended to reflect the unspanned macro risks in Joslin, Priebsch, and Singleton (2014) (henceforth JPS). Given that the observable variance factor can be unspanned by yields, the model can accommodate the two distinct types of unspanned risks in the term structure literature: unspanned macro risk factor (hidden factor) and unspanned stochastic volatility. The estimates of the model indicate that both unspanned macro risk and stochastic volatility drive expected returns. The stochastic volatility factor in the estimated model is not literally unspanned by yields, but its impact on the shape of the yield curve is noticeably small and can be effectively treated as an unspanned factor.

This paper also contributes to the recent discussion on unspanned macro risks in the macro-finance term structure literature. The unspanned macro risks are macroeconomic factors that are informative about macroeconomic fluctuations and term premia, but largely unrelated to the term structure movements. One open question with this strand of studies[2] is, among the hundreds of macroeconomic variables, which one *should* or *could* be treated as an unspanned macro risk? For example, Bauer and Rudebusch (2016) show that estimates of risk premia can differ significantly depending on whether a measure of the *level* or the *growth* in economic activity is used as unspanned risk. I show that the LPY ("linear projection of yields") criteria in Dai and Singleton (2002) provide informative guidance on this issue. The LPY criteria are descriptive statistics that measure whether a term structure model can match the pattern of violation of the expectations hypothesis as in Fama and Bliss (1987) or Cambpell and Shiller (1991). For the issue of choosing level or growth indicators of economic activity as an unspanned macro risk, the LPY criteria indicate that level variable is more relevant measure of economic activity in term structure modeling perspective. In other words, the models with level of economic activity as an unspanned macro risk are better at re-producing the pattern for the failure of expectations hypothesis in the data than the models with growth indicator as unspanned macro risk. Furthermore, in the LPY

[2]See e.g., Duffee (2011b), Chernov and Mueller (2012) and Joslin, Priebsch, and Singleton (2014).

dimension, the stochastic volatility models with/without unspanned macro risks outperform the corresponding Gaussian models. This shows that the observational variance risk is (i) properly identified and (ii) beneficial in explaining the time variation in risk premia.

This paper is related to several different strands of the literature. First, the construction of an observable proxy of variance risk in interest rates is based on the methodology of Mele and Obayashi (2013) and Choi, Mueller, and Vedolin (2016). However, while these papers utilize bond VIX^2s to study the price of variance risk or variance risk premium in a model-free manner, this paper (i) initially identifies the classes of asset pricing models under which the bond VIX^2 is equivalent to the interest rate variance risk of the models, (ii) and then *jointly* utilizes both the bond VIX^2s and the implication of the asset pricing models for a better understanding of expected excess returns on long-term bonds (rather than variance trading). In other words, given an asset pricing model within the class characterized by (i) affine short rate and (ii) affine state under the risk-neutral measure, variance risk takes the form of bond VIX^2, and this observable portfolio of options inherits all the properties and implications of the variance risk in the model. In this paper, the bond VIX^2s are utilized as instruments to identify such variance risks *within* the models. For structural asset pricing models with the two assumptions, the bond VIX^2s identify economic fundamentals that drive variance risks in the Treasury market. Hence, the bond VIX^2s should inherit all the asset pricing implications of the fundamentals.

This idea implies that within the long-run risk, rare disaster, and affine$^{\mathbb{Q}}$ habit formation frameworks, the bond VIX^2s should predict excess returns on bonds because the set of risk factors underlying variation in risk premia is the sole source of time-varying variances in bond yields. Hence, the return-forecasting factor in this paper is based on the theoretical prediction of those specific models, contrary to the return-forecasting factors from the yield curve as in Fama and Bliss (1987), Cambpell and Shiller (1991), and Cochrane and Piazzesi (2005). Furthermore, the bond VIX^2s can be measured in real time and contain forward-looking information, in contrast with infrequently-updated macro data as in Bansal and Shaliastovich (2013) or Ludvigson and Ng (2009).

The benefit of observable variance is also highlighted in the connection of the bond VIX^2 to the no-arbitrage affine dynamic term structure models (henceforth, ADTSM). Under the assumption of ADTSM, the bond VIX^2 directly identifies variance risk in ADTSM, which has been considered one of the most challenging tasks in the term structure literature. While previous term structure models also incorporate information from volatility-sensitive instruments into their estimation procedure for better identification of variance risk[3], the approach of this paper circumvents their computational difficulties. Specifically, previous studies match *individual* derivative prices from the models to actual derivative prices in their estimation procedures, but the calculations of the derivative prices are extremely cumbersome computationally. By formulating a specific option *portfolio* that directly reflects the changes in the underlying variance factor, the approach I propose simplifies the incorporation of information in volatility-sensitive instruments into ADTSM.

Furthermore, the observable variance risk enables ADTSM to be represented by observable and tradable factors, contrary to all the previous dynamic term structure models with stochastic volatility. Hence, the new representation of ADTSM that I posit here is based on the observable variance risk, and extends the representation for both spanned Gaussian ADTSM in Joslin, Singleton, and Zhu (2011) and Gaussian ADTSM with unspanned macro risk in Joslin, Priebsch, and Singleton (2014) into more general setting. Joslin and Le (2014) also utilize a parameterization scheme for ADTSM with stochastic volatility, in which the time-varying variance factor is approximated by observable portfolio of yields. Their volatility instrument can only be identified after the estimation of the model, while the bond VIX^2 identifies variance risk even before the estimation of ADTSM. Furthermore, the approach of this paper is robust to unspanned stochastic volatility (USV), because option prices are utilized to detect variance risk. On the other hand, yields do not span variance risk in the presence of USV, and hence one cannot construct a yield portfolio that captures time-varying variance as in Joslin and Le (2014).

Finally, while all the other USV models in the literature should be estimated with hard-wired constraints to generate USV effects[4], the approach here does not impose a priori constraints for the USV effect and lets the data speak about the presence of USV. With

[3]See e.g. Jagannathan, Kaplin, and Sun (2003), Bibkov and Chernov (2009), Trolle and Schwartz (2009), Bibkov and Chernov (2011), Almeida, Graveline, and Joslin (2011) and Joslin (2014) among many others.

[4]See e.g. Bibkov and Chernov (2009), Collin-Dufresne, Goldstein, and Jones (2009), Trolle and Schwartz (2009), Joslin (2015) and Creal and Wu (2015) among many others.

the bond VIX^2 at hand, the estimation of ADTSM reveals the relative importance of the already-identified variance factor in determining the shape of the yield curve. Once the bond VIX^2 turns out to play little role in explaining the cross-section of yields, then one can effectively treat it as an unspanned stochastic volatility factor.

The paper proceeds as follows. Section 1.2 theoretically shows how one can construct an observable proxy of the variance risk in the Treasury market by utilizing information in option markets. Section 1.3 argues why the observable measure of variance could capture the time-variation in risk premia, and investigates its predictive ability for excess returns. Section 1.4 analyzes the relation between the variance risk and the shape of the yield curves. Section 1.5 introduces a new representation for no-arbitrage term structure models in which the variance risk is identified as a portfolio of Treasury options. The representation is extended to accommodate unspanned macro risks in Section 1.6. In Section 1.7, the models in Section 1.6 but with different types of unspanned macro risks are evaluated based on the LPY criteria. Section 1.8 explores the properties of risk premia in more depth. Finally, Section 1.9 concludes. All proofs are deferred to the Appendices.

1.2 Observable Volatility

To start, let us assume the state variable $Z_t = (X_t', V_t')' \in \mathbb{R}^{N-m} \times \mathbb{R}_+^m$ follows the Ito diffusion under the risk-neutral measure $\mathbb{Q}$

$$d \begin{bmatrix} X_t \\ V_t \end{bmatrix} = \mu_{Z,t} dt + \Sigma_{Z,t} dB_t^{\mathbb{Q}} \tag{1.1}$$

where

$$\mu_{Z,t} = \begin{bmatrix} \mu_{X,t} \\ \mu_{V,t} \end{bmatrix} = \begin{bmatrix} K_{0X} \\ K_{0V} \end{bmatrix} + \begin{bmatrix} K_{1X} & K_{1XV} \\ K_{1VX} & K_{1V} \end{bmatrix} \begin{bmatrix} X_t \\ V_t \end{bmatrix}, \text{ and } \Sigma_{Z,t}\Sigma_{Z,t}' = \Sigma_{Z0} + \sum_{i=1}^{m} \Sigma_{Zi} V_{it}$$

with a set of restrictions on the parameters to ensure the non-negativity of the volatility factor V_t as in Duffie, Filipović, and Schachermayer (2003). $B_t^{\mathbb{Q}}$ is a N-dimensional Brownian motion under $\mathbb{Q}$. The short rate (the negative of the drift in a pricing kernel) is assumed to be linear in the state Z_t

$$r_t = \delta_0 + \delta_1 Z_t \tag{1.2}$$

In addition, denote the following portfolios of options as $\mathcal{V}_t$ which is a measure of model-free implied variance akin to the Chicago Board Options Exchange (CBOE) VIX^2 in equity markets:

$$\mathcal{V}_t\left(T,\mathbb{T}\right)=\frac{2}{P_{t,T}}\left[\int_0^{F_t(T,\mathbb{T})}\frac{\text{Put}_t\left(K,T,\mathbb{T}\right)}{K^2}dK+\int_{F_t(T,\mathbb{T})}^{\infty}\frac{\text{Call}_t\left(K,T,\mathbb{T}\right)}{K^2}dK\right] \tag{1.3}$$

where $P_{t,T}$ is the price of a zero-coupons bond expiring at T, and $F_t\left(T,\mathbb{T}\right)$ is the forward price at t, for delivery at T, of the bond maturing at $\mathbb{T}$. $\text{Put}_t\left(K,T,\mathbb{T}\right)$ and $\text{Call}_t\left(K,T,\mathbb{T}\right)$ are European options with strike price K and tenor T written on $P_{t,\mathbb{T}}$. It is well-known in the equity literature that cross-sectional information from options enables us to recover the risk-neutral probability density of underlying asset (Breeden and Litzenberger (1978)). The CBOE VIX is a specific application of this theory, to proxy the forward-looking risk-neutral volatility of the one-month return on S&P 500 index. Similarly, with T being equal to one-month, $\sqrt{\mathcal{V}_t\left(T,\mathbb{T}\right)}$ can be considered as a forward-looking measure of one-month volatility in $P_{t,\mathbb{T}}$ under the risk-neutral measure.

Under the two assumptions that the state is an affine process as in (1.1) and that the short rate is affine in the state Z_t as like (1.2), it can be shown that $F_t\left(T,\mathbb{T}\right)$ follows a diffusion process of which instantaneous variance is a linear function of the latent factor V_t. When $F_t\left(T,\mathbb{T}\right)$ follows a diffusion process, it is well-known that equation (1.3) represents the expected quadratic variation of the forward under $\mathbb{Q}_T$ measure of which numéraire is the bond $P_{t,T}$ (see, e.g., Carr and Madan (1998)). Furthermore, the change of measure between the forward measure $\mathbb{Q}_T$ and the risk-neutral measure $\mathbb{Q}$ is determined by the volatility of $F_t\left(T,\mathbb{T}\right)$ in a linear fashion: see for example Björk (2009). As a consequence, $\mathcal{V}_t$ can be expressed as a linear function of V_t, which means that one can observe the latent variance factor up to its linear transformation and its shocks via the option portfolio $\mathcal{V}_t$.

Proposition 1. *Suppose that the short rate is an affine function of the $\mathbb{Q}$ affine process in (1.1). Then,*

$$\mathcal{V}_t\left(T,\mathbb{T}\right)=\alpha\left(\Theta^{\mathbb{Q}};t,T,\mathbb{T}\right)+\beta\left(\Theta^{\mathbb{Q}};t,T,\mathbb{T}\right)\cdot V_t \tag{1.4}$$

where $\Theta^{\mathbb{Q}}$ is the set of parameters for (1.1) and (1.2).

Proof: See Appendix 1.

One of the key features of Proposition 1 is that it does not require any specification of the market price of risk (or the dynamics of Z_t under $\mathbb{P}$) to completely characterize a pricing kernel. In other words, the proposition can be utilized even though Z_t follows a non-linear process under $\mathbb{P}$. In sum, for large classes of asset pricing models, one can capture the innovations in variance factor through the portfolio of options, $\mathcal{V}_t$.

As can be seen from equation (1.3), the option portfolio $\sqrt{\mathcal{V}_t}$ is a Treasury market version of the CBOE VIX in the equity market. While the VIX has been intensively studied and utilized in the literature,[5] studies about its analogue for US Treasuries (henceforth, the bond VIX) started relatively recently. Mele and Obayashi (2013) develop theories on pricing Treasury volatility (i.e. expected value of Treasury volatility under a forward measure), and suggest a practical way of representing the price as a portfolio of Treasury futures options. Based on their methodology, CBOE launched the 10-year U.S. Treasury Note Volatility Index (TYVIX) in May 2013. Choi, Mueller, and Vedolin (2016) show how investors can make use of the bond VIX to get pure exposure to variance risk in the fixed income market and document the empirical properties of the trading strategy. They construct the bond VIX named as Treasury Implied Volatility index (TIV) for a 10-year T-note, plus TIV for a 5-year Treasury bill and a 30-year Treasury bond.

This paper utilizes their TIVs since the three measures of volatility with different underlying bonds enable us to identify multiple latent volatility factors via Proposition 1. For a detailed description of how to construct TIV, I refer the reader to Choi, Mueller, and Vedolin (2016). Figure 1.1 provides a plot of the CBOE VIX, the CBOE TYVIX, and the 10-year TIV; following the custom in practice, they are the square root of the annualized variances expressed in percent. The 10-year TIV is virtually identical to the TYVIX, and they are largely correlated with the VIX. The bond VIXs are driven by the variance factor in the discount rates (or the pricing kernel), while the VIX reflect the variance factor in both the discount rates and cash flow dynamics. The figure shows that the impact of the variance factor in the cash flow dynamics became less important from late 90s.

[5] See, e.g., Carr and Wu (2009), Drechsler and Yaron (2011), Bollerslev, Tauchen, and Zhou (2009), Ang, Hodrick, Xing, and Zhang (2006), Adrian and Shin (2010), Nagel (2012), Brunnermeier, Nagel, and Pedersen (2009), Bao, Pan, and Wang (2011), Amengual and Xiu (2014),Bekaert, Hoerova, and Duca (2013), Kelly, Pástor, and Veronesi (2014) among many others.

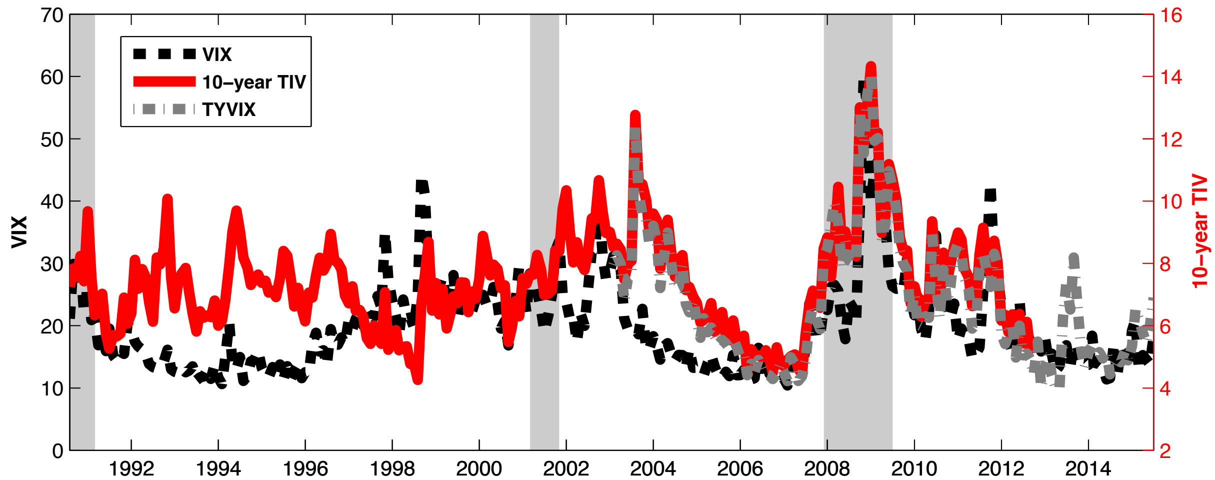

Figure 1.1. TIV,TYVIX and VIX

This figure plots the CBOE VIX, the CBOE TYVIX and the 10-year TIV from Choi, Mueller, and Vedolin (2016). Volatilities are the square root from variances as constructed using option prices via equation (1.3). Numbers are annualized and expressed in percent. Gray bars indicate NBER recessions. The data is monthly and runs from July 1990 to June 2015; The 10-year TIV data ends in August 2012, and the TYVIX data starts from January 2003.

To summarize, Proposition 1 gives the implication of the model-free measure of implied volatility in the Treasury market, the bond VIX, once it is combined with additional structures embedded in many economic models. Once the information in bond VIX is incorporated with an economic model in which the short rate is linear in $\mathbb{Q}$ affine diffusion state variables, the bond VIXs can completely identify the volatility factors. This result also implies that some economic fundamentals in macro-finance asset pricing models can be identified via the bond VIX^{2}s if the fundamentals determine the conditional variances of bond yields.

1.3 Predictability

The assumptions for Proposition 1 are that (i) the drift of a pricing kernel is affine in the state variable and (ii) the state variable follows affine diffusion under $\mathbb{Q}$. Three classes of well-known consumption-based asset pricing models incorporate this feature. They are the long-run risk framework of Bansal and Yaron (2004), the rare disasters framework of Wachter (2013), and the affine$^{\mathbb{Q}}$ habit formation model of Le, Singleton, and Dai (2010).

Importantly, in all of these models, the source of time variation in risk premia is entirely spanned by volatilities in yields only (see Le and Singleton (2013) for detailed explanations). Once this salient feature of those models is incorporated with Proposition 1, it means that the bond VIX^{2}s should predict future excess returns and they are the sole source of time variation in risk premia.

Specifically, the assumptions in Proposition 1 are canonical in most long-run risks models without jumps (or rare disasters) - see e.g. Bansal and Yaron (2004), Bollerslev, Tauchen, and Zhou (2009), Bansal and Shaliastovich (2013), Zhou and Zhu (2015). For example, in Bansal and Shaliastovich (2013), (i) the drift of the pricing kernel is a linear function of the subset of affine diffusion state variables - expected consumption growth, expected inflation and their variance factors, and (ii) the $\mathbb{P}$ affine state variable, in conjunction with their market price of risk, imply $\mathbb{Q}$ affine state variable. The time-varying volatilities in expected consumption growth and expected inflation are the two economic fundamentals that induce time variation in volatilities in yields. In this economy, Proposition 1 implies that the bond VIX^{2}s should be the manifestation of uncertainty about the two macroeconomic fundamentals: expected consumption growth and expected inflation (see Appendix 2 for a formal derivation). Note that in long-run risk economies, the time-varying quantity of macroeconomic risk is the only source of time variation in risk premia - the price of risk is pinned down by Epstein-Zin preference. As a consequence, the bond VIX^{2}s should capture the entire innovations in risk premia though the channel of time-varying quantity of risk.

The rare disaster framework with time-varying disaster probabilities is another example that fits the assumptions of Proposition 1 - see e.g. Wachter (2013), and Tsai (2016). In this framework, the short rate is linearly dependant on time-varying risk of disasters (intensity of a disaster more precisely). The intensity process follows affine diffusion under both the physical and the risk-neutral measures, and it also determines volatilities in yields. Then, the bond VIX^{2}s disclose the time-varying probability of a disaster because of Proposition 1. Moreover, in this economy, time variation in risk premia solely stems from the time-varying probability of a disaster. Hence, the bond VIX^{2}s should have the ability to predict future excess returns.

The habit formation model in Le, Singleton, and Dai (2010), henceforth LSD, is another class of asset pricing models in which the bond VIX^{2}s should explain the entire time variation in risk premia. The model, based on Campbell and Cochrane (1999) and Wachter (2006), uses the two assumptions in Proposition 1 to obtain affine pricing. By doing so, they specify the market price of risk as a non-linear function of the states as in Duarte (2004) and, as a result, the state variable follows a non-linear process under $\mathbb{P}$. LSD shows that their model approximately nests Wachter's model and closely resembles its prominent features. In this type of affine$^{\mathbb{Q}}$ habit formation models with external habit level H_t, the consumption surplus ratio $s_t = \log\left[\left(C_t - H_t\right)/C_t\right]$ is the sole source of time-varying risk premia since the shocks on consumption growth (that drives the quantity of risk in the economy) are assumed to be homoscedastic. Furthermore, the volatilities in yields are driven by the non-negative process $\varphi_t = s_{\max} - s_t$ where $s_{\max}$ is the upper bound of s_t.[6] Hence, the bond VIX^2 is linear in φ_t, the inverse consumption surplus ratio, and contains the full information on risk premia through the reflection of the time-varying price of risk.

Motivated by the implication of Proposition 1 for the three classes of asset pricing models, I examine whether the bond VIX^{2}s explain time variation in expected bond excess returns. To assess their predictive ability, I initially apply MA2 filters for the one-month bond VIX^{2}s (with 5yr, 10yr and 30yr bonds as underlying assets) constructed in Choi, Mueller, and Vedolin (2016) with the aim of removing transitory shocks potentially due to measurement errors and institutional effects (see, for example, Kim (2007)). Then, I regress one-year holding period excess returns of bonds with different maturities onto the space of the three (filtered) one-month bond VIX^2, henceforth denoted as TIV^{2}s following Choi, Mueller, and Vedolin (2016). The projections indicate that, across all maturities, the excess returns' loadings on the TIV^{2}s exhibit tent-shape pattern akin to the pattern in Cochrane and Piazzesi (2005). Hence, in the spirit of Cochrane and Piazzesi (2005), I construct a single factor by projecting the average (across maturity) excess returns onto the three TIV^{2}s:

$$\overline{rx}_{t+12} = \gamma_0 + \gamma_1 TIV^2_{t,5yr} + \gamma_2 TIV^2_{t,10yr} + \gamma_3 TIV^2_{t,30yr} + e_t$$

[6] Note that s_t is always negative.

Table 1.1
Predictive Regressions

Panel A reports adjusted R^2 from regressing twelve-month excess returns rx(n) of bonds with n years to maturity on CP factor, CIV, and both. Panel B presents estimated coefficients from predictive regressions from the mean of excess returns in Panel A onto CP factor or (and) CIV. Standard errors are in parentheses and adjusted according to Newey and West (1987). Data is monthly and runs from October 1990 to December 2007.

Panel A: Predictive Regressions and Adjusted R^2s

Excess returns from Fama-Bliss data				Excess returns from GSW data			
	CP	CIV	$CIV\&CP$		CP	CIV	$CIV\&CP$
rx(2)	0.20	0.27	0.36	rx(2)	0.15	0.26	0.31
rx(3)	0.22	0.28	0.38	rx(4)	0.21	0.28	0.36
rx(4)	0.24	0.28	0.40	rx(6)	0.25	0.27	0.38
rx(5)	0.22	0.28	0.38	rx(8)	0.27	0.26	0.39
				rx(10)	0.28	0.23	0.38

Panel B: Predictive Regression Coefficients

	CIV	CP	adj R^2
mean(FB rx)	0.53		0.28
	(0.10)		
mean(FB rx)	0.35	0.43	0.39
	(0.10)	(0.12)	
mean(GSW rx)	0.52		0.27
	(0.11)		
mean(GSW rx)	0.38	0.39	0.39
	(0.11)	(0.12)	

The time-series of the fitted values (henceforth, CIV) is the return-forecasting factor, and is utilized to predict realized excess returns on each bond with maturity of n.

$$rx_{t+12}^{(n)} = b_0^{(n)} + b_1^{(n)} \left(\hat{\gamma}_0 + \hat{\gamma}_1 TIV_{t,5yr}^2 + \hat{\gamma}_2 TIV_{t,10yr}^2 + \hat{\gamma}_3 TIV_{t,30yr}^2\right) + e_t^{(n)}$$

For comparison purposes, the Cochrane-Piazzesi return-forecasting factor (henceforth, CP) is constructed by regressing mean excess returns onto the *spreads* of five Fama-Bliss forward rates with maturities of 1 through 5 years as in Cochrane and Piazzesi (2008). Excess returns from the Gürkaynak, Sack, and Wright (GSW) data set (with maturities of one through 10 years) are also utilized to assess excess returns on long-term bonds since the longest time-to-maturity of yields in the Fama-Bliss (FB) data set is five years.

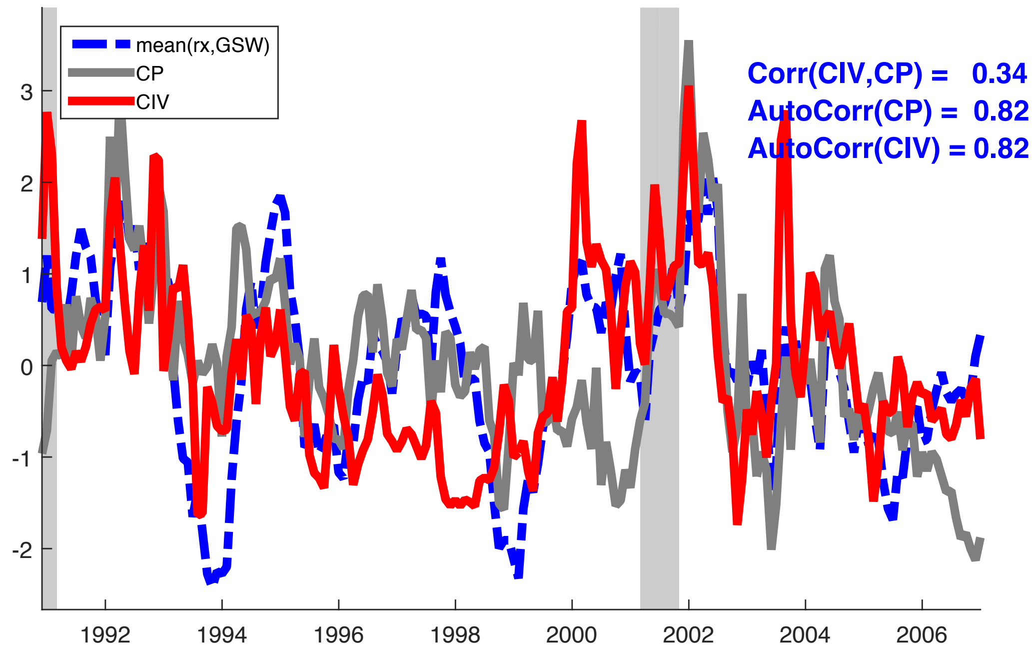

Figure 1.2. Excess Returns, CP and CIV

This figure plots the Cochrane-Piazzesi factor (CP), CIV, and the average of twelve-month excess returns on bonds with maturities of 1 through 10 years. Gray bars indicate NBER recessions, and blue bars represent financial crisis periods. The data is monthly and runs from October 1990 to December 2007.

Table 1.1 present adjusted R^2s and coefficients from predictive regressions of twelve-month bond excess returns on CP, CIV, and both CP and CIV jointly. Panel B shows that, for both sets (FB and GSW) of excess returns, each estimated coefficient on CIV is statistically significant during the sample period, and the variation in CIV explains more than 20% of the variation in realized mean excess returns. Once CIV and CP are jointly utilized, R^2s increase more than 10 percentage points in addition to the statistical significance of both coefficients. Panel A reports adjusted R^2 from regressing each excess return on the predictors, and it reveals that the predictability of the CIV stems mainly from excess returns of bonds with short-term maturities while the predictability of the CP comes from relatively long-term maturities. As a result, the adjusted R^2s are improved significantly once CIV and CP are utilized jointly to predict excess returns. Their joint significance can also be observed in Figure 1.2 where the time-series of CIV, CP and the mean realized excess returns from the GSW data set are plotted together. Between 1998 and 2002, for example, CIV and CP exhibit heterogeneous movements and can assist each other to explain the time variation in the excess returns.

1.4 Variance Risk and the Shape of the Yield Curve

How does volatility risk affect the shape of the yield curve? Do yields strongly/weakly load on volatility risk? Can we extract a reliable measure of interest rate volatility from the cross-section of bonds? The first step to addressing these questions is to identify volatility risk in a framework where volatility has a systematic impact on the yields across different maturities.

The class of affine dynamic term structure models (henceforth, ADTSM) is a typical example of such a framework, and has also served as the workhorse in the literature to assess the impact of volatility risk on the cross-section of bond yields; ADTSMs are fully characterized by (i) the two assumptions in Proposition 1, (ii) the specification of the market price of risk, and (iii) a set of parametric restrictions needed to identify the model. It is well established that the affine models successfully capture the cross-sectional properties of yields; see for example Dai and Singleton (2000). However, the ADTSMs' ability to capture variation in the volatility of interest rates is questionable and controversial, especially once volatility-sensitive derivatives are not incorporated into the estimation procedure of the model. For instance, using U.S. swap data only, Collin-Dufresne, Goldstein, and Jones (2009) show that the model implied volatilities from affine models seem unrelated to their non-parametric or semi-parametric counterparts (i.e. realized volatility estimates and GARCH estimates).

Because ADTSMs are built up on the two assumptions in Proposition 1, the bond VIX^2s directly represent variance risk in the model. In other words, the variance risk in ADTSM is readily identifiable via the bond VIX^2s as a consequence of Proposition 1. This identification strategy is beneficial in several ways. First, it is based directly on option prices that tend to be more sensitive to the changes in volatility than nominal bond prices. This is in line with previous studies pointing out that the introduction of volatility-sensitive instruments into the estimation procedure can significantly mitigate the difficulty in identifying volatility risk of affine models; see, for example, Bibkov and Chernov (2009), Almeida, Graveline, and Joslin

(2011), Jagannathan, Kaplin, and Sun (2003), and Joslin (2014). In addition, the approach doesn't require us to estimate a specific model, and allows variance risk to be measured in real time. Furthermore, since the variance measure is constructed in a model-free manner, the approach can be easily incorporated into the class of Gaussian quadratic term structure models, as in Ahn, Dittmar, and Gallant (2002). In this case, $\mathcal{V}_t$ is a quadratic function of the Gaussian state factors in the model [7] (See Appendix 3 for a detailed explanation).

Before estimating a fully-fledged model to analyze the impact of variance risk on the cross-section of yields, I conduct two simple regression-based tests on the relationship between variance risk and the cross-section of yield. First, I examine the marginal impact of the variance risk on the shape of the yield curve beyond the traditional term structure factors: level, slope and curvature of the yield curve. The results suggest that variance risk is largely unrelated to the shape of the yield curve and that at least three non-volatility factors are required to adequately explain the cross-section of yields. The second test investigates whether variance risk can be identified from the cross-section of yields: the unspanned stochastic volatility (USV) effect in Collin-Dufresne and Goldstein (2002). The USV effect can be or cannot be rejected, depending on the number of variance factors. The empirical evidence will be utilized as guidance for designing highly parameterized term structure models in later sections.

1.4.1 The Shape of the Yield Curve and $\mathcal{V}_t$

Provided that (i) the state follows affine diffusion and (ii) the short rate is an affine function of the state, the yield on a zero-coupon bond of maturity n is affine in the state variable Z_t:

$$y_{n,t} = A_n\left(\Theta^{\mathbb{Q}}\right) + B_n\left(\Theta^{\mathbb{Q}}\right) Z_t \tag{1.5}$$

where A_n and B_n are obtained from standard recursions as in Duffie and Kan (1996). The linear relationship between yields and factors in equation (1.5) implies that yields can be treated as state variables; given a set of maturities equal in number to the number of latent factors, one can rotate the underlying factor into the yields (see for example Pearson and Sun

[7] For Gaussian quadratic term structure models, the short rate equation is a quadratic function of Gaussian state vector. The conditional variance of yields is linear in the square of a subset of the state.

(1994), Chen and Scott (1993) and Duffie and Kan (1996) among many others). One can further rotate the risk factors into portfolios of yields, especially the principal component of yields, $\mathcal{P}_t$, as in Joslin, Singleton, and Zhu (2011) for example. As will be shown thoroughly in Section 1.5, Proposition 1 enables us to rotate the latent risk factors into portfolio of yields $\mathcal{P}_t$ and portfolio of options $\mathcal{V}_t$

$$y_t^o = \mathcal{A} + \mathcal{B}_{\mathcal{Z}}\mathcal{Z}_t = \mathcal{A} + \mathcal{B}_{\mathcal{P}}\mathcal{P}_t + \mathcal{B}_{\mathcal{V}}\mathcal{V}_t + e_t, \quad e_t \sim N\left(0, \sigma_e^2 I\right) \tag{1.6}$$

where y_t^o denotes a vector of stacked observed yields and e_t represents measurement error assumed to be an independent and homoscedastic Gaussian random variable (as commonly assumed in the literature). Because all the variables in equation (1.5) can be observable, the yields' loading on the factors $\mathcal{A}$, $\mathcal{B}_{\mathcal{P}}$ and $\mathcal{B}_{\mathcal{V}}$ can be estimated by linear regressions. The estimated model, then, can be treated as a standard linear factor model nesting the no-arbitrage affine models since $\mathcal{A}$, $\mathcal{B}_{\mathcal{P}}$ and $\mathcal{B}_{\mathcal{V}}$ are non-linear functions of $\Theta^{\mathbb{Q}}$ under the affine bond pricing models: see for example, Duffee (2011a), Hamilton and Wu (2012), Joslin and Le (2014) and Joslin, Le, and Singleton (2013).

The marginal impact of the variance risk beyond traditional yield factors like level, slope and curvature factors can be examined by comparing the likelihood of (1.6) with the following restricted version of it:

$$y_t^o = \mathcal{A}^* + \mathcal{B}_{\mathcal{P}}^*\mathcal{P}_t + e_t^*, \qquad\qquad e_t^* \sim N\left(0, \sigma_{e^*}^2 I\right) \tag{1.7}$$

Table 1.2 reports the test statistics of the likelihood ratio test for the hypothesis of the zero coefficients on the additional variable in the unrestricted version. The first column presents the right hand side variables in the restricted models where PC1-PC3 denotes the first, second and third principal components of yields on U.S. Treasury nominal zero-coupon bonds with maturities of six months and 1 through 10 years[8]. The remaining columns present an additional variable in each version of the unrestricted model and its corresponding LR statistics. VPC1 and VPC2 denotes the first and second principal components of the MA2 filtered 5, 10 and 30-year TIV^{2}s as in Section 1.3. Each of VPC1 and VPC2 capture

[8]The yields with maturities of two to ten years are from Gürkaynak, Sack, and Wright (2007). The six-month and one-year yields are bootstrapped from observed bond prices using the Fama-Bliss methodology. My thanks to Anh Le for allowing me to use this data set.

Table 1.2
Marginal Impact of Variance Risks onto the Shape of the Yield Curve

This table reports the test statistics of the likelihood ratio test. The first column presents the right hand side variables in the restricted models: equation (1.7). The remaining columns present an additional variable in each version of the unrestricted model and its corresponding LR statistics. The last column shows the 5% critical value of the test statistics, which follows $\chi^2(11)$ distribution. Data is monthly and runs from October 1990 to December 2007.

Restricted Model	Additional Variable in Unrestricted model				
	VPC1	VPC2	PC3	PC4	C.V.(5%)
PC1, PC2	4.2	14.2	5882.9		19.7
PC1, PC2, PC3	45.0	18.6		3678.7	19.7

respectively 94% and 5.6% of the variation in the three TIV^2s. The last column shows the 5% critical value of the test statistics which follows $\chi^2(11)$ distribution. The table indicates that for each version of the restricted model, its likelihood ratio is greatest when a yield factor (PC3 or PC4) is the additional variable in the unrestricted model. In other words, adding a PC factor to the restricted models is the best extension for the purpose of a better cross-sectional fit. Moreover, for the unrestricted models with the variance factors as the additional variables, the null can be rejected or not, but the magnitude of test statistics is not very large, regardless of their statistical significance. Similar results are obtained once two or three representative yields, instead of the yield PCs, are utilized as the right hand side variables of the restricted models (the results are omitted in the paper for the sake of brevity). In sum, the exercise indicates that the marginal benefit of adding variance factors is fairly limited and it is hard to identify variance risk from the cross-section of yields.

The exercise also implies that it is empirically difficult to extend the estimation approach of Hamilton and Wu (2012) into affine bond pricing models with stochastic volatilities. They propose a minimum-chi-square estimation procedure of Gaussian ADTSM in which the risk-neutral parameters of the model are inferred by minimizing the differences between the ordinary least square (OLS) estimates of the cross-sectional equation (1.7) and the corresponding yields' loadings from Gaussian ADTSM. In theory, their approach can be applied to equation (1.6) for the estimation of ADTSM with stochastic volatilities. However, the limited impact of the variance risk on the shape of the yield curve causes difficulties in its empirical implementation. The OLS estimates of $\mathcal{B}_{\mathcal{V}}$ in equation (1.6) are not informative enough to precisely pin down the risk-neutral parameters related to the volatility factors.

1.4.2 Unspanned Stochastic Volatility Effect

The difficulty of identifying volatility risk under ADTSM stems from the multiple roles of volatility risk in the class of affine models. The volatility risk affects (i) the second moments of yields, (ii) the expectation of future interest rates under both physical and risk-neutral measures, and (iii) the so-called convexity effect introduced by the non-linear relationship between bond prices and the latent factors. The various roles of volatility enables us to infer it through multiple channels, but this feature causes tension rather than a complimentary effect in identifying it (see Joslin and Le (2014) for a detailed explanation).

One potential resolution for the issue is to impose a set of model-based restrictions to remove the dependence of the cross-section of yields on volatility, a set of restrictions coined as an "unspanned stochastic volatility" (USV) restriction by Collin-Dufresne and Goldstein (2002). More broadly, the USV effects mean that the yields curve itself fails to span the volatilities in the changes in yields. In their seminar paper, Collin-Dufresne and Goldstein (2002) define the USV effect as the existence of a set of parameters $\{\phi_1, ..., \phi_N\}$ that are not all zero such that

$$\sum_{i=1}^{N} \phi_i B_{n,i} = 0 \quad \forall n > 0 \tag{1.8}$$

where N is the number of pricing factors and $B_{n,i}$ the i-th element of B_n in equation (1.5). The authors further show that, under the existence of such a set of parameters with $N \geq 3$, one can find a rotation such that the variance factor V_t has no effect on the price of bonds. As a consequence, the variance factor cannot be extracted from the cross-section of observed yields (see Collin-Dufresne and Goldstein (2002), and Joslin (2015) for further details).

The following studies, however, have accumulated conflicting evidence on the USV effect. Decoupling the dual role of volatility through the USV restrictions helps the model to produce more realistic model-implied volatility, even though the model's cross-sectional fit is slightly impeded (see for example Creal and Wu (2015) and Collin-Dufresne, Goldstein, and Jones (2009) among others). Andersen and Benzoni (2010) also show that their measure of intraday volatility in yields is largely unexplained by term structure factors, which is in line with the USV effect. On the other hand, the USV effect is rejected once the model-specific restrictions are directly tested by the likelihood-ratio or the Wald test (Bibkov and Chernov (2009), Joslin (2015)). Utilizing the observable volatility proxy, I devise a new test for the USV effects, which can shed new light on the debate.

The condition for the USV effect in equation (1.8) can be translated into the statement that the matrix $\mathcal{B}_{\mathcal{Z}} \equiv [\mathcal{B}_{\mathcal{P}}, \mathcal{B}_{\mathcal{V}}]$ in equation (1.6) is not full rank, regardless the maturities of the yields on the left hand side of the equation (see Appendix 5 for a formal derivation). As a result, a statistical test for the rank of the estimated matrix $\hat{\mathcal{B}}_{\mathcal{Z}}$ is a test of the USV effect. The null hypothesis is

$$H_0 : \text{rank}\left(\mathcal{B}_{\mathcal{Z}}\right) \leq N - 1 \tag{1.9}$$

where N is the total number of factors. I use the Kleibergen-Paap rank test, among many other rank tests. The test statistic follows χ^2 distribution: for details, see Kleibergen and Paap (2006).

The approach has several benefits not shared by other tests for the USV effect in the literature. First, it is a formal statistical test - many of others in the literature are not formal statistical tests as pointed out by Bibkov and Chernov (2009). Second, while Bibkov and Chernov (2009) and Joslin (2015) conduct formal tests for the set of restrictions generating the USV effect, the USV restrictions are not unique as pointed out by Joslin (2015). For example, two different sets of restriction on the $A_1(4)$ specification can induce the USV effect while the two models fit volatilities in significantly different manners; see for example Creal and Wu (2015). The rank test that I posit here is free from this issue. Finally, the test can be implemented even in the presence of hidden factors as in Duffee (2011b) or Joslin, Priebsch, and Singleton (2014) - a detailed explanation of the hidden factors can also be found in Section 1.6. The test only exploits the cross-sectional relationship between the yields and variance factors, so the test results should be identical even after taking into account hidden factors.

Table 1.3 reports the test statistics for specifications with one through two volatility factors in conjunction with two through three additional non-volatility factors. Following Dai and Singleton (2000), $A_m\left(N\right)$ denotes an N factor model with m factor driving volatility.

Table 1.3
Tests for the USV effect

This table reports the test statistics of Kleibergen and Paap (2006) for the rank of the matrix $\mathcal{B}_{\mathcal{Z}}$ in equation (1.6). The null of the test is equation (1.9). Data is monthly and runs from October 1990 to December 2007.

Specification	Stat.	d.f.	C.V.(5%)	p-val
$A_1\,(3)$	20.82	9	16.92	0.01
$A_1\,(4)$	20.72	8	15.51	0.01
$A_2\,(4)$	14.90	8	15.51	0.06
$A_2\,(5)$	14.89	7	14.07	0.04

The data set is the same as the one in Section 1.4.1, and the first m PCs of the MA2 filtered TIV^{2}s are used as the variance factors for $A_m\,(N)$ models. Specifications with up to two volatility factors are considered for the exercise because the first two PCs of TIV^{2}s explain 99% variation of the three TIV^{2}s as pointed out in Section 1.4.1. The table shows that, for all the specifications, the null (the presence of USV effects) is rejected at the 10% significance level. Hence, the conditions for the affine models to generate USV effect do not hold in the data.

In sum, it is true that variance risks are hard to identify from the cross-section of yields as shown in Section 1.4.1, however, the knife-edge conditions for the USV effect are rejected in the data. In other words, the variance risk is effectively unspanned by yields not because of the USV restrictions but because of its limited impact on the shape of the yield curve, and it can be hardly identified without help of option prices.

1.5 A New Representation of ADTSM

In this section, I suggest a new representation of ADTSM in which all factors are represented as portfolios of bonds and options. The representation inherits the spirit of Joslin, Singleton, and Zhu (2011), and the advantages of their representation. Since all the term structure factors (including volatility) are observable, the estimation procedure becomes greatly simplified and economic interpretation of the model is more straightforward compared to conventional latent factor approaches.

For econometric identification, I initially assume that the risk-neutral dynamics of the latent factor in equation (1.1) is drift normalized as in Joslin (2015) or Creal and Wu (2015). The yield on a zero-coupon bond of maturity n is affine in the states Z_t:

$$y_{n,t} = A_n \left(\Theta^{\mathbb{Q}}\right) + B_n \left(\Theta^{\mathbb{Q}}\right) Z_t$$

where A_n and B_n are obtained from standard recursions as in Duffie and Kan (1996). I let $(n_1, n_2, ..., n_J)$ be the set of maturities of the bonds used in estimation and y_t be the $(J \times 1)$ vector of corresponding yields. For any full-rank matrix $W \in \mathbb{R}^{(N-m)\times J}$, Wy_t represents the associated $(N-m)$-dimensional set of portfolios of J $(\geq N)$ yields. Following Joslin, Singleton, and Zhu (2011), I let $\mathcal{P}_t$ denote the first $(N-m)$ principal components (PCs) of J yields with W being the weighting matrix of the PCs:

$$\mathcal{P}_t = Wy_t = A_W \left(\Theta^{\mathbb{Q}}\right) + B_W \left(\Theta^{\mathbb{Q}}\right) Z_t = A_W \left(\Theta^{\mathbb{Q}}\right) + B_{W,X} \left(\Theta^{\mathbb{Q}}\right) X_t + B_{W,V} \left(\Theta^{\mathbb{Q}}\right) V_t$$

Invoking Proposition 1, then, we can define the N observable pricing factors $\mathcal{Z}_t$ such that

$$\mathcal{Z}_t \equiv (\mathcal{P}_t', \mathcal{V}_t')' = \left((Wy_t)', \mathcal{V}_t'\right)' = U_0 + U_1 (X_t', V_t')' \tag{1.10}$$

where

$$U_0 = \begin{bmatrix} A_W \\ \alpha \end{bmatrix}, \quad U_1 = \begin{bmatrix} B_{W,X} & B_{W,V} \\ 0_{m\times(N-m)} & \beta \end{bmatrix}$$

with α and β defined in Proposition 1. The dynamic of $\mathcal{Z}_t$ can be represented as a function of the observable factor $\mathcal{Z}_t$ after applying the invariant rotation of Dai and Singleton (2002) to the latent factor Z_t. Provided that the mapping between $\mathcal{Z}_t$ and Z_t is bijective (i.e. one-to-one mapping), the model with observable $\mathcal{Z}_t$ is observationally equivalent to the representation with the latent Z_t. The sufficient condition for the mapping to be bijective is a full rank matrix β. Once the Gaussian factor X_t is drift normalized, Joslin (2015) shows that the matrix $B_{W,X}$ should be full rank. Hence, the first $(N-m)$ columns of U_1 are linearly independent. With non-zero β, the last m columns of U_1 are not spanned by the first $(N-m)$ columns of U_1, which implies that a full rank matrix β guarantees U_1 to be not rank deficient.

The new representation of ADTSM with the observable factors $\mathcal{Z}_t$ in equation (1.10) follows the idea of Joslin, Singleton, and Zhu (2011), henceforth JSZ, and can be considered an extension of their work into general affine models. JSZ suggests a new representation of Gaussian ADTSM in which all the Gaussian pricing factors are observable as portfolios of yields, i.e. $\mathcal{P}_t$ in equation (1.10). Through their representation, the estimation of the Gaussian term structure model is extremely simplified, and becomes more reliable in terms of finding a global optimum in maximum likelihood estimation. In particular, simple ordinary least square estimation (OLS) can be utilized to estimate the $\mathbb{P}$ conditional mean parameters of the pricing factors which had been treated as one of the most challenging parts in estimating term structure models due to the high degree of persistence in yields.

In my representation, all pricing factors (including the variance risk), are observable. As a consequence, one can make use of generalized least square estimation (GLS) to pin down the drift of the pricing factor under $\mathbb{P}$. Since variance is directly observational up to its linear transformation via $\mathcal{V}_t$, it is also easy to estimate the parameters governing the time-series dynamics of $\mathcal{V}_t$. In addition, when volatility risk is identified from the cross-section of yields, one should solve a numerically unstable equation $AX = b$ where A is often nearly singular, with the possibility that the solution leads to negative values for volatility: see for example Piazzesi (2010) and Joslin (2014). Instead, the representation I posit here is unaffected by this issue. In sum, the representation helps us find the global optimum of maximum likelihood estimation by simplifying the two hardest parts of the $ADTSM$ with stochastic volatility estimation, namely, the identification of volatility, as well as the drift of the state under $\mathbb{P}$.

The parameterization scheme using portfolios of yields as pricing factors for $A_m(N)$ model is also explored in Joslin and Le (2014), where the variance factor in $A_m(N)$ is approximated by portfolios of yields. The model I posit here utilize portfolios of options rather than portfolios of yields, and the variance factor is known before the model estimation while their variance factors can only be identified after the model estimation. In addition, the approach here is robust even in the presence of unspanned stochastic volatility factors as in Section 1.4.2, while their approach only works for spanned stochastic volatility. Furthermore, as discussed extensively in their paper, extracting the variance factor from yields only (without options) results in undesirable properties in the factor dynamics under $\mathbb{P}$ - this issue is discussed further in Section 1.7.2.

1.6 Unspanned Macro Risk and the Likelihood Function

More recently, a large literature has been studying so-called hidden factors or unspanned macro factors, see e.g., Duffee (2011b), Chernov and Mueller (2012) and Joslin, Priebsch, and Singleton (2014), henceforth JPS. A factor is described as hidden if it plays an important role in determining investors' expectations for future yields, yet is not priced in the fixed income market. Hence, the hidden factor cannot be recovered from the cross-section of any fixed income assets. This section explains how to take into account the hidden factor inside the model described in the previous section.

Since all priced factors are observable due to the representation in the previous section, the same argument as in JPS can be applied in order to add hidden factors in the framework. Once both hidden and non-hidden factors are projected onto the space of fixed income asset returns as in JPS, we get the following factor dynamics under the physical measure $\mathbb{P}$ and the risk-neutral measure $\mathbb{Q}$. First, the factors are composed of (i) the priced risks in the fixed income market $\mathcal{Z}_t = (\mathcal{P}_t', \mathcal{V}_t)'$ and (ii) a non-priced (hidden) factor M_t. In discrete time setting, the dynamics of the non-variance factors, $(\mathcal{P}_t', M_t)$, can be represented as

$$\begin{bmatrix} \mathcal{P}_{t+1} \\ M_{t+1} \end{bmatrix} = \begin{bmatrix} K_{0\mathcal{P}}^{\mathbb{P}} \\ K_{0M}^{\mathbb{P}} \end{bmatrix} + \begin{bmatrix} K_{\mathcal{P}\mathcal{P}}^{\mathbb{P}} & K_{\mathcal{P}M}^{\mathbb{P}} \\ K_{M\mathcal{P}}^{\mathbb{P}} & K_{MM}^{\mathbb{P}} \end{bmatrix} \begin{bmatrix} \mathcal{P}_t \\ M_t \end{bmatrix} + \begin{bmatrix} K_{\mathcal{P}\mathcal{V}}^{\mathbb{P}} \\ K_{M\mathcal{V}}^{\mathbb{P}} \end{bmatrix} \mathcal{V}_t + \begin{bmatrix} \Sigma_{\mathcal{P}\mathcal{V}} \\ \Sigma_{M\mathcal{V}} \end{bmatrix} \epsilon_{\mathcal{V},t+1}^{\mathbb{P}} + \begin{bmatrix} \epsilon_{\mathcal{P},t+1} \\ \epsilon_{M,t+1} \end{bmatrix} \tag{1.11}$$

with

$$\begin{aligned} \left(\epsilon_{\mathcal{P},t+1}', \epsilon_{M,t+1}'\right)' &\sim N\left(0, \Sigma_t\right) \\ \Sigma_t &= \Sigma_0 + \Sigma_1 \left(\mathcal{V}_t - \alpha\right) \\ \epsilon_{\mathcal{V},t+1}^{\mathbb{P}} &= \mathcal{V}_{t+1} - E_t\left(\mathcal{V}_{t+1}\right) \end{aligned} \tag{1.12}$$

The variance factor $\mathcal{V}_{t+1}$ follows a compound autoregressive gamma process

$$\mathcal{V}_{t+1} | \mathcal{V}_t \sim CAR\left(\rho^{\mathbb{P}}, c^{\mathbb{P}}, \nu^{\mathbb{P}}, \alpha\right)$$

where $c^{\mathbb{P}}$ is a scale parameter, $\nu^{\mathbb{P}}$ is a shape parameter, and $\rho^{\mathbb{P}}$ determines the autocorrelation of $\mathcal{V}_t$. The lower bound of $\mathcal{V}_t$ is set α, contrary to the standard lower bound of zero for a variance process. Indeed, the lower bound of the latent variance factor V_t should be set to zero for econometric identification. Then, the linear relationship between the observable $\mathcal{V}_t$ and the latent V_t, $\mathcal{V}_t = \alpha + \beta V_t$ with $\alpha\left(\Theta^{\mathbb{Q}}\right)$ defined in Proposition 1, implies that $\mathcal{V}_t$ should be greater than α. Furthermore, α should be positive since both $\alpha\left(\Theta^{\mathbb{Q}}\right)$ and $\beta\left(\Theta^{\mathbb{Q}}\right)$ capture the convexity components of yields - see equation (A-8) and Appendix 4.2 for details. The non-zero lower bound of $\mathcal{V}_t$ also leads Σ_t in equation (1.12) to be $\Sigma_0 + \Sigma_1\left(\mathcal{V}_t - \alpha\right)$ rather than $\Sigma_0 + \Sigma_1 \mathcal{V}_t$. For a detailed explanation of compound autoregressive processes, see Gourieroux and Jasiak (2006), Le, Singleton, and Dai (2010) and Creal and Wu (2015).

Under the pricing measure $\mathbb{Q}$, the dynamics of $\mathcal{Z}_t$ are assumed to be

$$\begin{aligned} \mathcal{P}_{t+1} &= K_{0\mathcal{P}}^{\mathbb{Q}} + K_{\mathcal{P}\mathcal{P}}^{\mathbb{Q}}\mathcal{P}_t + K_{\mathcal{P}\mathcal{V}}^{\mathbb{Q}}\mathcal{V}_t + \Sigma_{\mathcal{P}\mathcal{V}}\epsilon_{\mathcal{V},t+1}^{\mathbb{Q}} + \epsilon_{\mathcal{P},t+1} \\ \mathcal{V}_{t+1}|\mathcal{V}_t &\sim CAR\left(\rho^{\mathbb{Q}}, c^{\mathbb{Q}}, \nu^{\mathbb{Q}}, \alpha\right) \end{aligned} \tag{1.13}$$

Hence, the specification of the price of risks follows that in Cheridito, Filipović, and Kimmel (2007), and yields can be represented as a linear function of $\left(\mathcal{P}_t', \mathcal{V}_t\right)'$ where yields' loadings on the pricing factors are determined by $\Theta^{\mathbb{Q}}$ (see Appendix 4.1).

Furthermore, in order to maintain (i) the diffusion invariance property of the variance process $\mathcal{V}_t$ and (ii) non-exploding market price of risk in the continuous time limit (see Appendix B.4 in Joslin and Le (2014) for explanations), I impose the following two restrictions on parameters for $\mathcal{V}_t$:

$$c^{\mathbb{P}} = c^{\mathbb{Q}}, \quad \nu^{\mathbb{P}} = \nu^{\mathbb{Q}}$$

For the fitting of the cross-section, I assume that higher-order PCs, denoted by $\mathcal{P}_{e,t}$, are observed with i.i.d. uncorrelated Gaussian measurement errors with a common variance:

$$\mathcal{P}_{e,t}^{o} = \mathcal{P}_{e,t} + e_t \quad \text{and} \quad e_t \sim N\left(0, I\sigma_e^2\right)$$

In sum, the likelihood function of the observed data, $\mathcal{L}$, is

$$\mathcal{L} = \sum_t f\left(\mathcal{P}_{t+1}, M_{t+1}|\mathcal{V}_{t+1}, \mathcal{I}_t\right) + f\left(\mathcal{V}_{t+1}|\mathcal{V}_t, \mathcal{I}_t\right) + f\left(\mathcal{P}_{e,t+1}|\mathcal{P}_{t+1}, \mathcal{V}_{t+1}\right)$$

where f denotes the log conditional density. The first two terms capture the density of the time-series dynamics, and the last term is the density of the cross-sectional fit on which the unspanned macro factors M_t have no impact. Particularly, the $\mathbb{P}$-feedback matrix of $\mathcal{P}_t$ and M_t can be concentrated out by running GLS of the following system:

$$\begin{bmatrix} \mathcal{P}_{t+1} - \Sigma_{\mathcal{P}\mathcal{V}}\epsilon^{\mathbb{P}}_{\mathcal{V},t+1} \\ M_{t+1} - \Sigma_{M\mathcal{V}}\epsilon^{\mathbb{P}}_{\mathcal{V},t+1} \end{bmatrix} = \begin{bmatrix} K^{\mathbb{P}}_{0\mathcal{P}} \\ K^{\mathbb{P}}_{0M} \end{bmatrix} + \begin{bmatrix} K^{\mathbb{P}}_{\mathcal{P}\mathcal{P}} & K^{\mathbb{P}}_{\mathcal{P}M} \\ K^{\mathbb{P}}_{M\mathcal{P}} & K^{\mathbb{P}}_{MM} \end{bmatrix} \begin{bmatrix} \mathcal{P}_t \\ M_t \end{bmatrix} + \begin{bmatrix} K^{\mathbb{P}}_{\mathcal{P}\mathcal{V}} \\ K^{\mathbb{P}}_{M\mathcal{V}} \end{bmatrix} \mathcal{V}_t + \begin{bmatrix} \epsilon_{\mathcal{P},t+1} \\ \epsilon_{M,t+1} \end{bmatrix}$$

The observable variance $\mathcal{V}_t$ can be either spanned by yields or unspanned (i.e. of the unspanned stochastic volatility type as in Section 1.4.2). However, this does not affect the estimation procedure, since the volatility factor is identified not via yields but via options, even before the estimation procedure. In contrast, without the observable volatility factor, one should choose a specific set of restrictions on the $\mathbb{Q}$ parameters (among many possible set of restrictions), in order to estimate a model with unspanned stochastic volatilities. Otherwise, the identification of the volatility factor is infeasible, because it has no effect on the price of bonds.

In the term structure literature, both the unspanned stochastic volatility and hidden factors have been considered important components driving the time variation in risk premia, although their mechanisms are totally different. The effect of hidden factors on changes in risk premia exactly cancels out its effect on expectations of future short rate while USV implies a cancelation of the convexity bias. The USV factor can be identified from interest rate derivatives while hidden factors cannot be identified from any financial instrument in the market. To the best of my knowledge, my model is the first one capable of accommodating both types of unspanned risks: the unspanned stochastic volatility factors as well as the hidden factors.

1.7 Model Comparison

1.7.1 Model Specifications and Data

The discussion in Section 1.4 indicates that the variance risks' explanatory power for the cross-section of yield is fairly limited when it is compared to the explanatory power of the three term structure factors, level, slope and curvature. Under the representation in Section 1.5, this implies that at least three Gaussian factors are required to adequately explain the shape of the yield curve. Hence, I study a model with three yield factors and one stochastic volatility, which I denote by $A_1(4)$ as in Section 1.4.2. Its corresponding specification with two unspanned macro risks, denoted as $UMA_1^2(6)$, is also investigated; $UMA_m^R(N)$ stands for the family of ADTSMs in Section 1.6, with $(\mathcal{P}_t', M_t', \mathcal{V}_t')'$ of dimension N, M_t of dimension R, and $\mathcal{V}_t$ of dimension m. Following JPS, I use measures of economic activity and inflation as the two unspanned macro risks. In particular, the three-month moving average of the Chicago Fed's National Activity Index (CFNAI), henceforth denoted as GRO, is used as the measure of the *growth* in real economic activity as in JPS. However, I use year-over-year growth in Consumer Price Index excluding food prices and energy prices (henceforth, CPI) for the measure of inflation, contrary to JPS in which the measure of inflation is the expected rate of inflation from Blue Chip Financial Forecasts (henceforth, INF).

Moreover, the same $UMA_1^2(6)$ specification but with a different measure of economic activity - the unemployment gap - is also studied. The unemployment gap (henceforth $UGAP$) is the difference between the actual unemployment rate and the estimate of the natural rate of unemployment from the Congressional Budget Office (CBO). Hence, it gauges the *level* of economic activity rather than the *growth* of activity. Bauer and Rudebusch (2016), henceforth BR, argue that level indicators of activity such like $UGAP$ are largely related to the movement of the yield curves (i.e. weakly unspanned by yields) because these variables are relevant for setting the short-term policy rates; the authors also point out that the empirical monetary policy rules literature has identified level rather than growth variables as those which are most important for determining monetary policy (e.g. Taylor (1993), Taylor (1999), Orphanides (2003), Bean (2005) and Rudebusch (2006) among others). On the other hand, measures of growth in economic activity such as GRO are largely uncorrelated with the level of activity; see for example $UGAP$ and GRO in Figure 1.3. Furthermore, BR show that growth variables accompany low R^2s when (i) they are projected onto term structure factors or (ii) fed fund rates are regressed on them. Hence, they are strongly unspanned by yields. The different spanning properties of $UGAP$ and GRO induce significantly different estimates of risk premia for the $UMA_0^2(5)$ models in BR. BR qualitatively assess

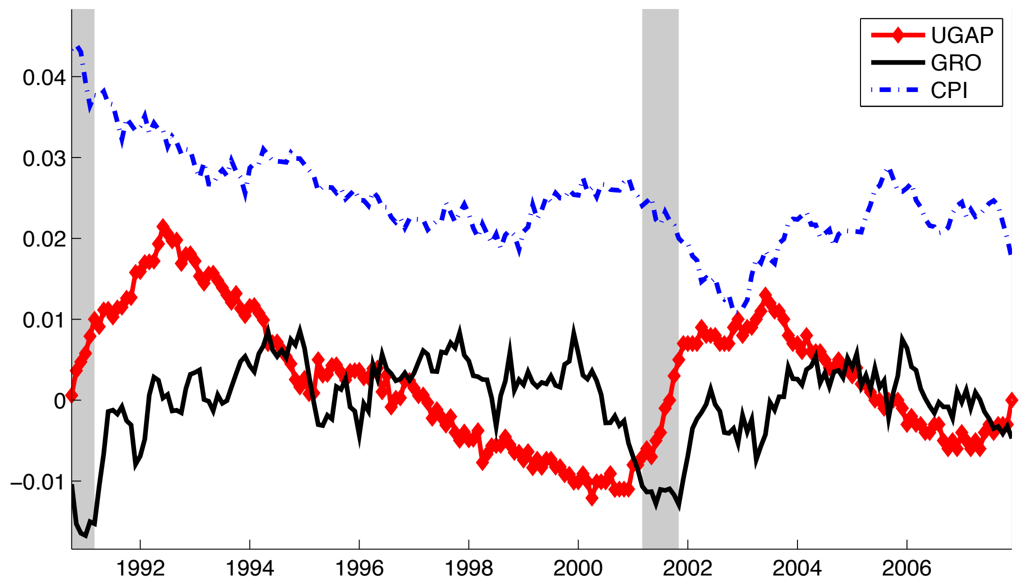

Figure 1.3. UGAP, GRO and CPI

This figure plots the unemployment gap ($UGAP$), the three-month moving average of the Chicago Fed's National Activity Index (GRO) and the year-over-year growth in Consumer Price Index, excluding food prices and energy prices (CPI). Gray bars indicate NBER recessions. The data is monthly and runs from October 1990 to December 2007.

the relevance of two different estimates of risk premia and claim that $UGAP$ is a better measure of economic activity. I access four unspanned models $UMA_0^2(5)$ and $UMA_1^2(6)$ and evaluate their relevance based on whether they can match the pattern for violation of the expectations hypothesis.

The yield data set is the same as that in Section 1.4, but I only use yields with maturities of six months, 1 through 3 years, 5, 7, 9 and 10 years. As a measure of observable volatility, I make use of the 30-year TIV.

1.7.2 The Campbell and Shiller Regression

The most well-known stylized fact in the fixed income market is the failure of the expectations hypothesis (see, for example, Fama and Bliss (1987) or Cambpell and Shiller (1991) among many others). As pointed out by Dai and Singleton (2002), this prominent pattern of return predictability can serve as a measure to access the goodness-of-fit of term structure models - one can investigate whether a model implied data-generating-process can re-produce the observed pattern in the data. In this section, I assess the ability of each model posited in Section 1.7.1 by comparing the extent to which each of them can match the important stylized fact of bond yields. The investigation reveals that the identification of a variance factor through the portfolio of options $\mathcal{V}_t$ strikingly enhances the models' ability to reproduce important patterns in the data. Furthermore, it is shown that the usage of different unspanned macro risks also determines the models' ability to generate the stylized fact. Hence, this property of models can serve as a useful guidance in selecting unspanned macro risks.

The expectation hypothesis implies that the changes in yields are solely attributed to the revision of future expected interest rates. As a result, high yield spreads should proceed increases in long rates, and changes in risk premia play no role in determining the shape of yield curves. One way of testing the expectations hypothesis is to regress realized changes in yields onto yield spreads

$$y_{n-h,t+h} - y_{n,t} = \phi_0^{(n)} + \phi_1^{(n)} \left(\frac{h}{n-h} \left(y_{n,t} - y_{h,t} \right) \right) \tag{1.14}$$

as in Cambpell and Shiller (1991). While $\phi_1^{(n)}$ should be one for $n > h$ under the null, the estimated $\phi_1^{(n)}$'s are typically negative and their magnitudes are increasing with n, i.e. for longer yields maturities. Dai and Singleton (2002) named this property of linear projections of yields as LPY, and suggested treating the projections as descriptive statistics that any empirically desirable term structure model should replicate. They find that the population coefficients $\phi_1^{(n)}$ implied by estimated Gaussian models closely match their data counterparts. In contrast, the affine models with stochastic volatilities are not capable of generating this pattern, and counter-factually imply that the expectations hypothesis nearly holds: $\phi_1^{(n)}$'s typically stay close to one across all maturities. In other words, the affine models with stochastic volatilities fail to match the key empirical relationship between expected returns and the slope of the yields curve. However, Almeida, Graveline, and Joslin (2011) document that the stochastic volatility models can be as good as Gaussian models in generating the LPY property, once options data are incorporated into the estimation procedure.

Table 1.4 reports the LPY property of each model. The stochastic volatility models with or without unspanned macro risks outperform the corresponding Gaussian models - this is in line with the findings in Almeida, Graveline, and Joslin (2011). The result first implies that the observational variance risk is beneficial in explaining the time variation in risk premia. Second, it indicates that the bond VIX is a properly identified measure of volatility risk in the Treasury market. As pointed out in Joslin and Le (2014), $\phi_1^{(n)}$ in equation (1.14) is mainly determined by the physical feedback matrix of factors. The population value of $\phi_1^{(n)}$ is

$$\phi_1^{(n)} = \frac{(n-h)}{h} \frac{\left(B_{n-h}\left(K_{1\mathcal{Z}}^{\mathbb{P}}\right)^h - B_n\right)\Sigma\left(B_n - B_h\right)'}{\left(B_n - B_h\right)\Sigma\left(B_n - B_h\right)'} \tag{1.15}$$

where Σ denote the unconditional covariance matrix of the time-series innovations and B_n is the yield's loadings on the observational factors $\mathcal{Z} = \left(\mathcal{P}_t', M_t', \mathcal{V}_t'\right)'$. The loadings are almost identical across all the models - the variance factor's marginal impact on the cross-section is minimal as shown in Section 1.4, and the yield's loadings on unspanned macro risks are zero by construction. Since the covariance matrix Σ is in both the numerator and denominator of equation (1.15), its impact cancels out. Hence, $K_{1\mathcal{Z}}^{\mathbb{P}}$ is the key that causes the variation in the LPY property of each model. In the case of Gaussian models, the physical feedback matrix of $\left(\mathcal{P}_t', M_t'\right)'$ is estimated by OLS. Hence, the estimates should be biased if the conditional volatility of $\mathcal{P}_t$ is time varying (which is strongly evident in the data). In the presence of $\mathcal{V}_t$, the bias of OLS estimates can be corrected because the physical feedback matrix of $\left(\mathcal{P}_t', M_t'\right)'$ is estimated by GLS. However, the correction works only if the the instrument of conditional volatility can truly resemble the data generating process. The outperformance provides evidence that the bond VIX is a well-identified measure of volatility risk in the Treasury market.

Figure 1.4 plots the $\phi_1^{(n)}$s from $A_0\,(3)$ and $A_1\,(4)$ models where six-month changes in yields are the dependent variables in the Campbell-Shiller regression. The $\phi_1^{(n)}$s from corresponding unspanned models, with GRO and CPI as the macro risks, are also displayed. They are notably worse than the three factor Gaussian model, $A_0\,(3)$. Figure 1.5 plots the same but with $UGAP$ as a measure of economic activity. Contrary to Figure 1.4, the unspanned models' performances are much improved, and it becomes even better than $A_0\,(3)$ model once the unspanned model incorporates the variance risk. Hence, the unemployment gap,

Table 1.4

Campbell-Shiller Regressions

Panel A reports the coefficients $\phi_1^{(n)}$ from the Campbell-Shiller regression in equation (1.14) with three month changes in yields as regressands. GRO and CPI are used to estimate $UMA_m^R(N)$ models. Panel B reports the coefficients $\phi_1^{(n)}$ from the Campbell-Shiller regression with six month changes in yields as regressands. $UGAP$ and CPI are used to estimate $UMA_m^R(N)$ models. Data is monthly and runs from October 1990 to December 2007.

Panel A: Campbell-Shiller Regression with Three-month Changes in Yields

Specification	Macro	Maturities										MSE
		1 yr	2 yr	3 yr	4 yr	5 yr	6 yr	7 yr	8 yr	9 yr	10 yr	
Data		0.40	-0.25	-0.67	-0.90	-1.03	-1.10	-1.15	-1.18	-1.20	-1.23	0.00
$A_0(3)$		0.40	0.01	-0.23	-0.37	-0.46	-0.53	-0.59	-0.64	-0.69	-0.74	2.31
$A_1(4)$		0.26	-0.29	-0.64	-0.84	-0.96	-1.07	-1.16	-1.25	-1.33	-1.41	0.09
$UMA_0^2(5)$	GRO,CPI	0.96	0.64	0.38	0.21	0.09	0.00	-0.06	-0.12	-0.17	-0.21	10.28
$UMA_1^2(6)$	GRO,CPI	0.96	0.60	0.27	0.04	-0.13	-0.25	-0.35	-0.43	-0.50	-0.57	6.45
$UMA_0^2(5)$	UGAP,CPI	0.56	0.16	-0.09	-0.25	-0.35	-0.42	-0.48	-0.54	-0.59	-0.64	3.44
$UMA_1^2(6)$	UGAP,CPI	0.33	-0.11	-0.40	-0.56	-0.67	-0.75	-0.83	-0.89	-0.96	-1.02	0.74

Panel B: Campbell-Shiller Regression with Six-month Changes in Yields

Specification	Macro	Maturities										MSE
		1 yr	2 yr	3 yr	4 yr	5 yr	6 yr	7 yr	8 yr	9 yr	10 yr	
Data		0.34	-0.39	-0.78	-0.98	-1.10	-1.17	-1.23	-1.27	-1.31	-1.34	0.00
$A_0(3)$		0.30	-0.03	-0.21	-0.30	-0.36	-0.41	-0.46	-0.50	-0.54	-0.59	4.36
$A_1(4)$		0.13	-0.37	-0.64	-0.78	-0.87	-0.96	-1.03	-1.11	-1.19	-1.26	0.29
$UMA_0^2(5)$	GRO,CPI	1.11	0.84	0.62	0.47	0.36	0.27	0.21	0.15	0.10	0.06	18.41
$UMA_1^2(6)$	GRO,CPI	1.06	0.76	0.46	0.25	0.10	-0.01	-0.11	-0.18	-0.25	-0.31	12.29
$UMA_0^2(5)$	UGAP,CPI	0.48	0.13	-0.07	-0.18	-0.24	-0.28	-0.33	-0.36	-0.40	-0.43	6.25
$UMA_1^2(6)$	UGAP,CPI	0.23	-0.18	-0.40	-0.51	-0.58	-0.64	-0.69	-0.75	-0.80	-0.85	2.04

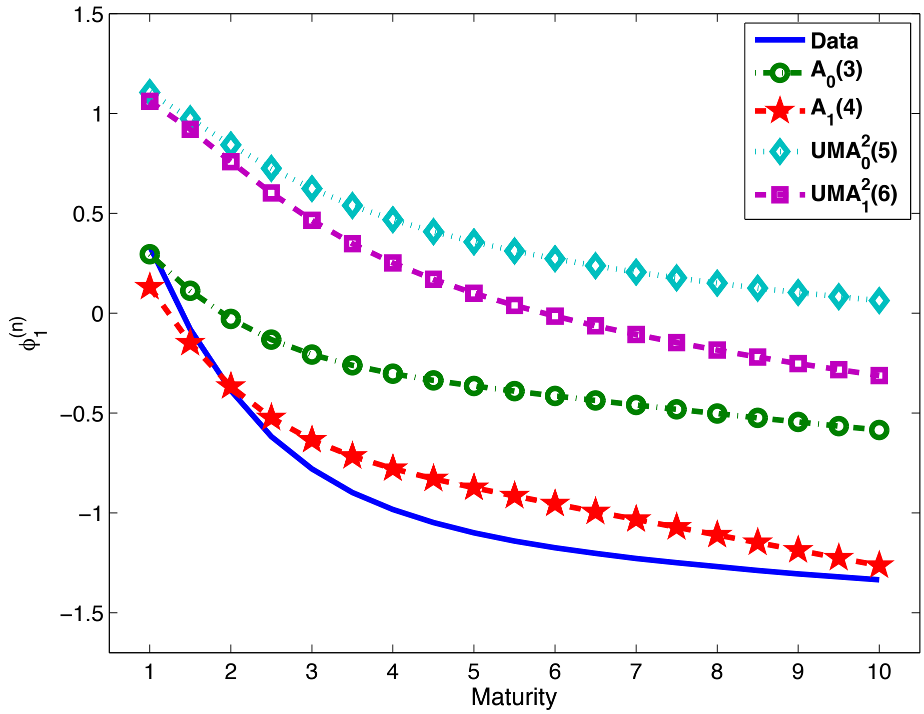

Figure 1.4. Campbell-Shiller Regression

This figure plots the coefficients $\phi_1^{(n)}$ from the Campbell-Shiller regression in equation (1.14) where the left-hand side variable is changes in yields over six months. The models $A_m(N)$ are models with $N-m$ Gaussian factors and m factors driving volatility. The models $UMA_m^R(N)$ are models with GRO and CPI as unspanned macro risks. Data is monthly and runs from October 1990 to December 2007.

a policy factor, can be considered a more relevant measure of real economic activity than GRO for a macro-finance term structure modeling. It also indicates that the impact of unspanned risk might not be as prominent as asserted by JPS in which CFNAI is utilized as a measure of output growth.

1.8 Risk Premia Accounting

This section studies the risk premia implied by $UMA_1^4(6)$ with UGAP and CPI as unspanned macro risks[9]. The risk premia on the risk factor $\mathcal{P}_t$ are the difference between the conditional expectation of $\mathcal{P}_{t+1}$ from the physical dynamics of equation (1.11) and the

[9]Since both UGAP and CPI are weakly unspanned, the model also can be treated as a shortcut of a spanned macro-finance term structure model as pointed out by Bauer and Rudebusch (2016).

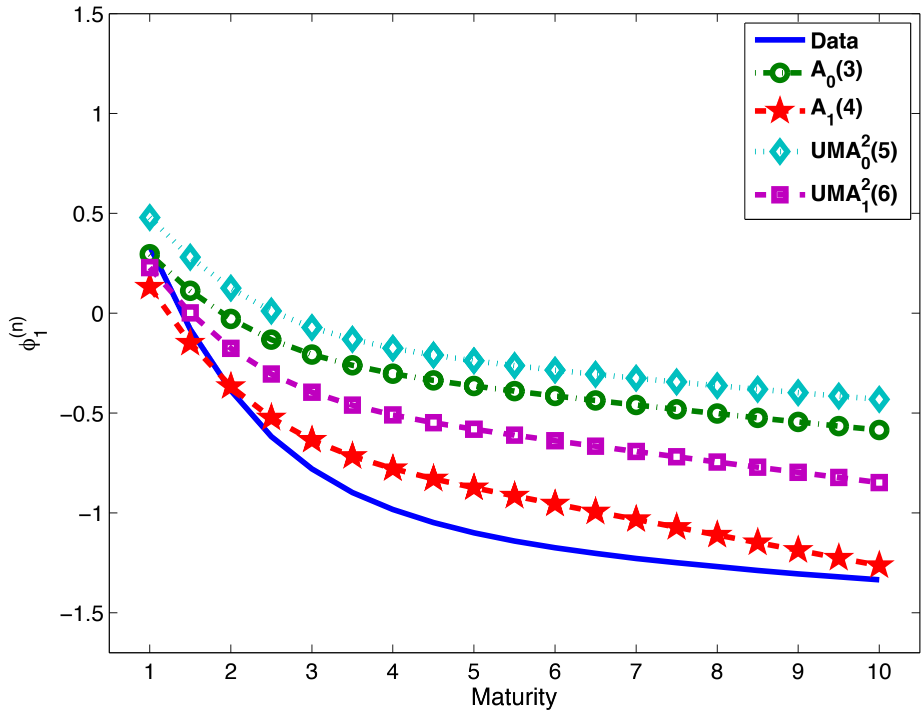

Figure 1.5. Campbell-Shiller Regression

This figure plots the coefficients $\phi_1^{(n)}$ from the Campbell-Shiller regression in equation (1.14) where the left-hand side variable is changes in yields over six months. The models $A_m(N)$ are models with $N-m$ Gaussian factors and m factors driving volatility. The models $UMA_m^R(N)$ are models with $UGAP$ and CPI as unspanned macro risks. Data is monthly and runs from October 1990 to December 2007.

risk-neutral dynamics of equation (1.13). They are determined by the full set of the state $\mathcal{Z}_t^* \equiv (\mathcal{P}_t', M_t', \mathcal{V}_t)'$ rather than solely by pricing factors $(\mathcal{P}_t', \mathcal{V}_t)'$:

$$E_t^{\mathbb{P}}(\mathcal{P}_{t+1}) - E_t^{\mathbb{Q}}(\mathcal{P}_{t+1}) = \left[K_{0\mathcal{P}}^{\mathbb{P}} - K_{0\mathcal{P}}^{\mathbb{Q}}\right] + \left[\begin{array}{ccc}\left(K_{1\mathcal{P}\mathcal{P}}^{\mathbb{P}} - K_{1\mathcal{P}\mathcal{P}}^{\mathbb{Q}}\right) & K_{1\mathcal{P}M}^{\mathbb{P}} & \left(K_{1\mathcal{P}\mathcal{V}}^{\mathbb{P}} - K_{1\mathcal{P}\mathcal{V}}^{\mathbb{Q}}\right)\end{array}\right]\mathcal{Z}_t^*$$

Furthermore, the risk premia on $\mathcal{P}_t$ are, to a first-order approximation, the scaled excess returns on the yield portfolios whose value change locally one-to-one with changes in $\mathcal{P}_t$. In other words, the first row of $E_t^{\mathbb{P}}(\mathcal{P}_{t+1}) - E_t^{\mathbb{Q}}(\mathcal{P}_{t+1})$ is the scaled excess return on the factor mimicking portfolio of $PC1$ while its value is unresponsive to changes in $PC2$, $PC3$, and $\mathcal{V}$ (see Appendix 6 for a detailed explanation - it extends the similar analysis of JPS for Gaussian unspanned models into affine models with stochastic volatilities.)

Table 1.5
Risk Premia Parameters

This table presents the maximum likelihood estimation of the parameters Λ_0 and Λ_1 governing expected excess returns on the PC-mimicking portfolios. Standard errors are given in parentheses.

$\mathcal{P}$	const	$PC1$	$PC2$	$PC3$	$UGAP$	CPI	TIV^2_{30yr}
$PC1$	0	-0.636	-1.305	0	0	1.647	1.509
		(0.213)	(0.326)			(0.650)	(0.683)
$PC2$	-0.018	0	0	0	0	0.699	0
	(0.007)					(0.265)	
$PC3$	0	0	0	0	0	0	0

Denoting these PC-mimicking portfolios as $xPC = \Lambda_0 + \Lambda_1 \mathcal{Z}_t^*$, I impose a set of zero restrictions on Λ_0 and Λ_1 due to the concerns about over-parameterization caused by the large number of parameters of the model (see, for example, Duffee (2010)). Furthermore, the constraint is economically interpretable since all the factors are tradable portfolios and macro variables. As pointed out by JPS (for their Gaussian models), no such model-free interpretation is feasible with a latent factor model. To figure out an adequate set of zero restrictions, I initially estimate the fully flexible version of a model in which no element of Λ_0 and Λ_1 is constrained to be zero. Then, the elements of the first estimates without statistical significance at 5% level are set as zero for the next step of estimation[10]. This procedure is repeated until I find that every non-constrained element of Λ_0 and Λ_1 is statistically significant. Table 1.5 displays the resulting estimates of Λ_0 and Λ_1. It indicates that exposure to both level and slope risks is priced, while exposure to curvature risk is not. Increase in the level of uncertainty, $\mathcal{V}_t$, induces higher expected return on the level mimicking portfolio even though $\mathcal{V}_t$'s impact on the cross-section of bonds is noticeably small as expected from the exercises in Section 1.4: see Appendix 7 for a detailed description of the risk-neutral parameters and cross-sectional fit of the model. Also, positive shocks on the measure of inflation (CPI) raise the risk premia on both level and slope risks. On the other hand, the level measure of economic activity, $UGAP$, does not contribute to the evolution of risk premia at all. Its unspanned component provides no relevant information for the

[10]Asymptotic standard errors are computed by numerical approximation to the Hessian and using the delta method.

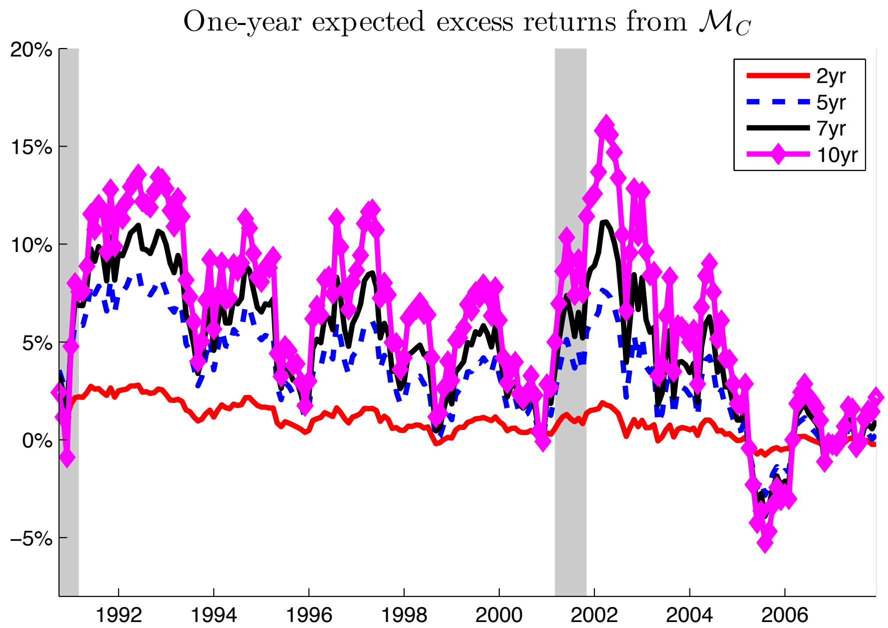

Figure 1.6. One-year Expected Excess Returns from $\mathcal{M}_C$

This figure plots the estimates of one-year simple expected excess returns from $\mathcal{M}_C$. The constraints on risk premia dynamics in Section 1.8 are imposed. The data is monthly and runs from October 1990 to December 2007.

time-variation in risk premia, because $UGAP$ is largely spanned by yield curve components as documented in Bauer and Rudebusch (2016).

Figure 1.6 plots the estimates of one-year simple expected excess returns from the model with constraints on Λ_0 and Λ_1 as in Table 1.5. Henceforth, this model is denoted as $\mathcal{M}_C$. For comparison purposes, I also estimate the preferred model in JPS (henceforth, $\mathcal{M}_J$). Based on information criteria, they conduct model selection searches over Λ_0 and Λ_1 of $UMA_0^2(5)$ with GRO and Blue Chip inflation forecasts as unspanned macro risks. One-year expected excess returns on 2-year and 10-year bonds from $\mathcal{M}_C$ and $\mathcal{M}_J$ are plotted in Figure 1.7. The term premia from $\mathcal{M}_C$ peak early in the recovery or near the end of the recession, and they are more volatile than the term premia from $\mathcal{M}_J$, especially for long-term bonds. For example, expected excess returns on a 10-year bond implied by $\mathcal{M}_J$, $E_t\left(RX^{(10yr)}_{t\rightarrow t+12}|\mathcal{M}_J\right)$, is much less time-varying than the expected excess return from $\mathcal{M}_C$, $E_t\left(RX^{(10yr)}_{t\rightarrow t+12}|\mathcal{M}_C\right)$. Table 1.6 reports R^2s from projecting one-year realized excess returns of n-year bonds onto their corresponding model implied expected excess returns from $\mathcal{M}_C$

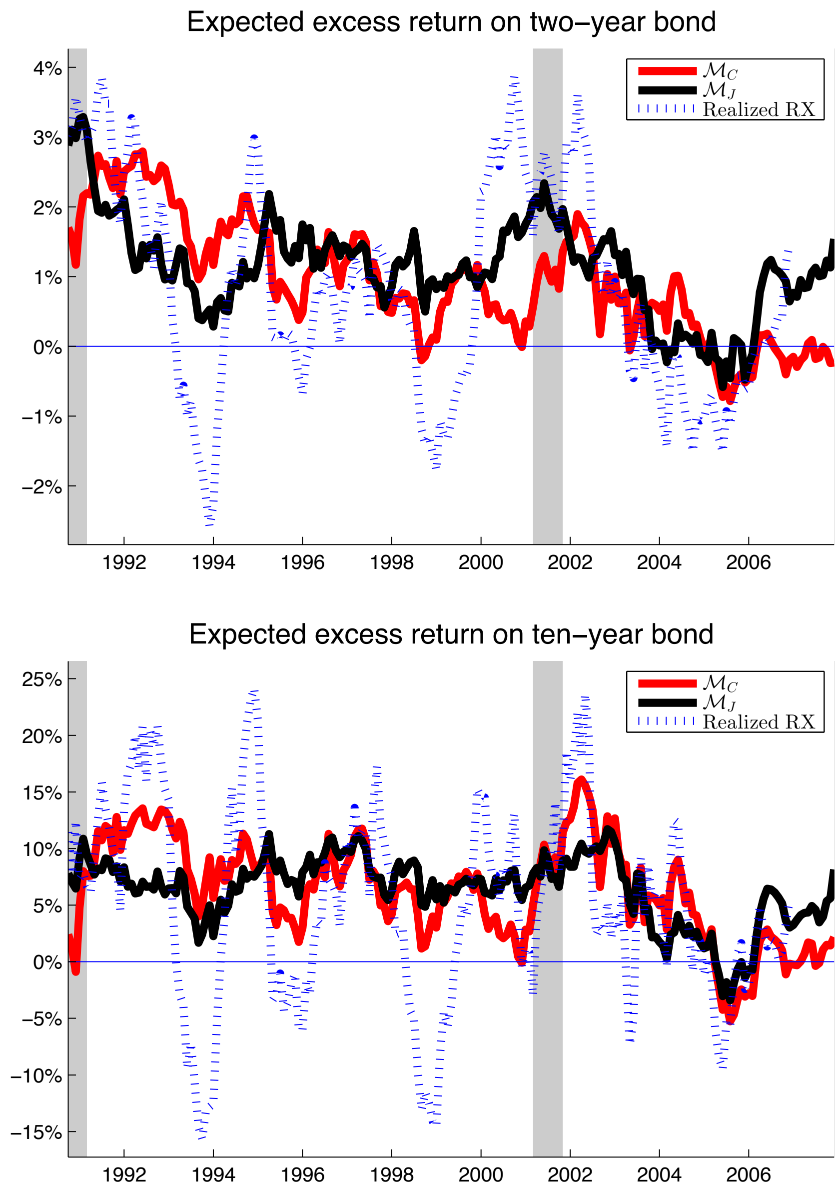

Figure 1.7. One-year Expected Excess Returns from $\mathcal{M}_C$ and $\mathcal{M}_J$

This figure plots the estimates of one-year simple expected excess returns from $\mathcal{M}_C$ and $\mathcal{M}_J$. The constraints on risk premia dynamics in Section 1.8 are imposed to estimate each specification. The data is monthly and runs from October 1990 to December 2007.

and $\mathcal{M}_J$. Expected excess returns from $\mathcal{M}_C$ explain, across all maturities, about 30% of time variation in realized excess returns. On the other hand, $\mathcal{M}_J$ is particularly good at capturing excess returns on short-term bonds, but its explanatory power diminishes along long-term bonds.

Table 1.6
Predictive Regressions

This table reports adjusted R^2 from regressing twelve-month excess returns xr(n) of bonds with n years to maturity on corresponding model implied expected excess returns.

Model	Specification	Macro	R^2			
			1 year	5 year	7 year	10 year
$\mathcal{M}_C$	$UMA_1^2(6)$	UGAP,CPI	0.26	0.28	0.29	0.30
$\mathcal{M}_J$	$UMA_0^2(5)$	GRO,CPI	0.38	0.22	0.18	0.14

The estimates of expected returns from June 2004 to June 2006 are of particular interest. During this period, the Federal Open Market Committee raised the policy rates 25 basis points for 17 consecutive meetings while the long-end of the yield curve remained relatively constant. The puzzling behaviour of long-rates has been labeled as a "conundrum" by the former Chairman Greenspan, and subsequent studies have attributed the phenomenon to declining risk premia. A comparison of the top and bottom panels in Figure 1.7 indicates that incorporating time-varying variance induces more a prominent reduction in risk premia during the conundrum period. Moreover, Figure 1.6 and Figure 1.7 show that *negative* one-year expected returns are associated with the conundrum period. The negative expected bond return, especially implied by $\mathcal{M}_C$, has an interesting implication for the design of structural asset pricing models. As shown in Martin (2015), any expected gross return R_T can be decomposed into

$$E_t^{\mathbb{P}}\left(R_T - R_{f,t}\right) = \frac{1}{R_{f,t}}\mathrm{var}_t^{\mathbb{Q}}\left(R_T\right) - \mathrm{cov}_t\left(\mathcal{M}_T R_T, R_T\right)$$

where $\mathcal{M}_T$ is the pricing kernel that prices time T payoffs from the perspective of time t. Hence, if R_T^{sp} is the return on the S&P 500 index and the second component of the decomposition, $\mathrm{cov}_t\left(\mathcal{M}_T R_T^{sp}, R_T^{sp}\right)$, is negative, then the risk-neutral variance of return, $\frac{1}{R_{f,t}}\mathrm{var}_t^{\mathbb{Q}}\left(R_T^{sp}\right)$, gives the lower bound on the equity premium. Martin (2015) also argues that $\mathrm{cov}_t\left(\mathcal{M}_T R_T^{sp}, R_T^{sp}\right)$ of the equity index is negative in most of macro-finance asset pricing models and estimates of covariance $\mathrm{cov}(\mathcal{M}_T R_T^{sp}, R_T^{sp})$ are negative across various sample periods. As a consequence, the measure of the risk-neutral variance constructed from S&P 500 index options can serve as the lower bound on the equity premium. On the other hand, the *negative* expected bond returns from $\mathcal{M}_C$ imply that $\text{cov}_t\left(\mathcal{M}_{C,T}R_T^{(n)}, R_T^{(n)}\right)$ for n-year bond return $R_T^{(n)}$ can take a positive value even after explicitly incorporating $\text{var}_t^{\mathbb{Q}}(R_T)$ into the stochastic discount factor. Hence, in a desirable equilibrium asset pricing model, $\text{cov}_t(\mathcal{M}_T R_T, R_T)$ for bonds should be able to switch its sign while $\text{cov}_t(\mathcal{M}_T R_T, R_T)$ for the equity market remains negative. Given a preference of representative agent, for example, this condition gives a clue for the factor structure of the state. Alternatively, it can be utilized to restrict the parameter space of an equilibrium model.

1.9 Conclusion

This paper studies the impact of variance risk in the Treasury market on both term premia and the shape of the yield curve. Variance risk in the Treasury market can be observed via a portfolio of options given the assumptions that (i) the state of the economy is determined by a state variable following an affine diffusion process under the risk-neutral measure and (ii) the drift of a pricing kernel is affine in the state variable. This unique approach for the identification of variance risk in interest rates enables me to treat the bond VIX^2 as a measure of fixed income variance risk.

Using the observable proxy of variance risk, labeled bond VIX^2, this paper first proposes a novel return-forecasting factor. The return-forecasting factor is motivated by leading consumption-based asset pricing models, the long-run risk, rare disaster, and affine$^{\mathbb{Q}}$ habit formation models. In these frameworks, the set of risk factors underlying variation in risk premia is the sole source of time-varying variances in bond yields and should be captured by the bond VIX^2s. Projection of realized excess bond returns onto the space of VIX^2s gives a single return-forecasting factor that describes time-variation in the expected bond returns. The return-forecasting factor predicts excess returns of relatively short-term bonds well, and complements the Cochrane-Piazzesi factor.

Second, the observable measure of variance risk can be utilized to analyze the relationship between variance risk and the shape of the yield curve in a simple and parsimonious way. Its marginal impact on the cross-section of bonds is limited once I control for standard term structure factors. Furthermore, the hypothesis of unspanned stochastic volatility can be directly assessed by testing the spanning conditions of affine models, and I find the hypothesis is rejected. In sum, though the knife-edge conditions for USV effects do not hold in the data, it is true that identifying variance risk from the yield curve movements is very hard, and the variance risk can be effectively considered as unspanned risk.

Third, I propose a new representation of affine dynamic term structure models with time-varying variance risks in yields. Due to the observable proxy of variance risk, affine bond pricing models can be represented by observable and tradable factors. This simplifies the estimation procedure significantly while the information in volatility-sensitive instruments is readily incorporated. The estimated risk premia show that it is important to take into account both types of unspanned risk: the unspanned stochastic volatility factor and the hidden non-volatility factor.

Chapter 2

Asset Prices and Pricing Measures with Alternative Investment Horizons

2.1 Introduction

Asset prices contain enormous amount of information about the stochastic discounting of possible future states, the pricing kernel or the stochastic discount factor (henceforth, SDF). For example, option prices can be utilized to extract the pricing kernel as in Aït-Sahalia and Lo (2000), Rosenberg and Engle (2002), and Ross (2015), among many others. Alternatively, the lower bound for the variance of the SDF can be recovered from asset prices as in Hansen and Jagannathan (1991). The prices of zero-coupon bonds also help to characterize the properties of the SDF at multiple intermediate horizons: see Backus, Chernov, and Zin (2014).

This paper also utilizes asset prices to identify the SDF, but I propose a particular zero-cost trading strategy, named strategy-**F**, that distinctively reveals the information on the dynamics of the pricing kernel. Through the strategy, the investor holds a long-term bond by borrowing at a series of short term interest rates, and importantly the payoff of the strategy mimics the Radon-Nikodym derivative between two pricing measures with alternative investment horizons (for example, the risk-neutral measure and a T-forward measure). The return on the strategy is hence characterized by the term structure of the SDF (or the dynamics of the SDF), and as a result, imposes a distinctive restriction on the set of admissible SDFs that price the strategy correctly.

I then incorporate the strategy-**F** into non-parametric estimation schemes for the SDF, specifically the Hasen-Jagannathan pricing kernel in Hansen and Jagannathan (1991) and the information kernel in Ghosh, Julliard, and Taylor (2016b). The non-parametrically estimated SDFs (henceforth, filtered SDFs) are then utilized to assess the benefit of the strategy in identifying the underlying SDF. In particular, I first compare the cross-sectional pricing abilities of the filtered SDFs with/without taking into account the strategy for their estimation procedure. With respect to several alternative measures to evaluate the SDFs' ability for the task, the SDFs characterized by the strategy-**F** (in conjunction with other asset returns) are better than the corresponding SDFs estimated without the strategy-**F**.

Given that the validity of the strategy-**F** is supported by the cross-sectional asset pricing exercises, I then assess the extent to which the strategy raises various lower bounds of asset pricing models. In particular, I show that incorporating the strategy-**F** into the Euler equation remarkably raises the Hansen-Jagannathan bound in Hansen and Jagannathan (1991) and various entropy bounds in Ghosh, Julliard, and Taylor (2016b). That is, the strategy-**F** helps to significantly tighten the space of admissible SDFs, and as a result it sets a more stringent hurdle for the equilibrium models.

The paper proceeds as follows. Section 2.2 theoretically shows how the particular strategy is related to the Radon-Nikodym derivative between the risk-neutral measure and a forward measure, hence the dynamics of the pricing kernel. Section 2.3 argues how one can incorporate the information from the strategy-**F** into non-parametric extraction of the pricing kernel. Section 2.4 analyzes the estimated pricing kernel, and Section 2.5 assesses their cross-sectional pricing abilities. Section 2.6 explores the new Hansen-Jagannathan bound and the entropy bounds stemming from the strategy-**F** . Finally, Section 2.7 concludes. All proofs are deferred to the Appendices.

2.2 Theory

Under the assumption of no arbitrage opportunity, there exists a positive pricing kernel (also known as a stochastic discount factor), $M_{t,t+n}$, that satisfies

$$E_t^{\mathbb{P}}\left(M_{t,t+n}R_{t,t+n}\right) = 1, \tag{2.1}$$

for any positive time interval n, and any vector of returns $R_{t,t+n} \in \mathbb{R}^N$. Here, $R_{t,t+n}$ is the gross return on traded assets over the period t to $t+n$. The conditional expectation of $M_{t,t+n}$ is denoted as $P_t^{(n)}$ which should be the date t price of a zero coupon bond with maturity of n. Under weak regularity conditions, dividing both sides of equation (2.1) with $P_t^{(n)}$

$$E_t^{\mathbb{P}}\left(\frac{M_{t,t+n}}{P_t^{(n)}}R_{t,t+n}\right) = \frac{1}{P_t^{(n)}},$$

yields an expression for the Radon-Nikodym derivative of $\mathbb{Q}^{(n)}$ with respect to $\mathbb{P}$

$$\frac{d\mathbb{Q}^{(n)}}{d\mathbb{P}} = \frac{M_{t,t+n}}{E_t^{\mathbb{P}}\left(M_{t,t+n}\right)} \tag{2.2}$$

where $\mathbb{Q}^{(n)}$ is the pricing measure of which numeraire is $P_t^{(n)}$. Equation (2.2) also gives us a representation of the Radon-Nikodym derivative of $\mathbb{Q}^{(n)}$ with respect to $\mathbb{Q}$,

$$\frac{d\mathbb{Q}^{(n)}}{d\mathbb{Q}} = \frac{d\mathbb{Q}^{(n)}}{d\mathbb{P}}\frac{d\mathbb{P}}{d\mathbb{Q}} = \frac{M_{t,t+n}}{P_t^{(n)}}\frac{P_t^{(1)}}{M_{t,t+1}} = M_{t+1,t+n}\frac{P_t^{(s)}}{P_t^{(n)}}, \tag{2.3}$$

where $M_{t+1,t+n} \equiv \prod_{i=0}^{n-1} M_{t+i,t+i+1}$. Equation (2.3) shows that, at time t, the upcoming evolution of the pricing kernel from time $t+1$ to time $t+n$, $\prod_{i=0}^{n-1} M_{t+i,t+i+1}$, is fully reflected to the Radon-Nikodym derivative $d\mathbb{Q}^{(n)}/d\mathbb{Q}$. Hence if we observe the realization of $d\mathbb{Q}^{(n)}/d\mathbb{Q}$, it can be an informative channel to infer the dynamics of the one-period pricing kernel $M_{t,t+1}$. As pointed out by Backus, Chernov, and Zin (2014), the dynamics of the pricing kernel $M_{t,t+1}$ convey important and useful information. The novel approach of this paper is to utilize the fact that the realization of $d\mathbb{Q}^{(n)}/d\mathbb{Q}$ can be observed from the price of zero coupon bonds (or zero coupon bond yields), because of the following equation:

$$\frac{1}{P_t^{(n)}} \equiv E_t^{\mathbb{Q}^{(n)}}\left(R_{t+n}\right) \equiv E_t^{\mathbb{Q}}\left(\frac{d\mathbb{Q}^{(n)}}{d\mathbb{Q}}R_{t+n}\right) = E_t^{\mathbb{Q}}\left(\frac{\prod_{i=0}^{n-1} P_{t+i}^{(1)}}{P_t^{(n)}}R_{t+n}\right)$$

In other words,

$$\frac{d\mathbb{Q}^{(n)}}{d\mathbb{Q}} = \frac{\prod_{i=0}^{n-1} P_{t+i}^{(1)}}{P_t^{(n)}} \tag{2.4}$$

which implies that $d\mathbb{Q}^{(n)}/d\mathbb{Q}$ is about the expected path of the pricing kernel and its realization can be observed by the price of a one-period zero coupon bond (or the risk-free rate). In addition, the logarithm of equation (2.4) gives a log return of the zero-costing strategy in which the investor holds a long-term bond by borrowing at a series of short-term interest rates

$$\ln \frac{d\mathbb{Q}^{(n)}}{d\mathbb{Q}} = ny_t^{(n)} - \sum_{i=0}^{n-1} y_{t+i}^{(1)}.$$

The conditional expectation of the logarithm of $d\mathbb{Q}^{(n)}/d\mathbb{Q}$ is the term premium for the zero bond $P_t^{(n)}$. Hence, the Radon-Nikodym derivative $d\mathbb{Q}^{(n)}/d\mathbb{Q}$ is also closely related to term premia in the term structure literature.

Note that the conditional expectation of the Radon-Nikodym derivative $d\mathbb{Q}^{(n)}/d\mathbb{Q}$ should be 1 under the risk-neutral measure $\mathbb{Q}$, and this implies[1] the following restriction on the pricing kernel $M_{t,t,+1}$

$$E_t \left[M_{t,t+1} \left(\frac{\prod_{i=1}^{n-1} P_{t+i}^{(1)}}{P_t^{(n)}} - 1 \right) \right] = 0, \tag{2.5}$$

which helps to identify the dynamic property of the pricing kernel. The vast literature exploits the unconditional form of equation (2.1) to extract the stochastic discount factor from asset prices in a model-free manner. On the other hand, this paper *jointly* utilizes the unconditional version of equation (2.5) as well as equation (2.1) to construct the pricing kernel from asset returns in a non-parametric way. In particular, I incorporate the two restrictions into the extraction of (i) the Hansen-Jaganathan minimum variance kernel in Hansen and Jagannathan (1991) and (ii) the information kernel in Ghosh, Julliard, and Taylor (2016b). The approach that I posit here is that of unifying the information contents of (i) measure of dispersion entropy bound and (ii) horizon dependence in Backus, Chernov, and Zin (2014).

2.3 Methodology

Since the seminal work by Hansen and Jagannathan (1991), the literature has developed approaches to extract a stochastic discount factor (hereafter, SDF) from a given set of asset returns: for example, Ferson and Siegel (2003), Bekaert and Liu (2004), Chabi-Yo (2008)

[1] It is straightforward to show that

$$E_t^{\mathbb{Q}} \left(\frac{d\mathbb{Q}^{(n)}}{d\mathbb{Q}} - 1 \right) = E_t \left[\frac{d\mathbb{Q}}{d\mathbb{P}} \left(\frac{d\mathbb{Q}^{(n)}}{d\mathbb{Q}} - 1 \right) \right] = E_t \left[\frac{M_{t,t+1}}{E_t\left(M_{t,t+1}\right)} \left(\frac{\prod_{i=1}^{n-1} P_{t+i}^{(1)}}{P_t^{(n)}} - 1 \right) \right] = 0,$$

and equation (2.5) comes from the last equality.

and Ghosh, Julliard, and Taylor (2016b). This section first demonstrates how to incorporate the theory in Section 2.2 into the Hansen-Jagannathan minimum variance kernel in Hansen and Jagannathan (1991), and into the information kernel in Ghosh, Julliard, and Taylor (2016b). It then compares their properties with/without the new restriction in Section 2.2.

2.3.1 Hansen-Jagannathan minimum variance kernel

The canonical Hansen-Jagannathan pricing kernel is the minimum variance pricing kernel among admissible SDFs that perfectly price a given set of asset returns to construct it. More formally, for a given value of $E(M_t) = \overline{M}$, the Hansen-Jagannathan minimum variance kernel is

$$M_t^{HJ} \equiv \underset{\left\{M_t\left(\overline{M}\right)\right\}_{t=1}^{T}}{\arg\min}\ Var\left(M_t\right) \quad \text{s.t.} \quad E^{\mathbb{P}}\left(M_t \mathbf{R}_t^e\right) = \mathbf{0} \tag{2.6}$$

where $\mathbf{R}_t^e$ denotes a set of asset returns that defines the 'admissibility' of SDFs - among all the SDFs that are orthogonal to $\mathbf{R}_t^e$, the canonical Hansen-Jagannathan kernel (henceforth, I-SDF) is the one that achieves the smallest variance.

2.3.2 Information kernel in Ghosh, Julliard and Taylor (2016)

While the Hansen-Jagannathan minimum variance kernel minimizes the second moment deviation, the information kernel in Ghosh, Julliard, and Taylor (2016b) minimizes Kullback-Leibler Information Criterion (KLIC) divergence between the physical and the risk-neutral measure. More precisely, for each $E(M_t) = \overline{M}$, their information kernel is

$$M_t^I \equiv \underset{\{M_t(\overline{M})\}_{t=1}^{T}}{\arg\min}\ E^{\mathbb{P}}\left(M_t \ln M_t\right) \quad \text{s.t.} \quad E^{\mathbb{P}}\left(M_t \mathbf{R}_t^e\right) = \mathbf{0} \tag{2.7}$$

Because of $\frac{M_t}{\overline{M}} = \frac{d\mathbb{Q}}{d\mathbb{P}}$, it is straightforward to show that the optimization in equation (2.7) is equivalent to

$$\underset{\mathbb{Q}}{\arg\min}\, D\left(\mathbb{Q}||\mathbb{P}\right) = \underset{\mathbb{Q}}{\arg\min} \int \ln\left(\frac{d\mathbb{Q}}{d\mathbb{P}}\right) d\mathbb{Q} \quad \text{s.t.} \quad \int \mathbf{R}_t^e d\mathbb{Q} = \mathbf{0}$$

where $D\left(\mathbb{A}||\mathbb{B}\right) \equiv \int \ln\left(\frac{d\mathbb{A}}{d\mathbb{B}}\right) d\mathbb{A} \equiv \int \frac{d\mathbb{A}}{d\mathbb{B}} \ln\left(\frac{d\mathbb{A}}{d\mathbb{B}}\right) d\mathbb{B}$ denotes the Kullback-Leibler Information Criterion (KLIC) divergence between $\mathbb{A}$ and $\mathbb{B}$ (or the relative entropy of $\mathbb{A}$ with respect to $\mathbb{B}$). Hence, the information kernel (henceforth, I-SDF) is the minimum relative entropy (between the physical and the risk-neutral measure) among admissible SDFs that price $\mathbf{R}_t^e$ perfectly (or are orthogonal to the space of $\mathbf{R}_t^e$).

2.3.3 *Strategy-*F

For ease of explanation, I refer to the strategy that mimics the evolution of the Radon-Nikodym derivatives $\frac{d\mathbb{Q}^{(n)}}{d\mathbb{Q}}$ in equation (2.4) as 'strategy-**F**', and denote its excess return as $\mathbf{F}_t^e$:

$$\mathbf{F}_t^e \equiv \frac{\prod_{i=1}^{n-1} P_{t+i}^{(1)}}{P_t^{(n)}} - 1 \tag{2.8}$$

As argued in Section 2.2, $\mathbf{F}_t^e$ reveals the dynamics of the pricing kernel, and moreover it is orthogonal to M_t: see equation (2.5). Hence, it can be readily incorporated into the estimation procedure of the HJ-SDF and the I-SDF. The orthogonality of M_t with respect to $\mathbf{F}_t^e$ implies

$$E^{\mathbb{P}}\left(M_t \mathbf{F}_t^e\right) = 0, \tag{2.9}$$

and the HJ-SDF and the I-SDF with the strategy-**F** can be respectively implemented by

$$M_t^{HJ,\mathbf{F}} \equiv \underset{\left\{M_t(\overline{M})\right\}_{t=1}^T}{\arg\min} \; Var\left(M_t\right) \quad \text{s.t.} \quad E^{\mathbb{P}}\left(M_t \mathbf{R}_t^e\right) = \mathbf{0} \text{ and } E^{\mathbb{P}}\left(M_t \mathbf{F}_t^e\right) = 0$$

and

$$M_t^{I,\mathbf{F}} \equiv \underset{\{M_t(\overline{M})\}_{t=1}^T}{\arg\min} \; E^{\mathbb{P}}\left(M_t \ln M_t\right) \quad \text{s.t.} \quad E^{\mathbb{P}}\left(M_t \mathbf{R}_t^e\right) = \mathbf{0} \text{ and } E^{\mathbb{P}}\left(M_t \mathbf{F}_t^e\right) = 0$$

Thus, simply expanding the return space (to construct SDFs) from $\{\mathbf{R}_t^e | t = 1, ..., T\}$ to augmented $\{\mathbf{R}_t^e, \mathbf{F}_t^e | t = 1, ..., T\}$ enables the filtered SDFs to explicitly take into account the properties in the dynamics of the pricing kernel. In other words, the space of admissible SDFs is tightened from $\{\mathbf{R}_t^e | t = 1, ..., T\}^{\perp}$ to $\{\mathbf{R}_t^e, \mathbf{F}_t^e | t = 1, ..., T\}^{\perp}$ in order to explain the realization of $\frac{d\mathbb{Q}^{(n)}}{d\mathbb{Q}}$ in the data.

2.4 Pricing Kernels and Strategy-F

This section non-parametrically extracts *real* SDFs with/without strategy-**F** via the methodology in Hansen and Jagannathan (1991) and Ghosh, Julliard, and Taylor (2016b). Then, the impact of introducing strategy-$\mathbf{F}$ to the estimation procedure is examined. The restriction from strategy-$\mathbf{F}$, i.e. $E^{\mathbb{P}}\left(M_t \mathbf{F}_t^e\right) = 0$, causes noticeable changes, especially the persistency of filtered SDFs, and also increases the KLIC of information SDF.

2.4.1 Data

To construct real HJ-SDF and I-SDF, I utilize return space consisting of both excess returns on zero-coupon bonds and several equity portfolios. In particular, 3 and 8 year bond returns are calculated using the zero-coupon bond yields of Gürkaynak, Sack, and Wright (2007). For the equity portfolios, I make use of the 6 size and book-to-market-equity sorted portfolios of Fama-French, 10 industry-sorted portfolios, and 10 momentum sorted portfolios. The data can be obtained from Kenneth French's data library at a monthly frequency, and the quarterly returns on these data are formed by compounding the monthly returns within each quarter. Excess returns are then computed by subtracting the corresponding risk-free rate. I use the three-month Treasury bill rate from the Center for Research in Security Prices (CRSP) as a proxy for the risk-free rate, and the three-month rate is also used to form the excess returns on strategy-$\mathbf{F}$. The nominal returns are then converted into real returns using the Consumer Price Index (CPI). I also set the holding-period of the strategy-$\mathbf{F}$ at two years (hence, n in equation (2.8) is equal to 8 at a quarterly frequency). The data runs from 1963:1 to 2013:4 for the SDFs without strategy-$\mathbf{F}$, whereas the period between 2014:1 and 2015:4 is covered for the SDFs with strategy-$\mathbf{F}$ in order to observe realized excess returns on the strategy, $\mathbf{F}_t^e$.

2.4.2 Estimated pricing kernels

Figure 2.1 plots the filtered kernels with/without incorporating the strategy-$\mathbf{F}$ into their asset return spaces. Their levels are pinned down by matching their means to the average 3-month T-Bill rate - the filtered time-series of HJ-SDF exhibits frequently observed negative values due to this normalization. One can also observe more prominent spikes in the filtered time-series after introducing the strategy-$\mathbf{F}$ into the asset return space.

Table 2.1 presents the summary statistics of the SDFs. The most noticeable difference can be found in their first-order serial correlations, as the strategy-$\mathbf{F}$ conveys information on

Figure 2.1. Filtered Pricing Kernels

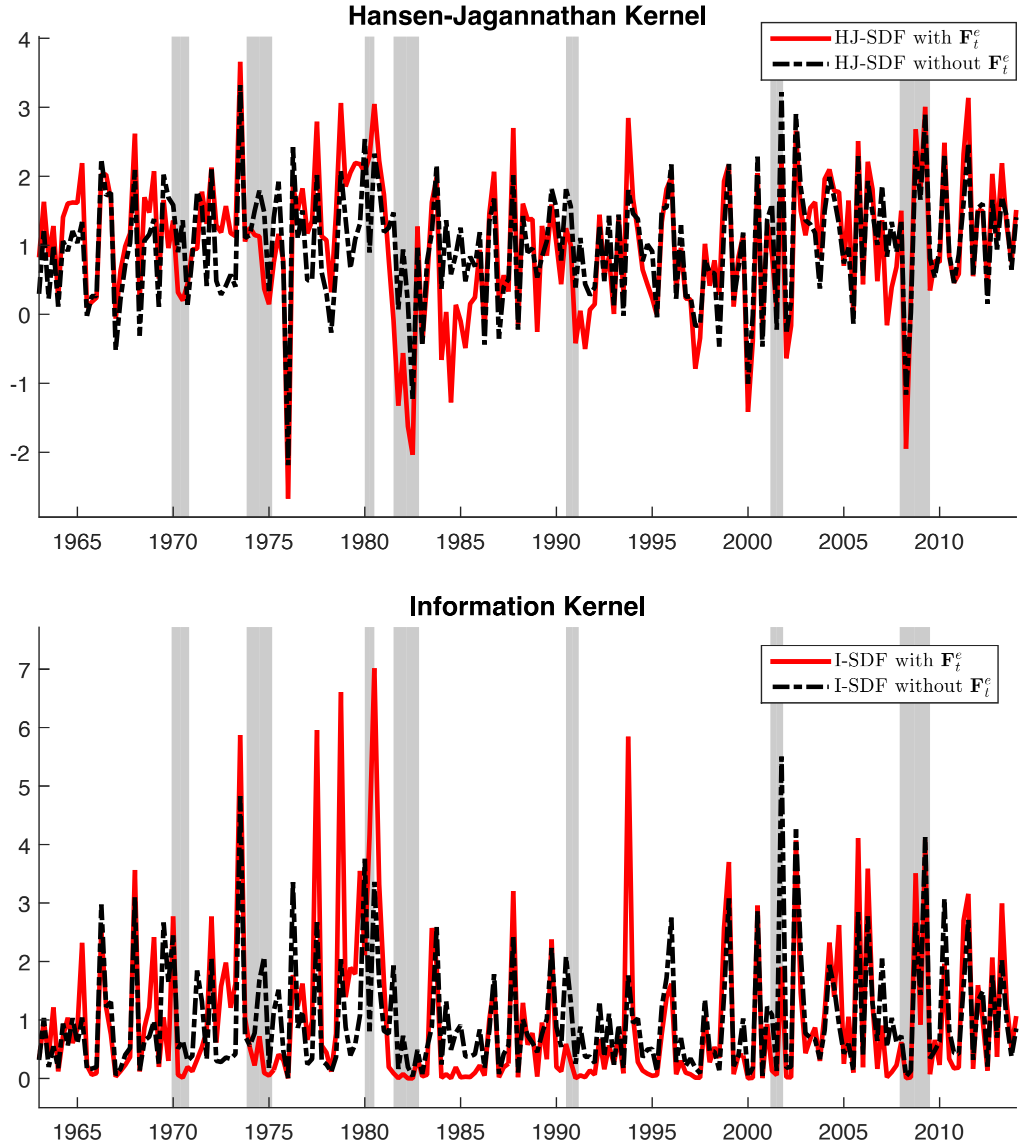

This figure plots the estimated HJ-SDFs and I-SDFs with/without the strategy-**F** i.e. with/without the restriction in (2.9). Data is quarterly and runs from 1963:1 to 2015:4.

the dynamics of the SDF. Imposing the condition that an SDF perfectly prices the strategy-**F** (i) induces the estimated SDF to switch the sign of their auto-correlation coefficients from negative to positive numbers, and (ii) drastically increases the persistency of each HJ-SDF and I-SDF. Second, the SDFs with the strategy-**F** are more volatile and more skewed, while only the I-SDF becomes more heavy-tailed after reflecting the strategy-**F**. Finally,

Table 2.1
Summary Statistics of Filtered Kernels

This table reports the summary statistics of filtered kernels. Data is quarterly and runs from 1963:1 to 2015:4.

	Strategy-**F**	Mean	Std	Min	Max	Skew	Kurt	A.C.	KLIC
HJ-SDF	No	0.99	0.83	-2.19	3.32	-0.21	3.75	-0.07	
	Yes	0.99	1.01	-2.67	3.66	-0.41	3.70	0.23	
I-SDF	No	0.99	0.97	0.01	5.50	1.81	6.63	-0.09	0.40
	Yes	0.99	1.32	0.00	7.01	2.14	8.03	0.15	0.70

the relative entropy of $\mathbb{Q}$ with respect to $\mathbb{P}$ becomes twice larger than before once equation (2.9) is incorporated into the Euler equation.

2.5 Cross-Sectional Pricing

The summary statistics of the Hansen-Jagannathan minimum variance kernels and information kernels in Table 2.1 exhibit noticeable differences once we take into account the strategy-**F**. The difference itself cannot tell us much, and it matters only to the extent that it has meaningful asset pricing implications. This section examines whether the differences entail meaningful asset pricing implications, especially in the context of cross-sectional pricing of asset returns. In other words, I assess the cross-sectional pricing performance of the two pricing kernels before/after adding the strategy-**F** to the return space to construct the SDFs[2].

To this end, I use the two-step procedure of Fama and MacBeth (1973) to evaluate the ability of each SDF in pricing the cross-section of asset returns. In the first step, I obtain the factor loadings for the test assets from a time-series regression of the excess asset returns on the pricing kernel:

$$\mathbf{R}_t^e = a + BM_t + \varepsilon_t$$

[2]Note that the cross-sectional asset pricing ability of the I-SDF (without the strategy-**F**) is comprehensively documented in Ghosh, Julliard, and Taylor (2016a)

In the second step, the risk premium is estimated from a cross-sectional regression of the average excess returns, $E_T(\mathbf{R}_t^e) \in \mathbb{R}^N$, on the estimated loadings B of the first stage regression:

$$\mu = X\lambda + \alpha = z\iota + B\gamma + \alpha, \quad X \equiv \begin{bmatrix} \iota & B \end{bmatrix}, \quad \lambda' \equiv \begin{bmatrix} z & \gamma \end{bmatrix} \tag{2.10}$$

where ι denotes a conformable vector of ones, γ denotes a regression slope, and α is an $N \times 1$ vector of pricing errors. z is a scalar constant which should be zero once the zero beta rate matches the average risk-free rate.

Following the critique of Lewellen, Nagel, and Shanken (2010), I make use of a mixed cross-section of test assets to avoid a strong factor structure in test asset returns. In particular, I consider the 25 size and book-to-market-equity sorted portfolios, the 10 momentum-sorted portfolios, and the 10 portfolios formed on long-term reversal, in addition to excess returns on zero-coupon bonds with maturities of 2, 5, 7 and 9 years.

I report several measures of performance for the cross-sectional regressions based on the suggestions of Lewellen, Nagel, and Shanken (2010), henceforth LNS. I first present the standard ordinary least squares (OLS) cross-sectional adjusted R^2 (henceforth $\bar{R}^2_{OLS}$) with its associated confidence interval. I also report the generalized least squares (GLS) cross-sectional adjusted R^2 (henceforth $\bar{R}^2_{GLS}$). This measure represents a factor's proximity to the minimum-variance boundary and, hence, can be considered a more relevant measure than OLS R^2 to assess the factor's ability to explain the risk-return opportunities (see Lewellen, Nagel, and Shanken (2010) for details). The T^2 statistics in Shanken (1985) are also reported. The statistics are a weighted sum of squared pricing errors and defined as $\hat{\alpha}' S_\alpha^+ \hat{\alpha}$ where S_α^+ is the pseudoinverse of the consistent estimates of the covariance matrix of pricing errors - henceforth denoted as Σ_α. The pricing errors, $\hat{\alpha}$, has asymptotic variance $\Sigma_\alpha = \psi y \Sigma y / T$ where $y \equiv I_N - X(X'X)^{-1} X'$ and Σ is the covariance matrix of the error term, ε_t, in the first stage regression. Here, $\psi = \left(1 + \gamma' \Sigma_f^{-1} \gamma\right)$ is the multiplicative correction in Shanken (1992) to reflect estimation error in B where Σ_f is the covariance matrix of risk factors, hence the variance of the pricing kernel in the specific application here. The statistics asymptotically follow χ^2 with degrees of freedom $N - K - 1$. Finally, the quadratic, $q \equiv \hat{\alpha}' (y \Sigma y)^+ \hat{\alpha}$, is presented. This measure shows how far factor mimicking portfolios are from the mean-variance frontier - it assesses the difference between the maximum generalized

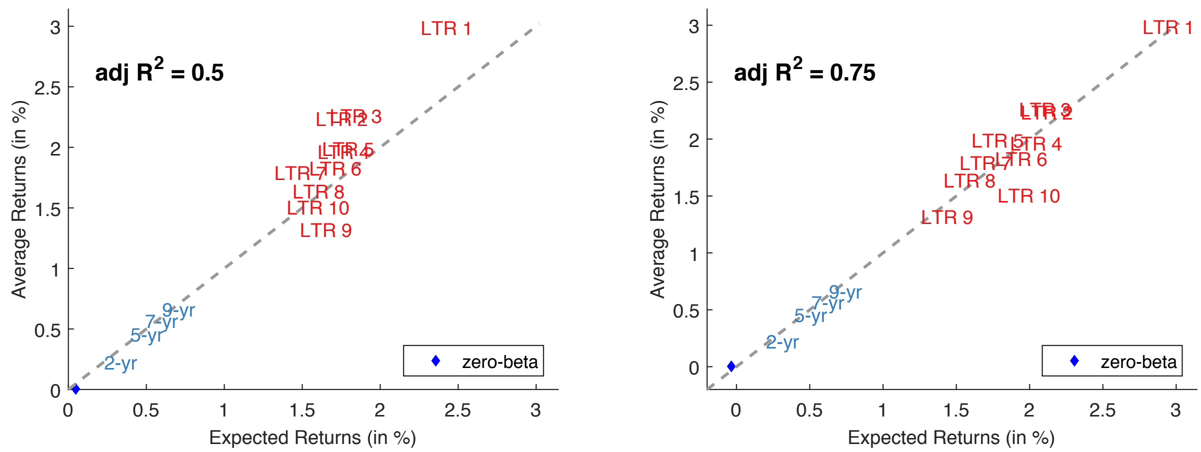

(a) Without the Strategy-**F** (b) With the Strategy-**F**

Figure 2.2. The figures plot the average and expected returns on 10 long-term reversal portfolios and bonds with maturities of 2, 5, 7 and 9 years. Expected returns and adjusted R^2s are from the GLS of the cross-sectional equation (2.10), and the factor for Panel B (A) is the I-SDF (not) using the strategy-**F** for its construction.

squared Sharpe ratio on any portfolio and that attainable from a portfolio constructed by the test assets that is maximally correlated with the SDF.

Panel A in Table 2.2 reports the cross-sectional pricing results when the test assets are composed of the 25 Fama-French portfolios, the 10 momentum portfolios, and the 30 industry portfolios, in addition to the zero-coupon bonds with maturities of 2, 5, 7 and 9 years. The estimated price of risk in each SDF is strongly statistically significant with an absolute value of the t-statistic (reported in parentheses) greater than 3, and the estimated zero-beta rates are not statistically different from the observed average three-month Treasury bill rate. However, with the strategy-**F** in their Euler equations, the price of risk in each HJ-SDF and I-SDF is more than doubled. Put differently, the price of risk is significantly underestimated when we do not take into account the dynamic property of an SDF. Across all the considered measures of cross-sectional asset pricing performance, the filtered SDFs with the strategy-**F** outperform the corresponding SDFs without the strategy-**F**. In particular, the T^2 statistic shows that the models with the strategy-**F** are not rejected at conventional significance levels, whereas the corresponding SDFs without the strategy-**F** are rejected at the 5% significance level (p-values are reported in parentheses).

Table 2.2
Cross-Sectional Regressions

Panel A reports the cross-sectional pricing ability of the HJ-SDFs and I-SDFs where the cross-sectional test assets are the 25 size and book-to-market-equity sorted portfolios (25 FF), the 10 momentum-sorted portfolios (10 MOM), and the 30 industry portfolios (30 IND), in addition to excess returns on zero-coupon bonds with maturities of 2, 5, 7 and 9 years. Panel B also reports the cross-sectional pricing ability but with 10 long-term reversal portfolios and 2, 5, 7 and 9-year zero-coupon bonds as test assets.

Panel A: 25 FF, 10 MOM, 30 IND Portfolios and Bonds

Kernel	Strategy-**F**	const.	λ	$\bar{R}^2_{OLS}$	$\bar{R}^2_{GLS}$	T^2	q
HJ-SDF	No	0.00	-0.65	0.84	0.43	104.61	0.83
		(0.27)	(-4.70)			(0.00)	
	Yes	0.00	-0.97	0.86	0.53	74.12	0.69
		(0.27)	(-4.26)			(0.26)	
I-SDF	No	0.00	-0.85	0.78	0.41	98.62	0.85
		(0.40)	(-4.48)			(0.01)	
	Yes	0.00	-1.60	0.85	0.54	55.41	0.67
		(0.27)	(-3.70)			(0.84)	

Panel B: 10 Long-term-reversal Portfolios and Bonds

Kernel	Strategy-**F**	const.	λ	$\bar{R}^2_{OLS}$	$\bar{R}^2_{GLS}$	T^2	q
HJ-SDF	No	0.00	-0.70	0.95	0.53	7.23	0.06
		(0.02)	(-2.02)			(0.84)	
	Yes	0.00	-1.02	0.97	0.63	4.87	0.05
		(0.10)	(-1.88)			(0.96)	
I-SDF	No	-0.00	-1.00	0.94	0.50	6.44	0.06
		(-0.12)	(-1.82)			(0.89)	
	Yes	-0.00	-1.85	0.96	0.75	2.21	0.03
		(-0.02)	(-1.53)			(1.00)	

The flexibility of non-parametric approaches tends to accompany over-fitting problems. Panel B in Table 2.2 investigates this issue by deploying fresh cross-sectional test assets that are not utilized in the construction of the SDF. Note that, for the previous exercise, the set of asset returns for the construction of the SDF is a subset of the test assets for the cross-sectional tests (although it is a very small subset). For Panel B, I use the same SDF as in Panel A, but instead run cross-sectional regressions both with the 10 portfolios formed on the basis of long-term reversal and with the 2, 5, 7 and 9-year bonds (note that 3 and 8-year bond returns are employed to filter the SDF). All the SDFs do reasonably good jobs with respect to all the measures of performance. However, the SDFs with the strategy-**F** noticeably outperform the SDFs without the strategy-**F** in their Euler equations. That is clearly demonstrated in Figure 2.2, where the average returns and expected returns fitted by the GLS of the cross-sectional equation (2.10) are plotted.

2.6 Evaluation of Asset Pricing Models

Both the variance of the Hansen-Jagannathan SDF and the Kullback-Leibler Information Criterion (KLIC) of the information SDF in Table 2.1 are significantly enlarged after incorporating the strategy-**F** into their Euler equations. This is in line with the observation in Backus, Chernov, and Zin (2014) that a tension exists for standard structural asset pricing models to explain the magnitude of risk premia and the dynamics of a pricing kernel (summarized as the shape of the yield curve in their approach). The indicative increases in the variance of the Hansen-Jagannathan SDF and the KLIC of the information SDF demonstrate that the tension is pervasive even in the non-parametrically estimated SDFs. Importantly, the strategy-**F** (which contains the information on the dynamics of the SDF) helps to better identify the underlying SDF, in the sense that the average asset returns are better explained by their covariance with the SDF embodying the strategy-**F** than by the conventionally estimated SDFs in the absence of the strategy-**F**. Hence, adding the return, $\mathbf{F}_t^e$, to the return space (for the construction of the SDF) should be considered as an important asset pricing constraint, rather than an artificial restriction to tighten the set of admissible SDFs. In other words, this section assesses the extent to which the strategy-**F** in the Euler equation (i.e. the information on the dynamics of the SDF) affects various bounds for the SDF.

To this end, I initially start with the definition of the famous Hansen-Jagannathan bound (henceforth, HJ bound), in conjunction with the entropy bounds of SDFs in Ghosh, Julliard, and Taylor (2016b) - henceforth GJT. It is well known that the HJ bound denotes the variance of the HJ-SDF in equation (2.6). Similarly, the Q-bound in GJT is defined as the KLIC of the I-SDF in equation (2.7) - hence, the KLICs in Table 2.1 are the Q bound of I-SDFs with/without the strategy-**F**. GJT show that the HJ bound can be considered a second-order approximation of the Q bound, which implies that the Q bound presents a more stringent hurdle for the equilibrium models to be evaluated. GJT also propose two additional entropy bounds based on the observation that a large class of equilibrium asset pricing models can be decomposed into observable and unobservable components. More precisely, most of the consumption-based asset pricing models can be factorized into two components, i.e. $M_t = m(\theta, t) \times \psi_t$ where $m(\theta, t)$ is a known function of observable variables (mostly consumption growth) with the parameter vector θ, and an unobservable component ψ_t. GJT define the following two entropy bounds for an SDF. M-bound is the KLIC between $m(\theta, t) \times \psi_t$ and the physical probability,

$$D\left(\frac{M_t}{\bar{M}}||\mathbb{P}\right) \geq D\left(\frac{m(\theta, t)\,\psi_t^*}{\overline{m(\theta, t)\,\psi_t^*}}||\mathbb{P}\right),$$

where ψ_t^* solves equation (2.7). Similarly, their Ψ-bound sets the lower bound for the relative entropy of ψ_t with respect to $\mathbb{P}$,

$$D\left(\frac{\psi_t}{\bar{\psi}_t}||\mathbb{P}\right) \geq D\left(\frac{\psi_t^*}{\bar{\psi}_t^*}||\mathbb{P}\right),$$

where ψ_t^* solves equation (2.7).

For the Consumption-CAPM (henceforth, C-CAPM) of Breeden (1979), Lucas (1978) and Rubinstein (1976), the pricing kernel consists of observable components only. The utility function is one with constant relative risk aversion and hence the SDF is specified as

$$M_{t+1} = \delta\left(C_{t+1}/C_t\right)^{-\gamma}$$

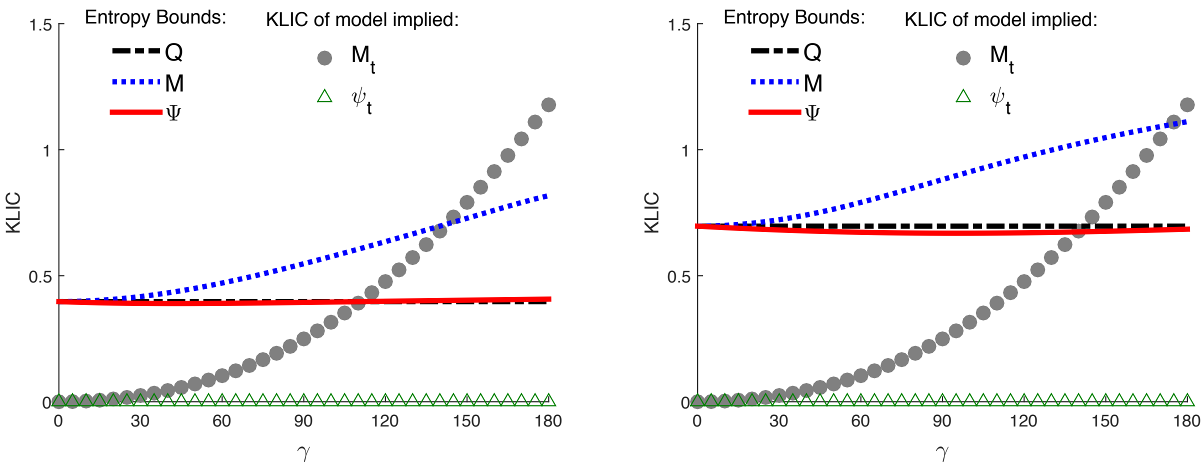

(a) Entropy Bounds without the Strategy-**F**

(b) Entropy Bounds with the strategy-**F**

Figure 2.3. The figures plot the entropy bounds of C-CAPM with/without the strategy-**F**

where γ is the coefficient of relative aversion, δ denotes the subjective time discount factor, and C_{t+1}/C_t is the real per capita aggregate consumption growth. The SDF itself is observable given the value of the parameters, and its unobservable component is set as $\psi_t = 1$ for all t. Hence, the entropy bounds for the C-CAPM provide useful benchmarks for more elaborate asset pricing models. More precisely, most recent developments in consumption-based asset pricing models are about refining the constant ψ_t in the C-CAPM so that a new unobserved component is time-varying and elucidates observed asset prices and macroeconomic quantities - see for example Campbell and Cochrane (1999), Bansal and Yaron (2004), Menzly, Santos, and Veronesi (2004) and Piazzesi, Schneider, and Tuzel (2007), among many others. Hence, the Ψ-bound of the C-CAPM designates what remains to be explained by any SDFs with the form of $M_t = m(\theta, t) \times \psi_t$ after taking into account the (observable) consumption growth.

Figure 2.3 summarizes the Q, M and Ψ bounds for the C-CAPM: the same set of asset returns in the previous two sections is used to calculate the entropy and HJ bounds, especially in conjunction with the strategy **F** for Panel (b). Because the missing component of the SDF, ψ_t, is a constant, its relative entropy is zero regardless of the value of the risk aversion coefficient γ (see the green line with triangles). Hence, the model-implied SDF is equivalent to the observable component $m(\theta, t)$ i.e. the consumption growth raised to power in $-\gamma$, $\delta(C_{t+1}/C_t)^{-\gamma}$. The black curve with circles shows the KLICs of the model implied

SDF along with a different degree of risk aversion. In addition, the blue dotted curve and the red solid curve represent, respectively, the relative entropy of the filtered SDF and its missing component as a function of the risk aversion coefficient.

Panel (a) in the figure shows the entropy bounds and the KLICs of the C-CAMP implied SDF without incorporating the strategy-$\mathbf{F}$ into the Euler equation. The model satisfies the Q bound for $\gamma \geq 111$ which corresponds to the intersection of the horizontal dash-dot black line and the black curve with circles. The minimum value of γ required to satisfy the M bound is 142, as shown by the intersection of the dotted blue line and the black curve with circles. The Ψ bound is not satisfied for any value of γ since the model's missing component, ψ_t, is constant and is not affected by the value of γ. Finally, the HJ bound is met for $\gamma \geq 89$. After incorporating the strategy-$\mathbf{F}$ into the Euler equation as in equation (2.9), then the Q, M and HJ increase up to 142, 174 and 101 respectively. The increases are remarkable, especially given the size of test assets - the return space in the Euler equation, i.e. $\{\mathbf{R}^e_t | t = 1, ..., T\}$ in equation (2.7), consists of 28 assets (6 size and book-to-market sorted portfolios, 10 industry portfolios, and 10 momentum portfolios, in addition to two bonds). By adding one single return $\mathbf{F}^e_t$ (the excess return on the strategy-$\mathbf{F}$) to the returns space, remarkably tighter bounds on SDFs can be achieved. Put differently, the single restriction on the SDF, $E^{\mathbb{P}}\left(M_t \mathbf{F}^e_t\right) = 0$, significantly refines the admissible space of M_t, and the validity of the restriction is even supported by the cross-sectional implications as in Section 2.5.

2.7 Conclusion

This paper utilize a particular zero-cost strategy to better identify the underlying SDF. The strategy tracks the evolution of the Radon-Nikodym derivative between two pricing measures with alternative investment horizons, hence reveal the characteristic of the dynamics of the SDF.

Incorporating the strategy into the Euler equation significantly enhances the ability of the non-parametrically filtered SDF to explain cross-sectional variation of expected asset returns.

Furthermore, the strategy remarkably tightens various lower bounds for the stochastic discount factor, hence setting a more stringent hurdle for equilibrium asset pricing models. Because the strategy-**F** is particularly informative about the term structure of the SDF, the approach that I posit here is that of unifying the information contents of (i) measure of dispersion entropy bound and (ii) horizon dependence in Backus, Chernov, and Zin (2014).

Chapter 3

Bond Variance Risk Premiums

Markets for volatility derivatives have grown a lot over the last decade. Nowadays, investors have both exchange-traded and over-the-counter instruments available to hedge and trade volatility in equity markets. While a plethora of research has focused on understanding and pricing equity variance risk, the same risk is much less well understood in fixed income markets.

In this paper, we propose the generalized Treasury variance swap (GTVS) which offers a pure exposure to fixed income variance.[1] We theoretically derive the fair value of the contract and empirically document significant returns to variance trading in Treasury markets that are comparable to those earned in the equity variance market. We quantitatively compare our GTVS strategy to alternative volatility trading strategies based on fixed income options that may not leave the investor with a clean volatility exposure. Calculating the variance swap rate for various horizons, we obtain a term structure of Treasury implied variances and ex post bond variance risk premiums. Finally, we show that the term structure of Treasury implied variances is significantly related to economic activity and stress indicators for financial markets.

The main contribution of this paper is twofold. First, we study the theoretical properties of variance swaps in Treasury markets. Different from standard variance contracts, our strategy is model-free and allows for stochastic interest rates. Variance swaps consist of two legs: (i) a realized leg and (ii) the fair strike. The latter is defined as the strike that makes the net present value of the swap equal to zero at initiation. Building on the pioneering

[1] Data is available from the authors' webpages.

work of Neuberger (1994) and Demeterfi, Derman, Kamal, and Zou (1999), we show that the fair strike can easily be approximated using a portfolio of puts and calls. The way we define the realized leg is crucial. For the generalized Treasury variance swap we use a generalized measure of realized variance which allows for perfect replication of the contract even in the presence of jumps (see Neuberger (2012) for an application in the equity market). This property is particularly useful in face of the recent extreme events as standard variance swaps (the log Treasury variance swap or LTVS) rely on squared log returns and are therefore exposed to cubed returns, resulting in extremely inefficient hedges for investors.

Second, using a large panel data set of daily option prices on Treasury futures with different tenors, we study the payoffs of the GTVS and the associated ex post variance risk premiums, defined as the difference between the realized variance and the fair strike or variance swap rate. Our main findings can be summarized as follows: Consistent with the literature in the equity market we find that the variance risk premiums are negative and economically large. The average excess returns for a strategy that shorts the variance swap with a one-month maturity, independent of the tenor of the underlying, is around 20% per month (21.2% for the 30y Treasury futures (t-statistic of 8.98), 27.6% for the 10 year Treasury futures (t-statistic of 12.71) and 18.7% for the 5y Treasury futures (t-statistic of 6.61), respectively) and the associated annual Sharpe ratio is just below two.[2] Since our theory is based on (European) options on forwards, we need to convert the (American) options on futures to European options on forwards. We show that the adjustment is overall very small, especially for options with maturities of less than one year.

Traditionally, the most common strategy to exploit variance is to invest in a delta-hedged at-the-money straddle. A drawback of straddles compared to variance swaps is that the sensitivity of the trading strategy with respect to volatility or variance is non-linear. Moreover, the sensitivity to other factors is typically non-zero (volatility of volatility effect). We find that trading a one-month delta-hedged at-the-money straddle generates significantly smaller rewards compared to a position in a one-month variance swap over our sample period: The average return on the 10y Treasury futures straddle is around -7%

[2]The excess return is calculated as a simple return, meaning it is the variance risk premium normalized by the fair strike price or the implied variance.

(with an associated t-statistic of 4.91) and the annualized Sharpe ratio is around 1.05, which is about half in size compared to the GTVS.

We then study whether the alpha of these strategies can be explained by common factors found to explain variance trading strategies. Not surprisingly, we find that the market return explains very little of the excess returns on variance swaps. However, even if we include a battery of additional factors such as size, book-to-market, or bond market liquidity, the time-series variation of these risk premiums remains largely unexplained. Moreover, the alpha of the variance swap strategy continues to be highly significant. For example, for the 30y Treasury futures, we find that the alpha is 26% per month, even when including all controls.

As a by-product of our analysis, we construct the unconditional term structure of implied variances and variance risk premiums. The unconditional variance risk premiums are obtained as the unconditional average of the difference between the realized variance under the physical measure and the expected variance under the risk-neutral measure. We find that, on average and for all tenors, the term structure of implied variances is downward sloping. The slope, defined as the difference between the implied variance of a one-year and a one-month option is strongly pro-cyclical. During crisis periods, the slope becomes extremely negative. Moreover, the slope has strong predictive power for future economic growth and proxies of economic stress, especially at short horizons: A steeper slope predicts higher growth or lower stress up to eight months ahead. When we add the slope of the term structure of Treasury yields, which itself is considered a good predictor of future growth, we find that the predictive power of the implied variance slope remains unchanged. Moreover, at short horizons, the slope of the yield curve has no power at all.

The term structure of variance risk premiums is also downward sloping in absolute terms. Expressed in monthly squared percent, the variance risk premiums for the 30y Treasury futures are ranging from -4.9 for a ten-day horizon to essentially zero for the one-year horizon. For 10y Treasury futures, the short- and longer-term variance risk premiums are -3.1 and -0.4, while for the 5y Treasury futures they are -1.1 and -0.4, respectively. Note that unlike the excess returns on the generalized Treasury variance swap, the variance risk premiums are not normalized by the implied variance and they are, thus, distinctly increasing in the variance level (which is increasing in the underlying tenors).

We run different robustness checks to challenge our findings. First, one might be worried about statistical significance as option returns are known to be non Gaussian. We therefore use a studentized bootstrap to obtain confidence intervals on the means, alphas, and Sharpe ratios of our trading strategies. We find all performance measures to be statistically different from zero.

Second, we also study the impact of bid and ask spreads on the profitability of the variance swap strategy and find that while average returns decrease by around 3%, shorting variance is still a very attractive strategy: Annualized Sharpe ratios remain above one and the alpha is highly significant.

Third, margin requirements limit the notional amount of capital that can be invested in trading strategies. Moreover, large losses in a position can force investors out of a trade, potentially at the worst possible time. Thus, in order to realistically assess the profitability of our proposed variance swap trading strategy we take into account the impact of margins on the realized returns. Earlier literature suggests that margin requirements can have an economically significant effect when investors do not have access to unlimited capital when the market is in a downturn. We confirm that there is a difference between average returns for an unrestricted and a margined variance swap strategy but we find that average returns and Sharpe ratios remain significant and economically large.

Fourth, we ask to what extent it matters how the variance swap is constructed. As mentioned, a perfect replication of the variance swap that is robust to jumps in the underlying and the choice of the re-balancing frequency requires a particular definition to measure the realized variance. This definition differs from the standard way to calculate realized variance, namely to use squared log returns. On average, we find that the payoffs of the GTVS and those from a contract based on realized variance measured using log returns (LTVS) are very similar. This stems from the fact that positive and negative jumps cancel each other out. The devil is in the details, however. We show that whenever there are jumps in the underlying, regardless of whether they are positive or negative, the returns to the two trading strategies are distinctly different. Negative (positive) jumps render the payoff to the

LTVS larger (smaller) compared to the GTVS. More importantly, while the payoff to the GTVS can still be perfectly hedged, this is no longer the case for the LTVS. This is in line with the findings of Broadie and Jain (2008) who study variance swaps in the presence of jumps in the equity index market and document that negative jumps have the most severe impact on replication errors.

Finally, one might wonder whether it matters for the profitability of our trading strategies whether we include crisis periods or not. Hence, a natural additional robustness check is to study variance swap returns for different time periods. We confirm that trading strategies that go short variance during periods that include the October 1987 crash or the height of the credit crisis in August 2008 are still very profitable over time. Overall, we find almost no quantitatively relevant differences across various sub-periods.

Related Literature: Our paper draws from the large literature on variance trading in equity markets starting with the work of Dupire (1994) and Neuberger (1994). For example, Carr and Wu (2009) use portfolios of options to approximate the value of the variance swap rate for different stock indices and individual stocks. They then compare the synthetic variance swap rates to the ex post realized variance to determine the size of the variance risk premium. Wu (2010) estimates variance risk dynamics by combining the information in realized variance estimators from high frequency returns and the VIX. Egloff, Leippold, and Wu (2010) directly use variance swap quotes and study the term structure of variance swap rates. Motivated by the recent financial crisis, much attention has been paid to the pricing and hedging of equity variance swaps in the presence of jumps. Both Schneider (2015) and Schneider and Trojani (2015) study tradeable properties of volatility risk, where the latter focus on higher-order risk premiums attached to time-varying disaster risk.

Other papers investigate the term structure of variance risk premiums and prices in the equity index market. Dew-Becker, Giglio, Le, and Rodriguez (2016) estimate an affine term structure model using variance swap data and find that realized volatility is the only priced risk factor which implies a term structure that is steeply negative at the short-end but flat beyond a one-month maturity. They conclude that these stylized facts are hard to reconcile within standard asset pricing models. Similarly, Andries, Eisenbach, Schmalz, and Wang (2015) study the term structure of variance risk premiums and find that a model where investors feature horizon-dependent risk aversion matches the data well. Our paper is different from this strand of the literature, as our approach is completely model-free and we do not take a stand on the microfoundations.

We are not the first to explore variance contracts in fixed income markets. Trolle (2009) estimates variance risk premiums in two ways: First, he estimates a dynamic term structure model that allows for unspanned stochastic volatility and, second, he corroborates his findings using a model-independent approach similar to ours. Both approaches lead him to the conclusion that the market price of variance risk is highly negative. While his focus is on a dynamic portfolio choice problem which includes interest rate derivatives, his approach also differs from ours as he derives the model-free variance risk premiums under different assumptions. Furthermore, he does not study the term structure of variance risk premiums. Trolle and Schwartz (2014) study variance and skewness across different swap maturities and option tenors in the swaptions market. They find that both variance and skewness risk premiums are negative and highly time-varying. The authors then propose a dynamic term structure model that fits the dynamics of these risk premiums. We see our work complimentary to theirs as our focus is on documenting empirical facts about a variance trading strategy rather than asking what model is best suited to capture the dynamic behavior of conditional swap rate moments. Recently, Cieslak and Povala (2016) study yield volatility risk and suggest that investors willingness to pay large premiums to hedge volatility can be linked to uncertainty about the future path of monetary policy.

In contemporaneous theoretical work, Mele and Obayashi (2013) explore variance contracts on Treasury futures similar to ours. There are important differences, however. While their fair strike is constructed in the same way than ours, their definition of the realized leg is different as it is constructed using log returns. In our paper, we show that using realized variance based on log returns leaves an exposure to cubed returns. This has important consequences during periods where the underlying exhibits jumps. Moreover, we also conduct an empirical analysis of the proposed contract while they are mainly interested in the theoretical derivation of the option-implied leg. Finally, Merener (2012) constructs a variance strategy on forward swap rates and studies dynamic hedging strategies. He assumes that the forward curve is flat which implies that the no-arbitrage condition is violated. In our setting, we study a variance contract where the underlying is a traded asset and derive explicit solutions for the hedge positions. Our assumptions are very general and, in particular, we only assume that no-arbitrage holds.

Our paper is also related to Aït-Sahalia, Karaman, and Mancini (2015) who estimate a two-factor affine stochastic volatility model to study the term structure of variance swaps in the equity index market. They show that the risk premiums contain a large jump risk component, especially at short horizons. Filipović, Gourier, and Mancini (2015) propose a quadratic term structure model for equity variance swaps. Using data on over-the-counter variance swaps, they also find a downward sloping term structure of variance swap payoffs.

The findings in this paper are also related to Duarte, Longstaff, and Yu (2007) who study risk and return for different fixed income arbitrage strategies featuring, among others, a volatility trading strategy through delta-hedged caps. Depending on the cap maturity, the (annualized) Sharpe ratios can be quite attractive, reaching 0.82 for a four-year maturity cap. The difference between their study and ours is that their results depend on a particular model. Hedge ratios to calculate the delta of caps are based on Black (1976). Our results are model-independent, moreover, shorting delta-hedged caps leaves the investor with Gamma exposure, similar to the straddle strategies that we consider as an alternative to the variance swap.

The rest of the paper is organized as follows. Section 3.1 provides the expressions to price variance swaps in Treasury futures markets and introduces the generalized Treasury variance swap. Section 3.2 describes our data set, explains the calculation of the variance risk premiums, and documents the term structures of implied volatilities and variance risk premiums in the fixed income market. Section 3.3 outlines the construction of the various trading strategies and presents the results of our empirical study, and Section 3.4 concludes. The Appendix contains some proofs and derivations; a detailed description of data filters and additional robustness checks are deferred to an Online Appendix.

3.1 Theory

This section theoretically derives the payoff of a variance swap in the Treasury bond market. While the contract we propose is robust to jumps in the underlying, we start with a standard contract that assumes a continuous process for the underlying before relaxing this assumption to present our main result. The model-free implementation requires that we use forward contracts instead of the futures that we have available in the data. Thus, in the empirical implementation, we also show how to convert American options on futures (as observed empirically) to European options on forwards (that are used in the theoretical derivation). Quantitatively, we find the differences between the two types of options to be negligible, especially for options with short maturities.

Variance swaps consist of two different legs: The floating leg (realized variance) and the fixed leg (expected variance). The difference between the two is then the variance risk premium (see, e.g., Carr and Wu (2009)). We fix the current time $t = 0$ and study contracts which pay at some future date T. We denote by $F_{t,T}$ the price of a forward contracted at time t with maturity T on the underlying X_T.

3.1.1 Log Treasury Variance Swap

To start, we assume that $F_{t,T}$ follows a continuous process. By no-arbitrage, this implies that the dynamics of $F_{t,T}$ are

$$\frac{dF_{t,T}}{F_{t,T}} = \sigma_t dW^{\mathbb{Q}_T},$$

where $W^{\mathbb{Q}_T}$ is a standard Brownian motion under the T-forward measure $\mathbb{Q}_T$ and σ_t is the instantaneous volatility.

A standard variance swap exchanges the realized variance defined as

$$\mathrm{RV}_{t,T}^{\log} := \left(\log \frac{F_{t+1,T}}{F_{t,T}}\right)^2 + \left(\log \frac{F_{t+2,T}}{F_{t+1,T}}\right)^2 + \ldots + \left(\log \frac{F_{T,T}}{F_{T-1,T}}\right)^2$$

with some fair strike $\tilde{F}_{t,T}$ at maturity T. We assume that the realized variance is constructed using some sampling partition $\mathcal{T} = [t_0, t_1, \ldots, t_n]$ with trading dates $0 = t_0 < t_1 < \ldots < t_n = T$. For a variance swap, the fair strike $\tilde{F}_{t,T}$ is chosen such that at initiation of the contract at time t no money is exchanged. Dupire (1994) and Neuberger (1994) define the variance swap rate $\tilde{F}_{t,T}$ in terms of the expected payoff under the risk-neutral measure $\mathbb{Q}$ from a so-called Log contract. To account for stochastic interest rates, we make use of the $\mathbb{Q}_T$ measure

$$\tilde{F}_{t,T} := -2\mathbb{E}_t^{\mathbb{Q}_T}\left[\log\frac{F_{T,T}}{F_{t,T}}\right]. \tag{3.1}$$

We now want to derive an expression for Equation (3.1). Let us start with the fundamental theorem of asset pricing, which implies that for any traded asset X_t

$$\frac{X_t}{p(t,T)} = \mathbb{E}_t^{\mathbb{Q}_T}\left[X_T\right], \tag{3.2}$$

where $p(t,T)$ is the price of a zero-coupon bond. Relation (3.2) holds in general and in particular under stochastic interest rates. Consider now the payoff $\log F_{T,T}$. Using the results in Carr and Madan (1998), we can write

$$\log F_{T,T} = \log F_{t,T} + \frac{F_{T,T}-F_{t,T}}{F_{t,T}} - \int_0^{F_{t,T}}\frac{(K-F_{T,T})^+}{K^2}dK - \int_{F_{t,T}}^{\infty}\frac{(F_{T,T}-K)^+}{K^2}dK.$$

Re-arranging yields that

$$-2(\log F_{T,T}-\log F_{t,T}) = 2\left(-\frac{F_{T,T}-F_{t,T}}{F_{t,T}} + \int_0^{F_{t,T}}\frac{(K-F_{T,T})^+}{K^2}dK + \int_{F_{t,T}}^{\infty}\frac{(F_{T,T}-K)^+}{K^2}dK\right).$$

The forward $F_{t,T}$ is a $\mathbb{Q}_T$-martingale.[3] By taking $\mathbb{Q}_T$ expectations on both sides, we get

$$\begin{aligned}-2\mathbb{E}_t^{\mathbb{Q}_T}\left[\log F_{T,T}-\log F_{t,T}\right] &= \mathbb{E}_t^{\mathbb{Q}_T}\left[\int_0^{F_{t,T}}\frac{(K-F_{T,T})^+}{K^2}dK + \int_{F_{t,T}}^{\infty}\frac{(F_{T,T}-K)^+}{K^2}dK\right]\\ &= \frac{2}{p_{t,T}}\left(\int_0^{F_{t,T}}\frac{P_{t,T}(K)}{K^2}dK + \int_{F_{t,T}}^{\infty}\frac{C_{t,T}(K)}{K^2}dK\right),\end{aligned} \tag{3.3}$$

[3] It is easy to see that

$$\mathbb{E}_t^{\mathbb{Q}_T}\left[F_{T,T}\right] = \mathbb{E}_t^{\mathbb{Q}_T}\left[S_T\right] = \mathbb{E}_t^{\mathbb{Q}}\left[\frac{\exp\left(-\int_t^T r_s ds\right)}{p_{t,T}}S_T\right] = F_{t,T}.$$

where $P_{t,T}(K)$ and $C_{t,T}(K)$ are puts and calls, since by Equation (3.2)

$$\frac{P_{t,T}(K)}{p(t,T)} = \mathbb{E}_t^{\mathbb{Q}_T}\left[(K - F_{T,T})^+\right]$$
$$\frac{C_{t,T}(K)}{p(t,T)} = \mathbb{E}_t^{\mathbb{Q}_T}\left[(F_{T,T} - K)^+\right].$$

Hence, it follows that the Log contract can be written as a portfolio of puts and calls with the same strike K and the same maturity T. To see that Equation (3.3) indeed represents expected variance, note that by applying Itô's Lemma, we get that

$$\log F_{T,T} - \log F_{t,T} = -\frac{1}{2}\int_t^T \sigma_u^2 du + \int_t^T \sigma_u dW_u^{\mathbb{Q}_T}.$$

Note that our objective is to price a contract on forward volatility which is the same as the volatility on futures. Forwards are martingales under the $\mathbb{Q}_T$-measure while futures are martingales under the risk-neutral measure. By Girsanov's theorem, the volatilities are the same while their drifts are not. We summarize our findings in a first Proposition.

Proposition 2. *Assume that $F_{t,T}$ is continuous. Then, the payoff $\int_t^T \sigma_u^2 du$ can be perfectly replicated by a static position in*

$$\tilde{F}_{t,T} = \frac{2}{p_{t,T}}\left(\int_0^{F_{t,T}} \frac{P_{t,T}(K)}{K^2}dK + \int_{F_{t,T}}^{\infty} \frac{C_{t,T}(K)}{K^2}dK\right),$$

and a dynamic position in the underlying which at any time $s \in [t,T]$ holds $2\left(\frac{1}{F_{s,T}} - \frac{1}{F_{t,T}}\right)$. Hence,

$$\tilde{F}_{t,T} = \mathbb{E}_t^{\mathbb{Q}_T}\left[\int_t^T \sigma_u^2 du\right].$$

Proof: See Appendix 8.

Equation (3.3) resembles the definition of the VIX but instead of deriving everything under the risk-neutral measure, we derive it under the T-forward measure. This is important as it allows for stochastic interest rates, which is needed since we want to study contracts in the fixed income market. The VIX is in general referred to as being model-free, because we have not made any assumption on $F_{t,T}$ other than it being an Itô process. However, once we deviate from this assumption, i.e., once we allow for jumps, VIX^2 is no longer the fair strike of a variance swap. For example, Carr and Wu (2009) note that

$$\mathbb{E}_t^{\mathbb{Q}_T}\left[\int_t^T \sigma_u^2 du\right] = \mathrm{VIX}_t^2 + \text{error from jumps}.$$

Negative (positive) jumps induce an upward (downward) bias in VIX^2. The sensitivity of the standard variance swap presented in Proposition 2 has been extensively documented in the literature, see for example Broadie and Jain (2008).

3.1.2 Generalized Treasury Variance Swap

We now proceed to relax the assumption that the underlying follows a continuous process. Martin (2013) studies so called simple variance swaps which are robust to jumps. He does this by altering the fair strike, $\tilde{F}_{t,T}$. In the following, we are deviating from this approach by changing the realized variance leg of the contract instead. The idea is that rather than focussing on the unobservable quantity $\tilde{F}_{t,T}$, we concentrate on the observed realized variance. This closely follows Neuberger (2012) and Bondarenko (2014) who study generalized variance swaps in the equity market. We extend their approach by allowing for stochastic interest rates.

We define the generalized Treasury variance swap (GTVS) as an agreement to exchange

$$\widetilde{RV}_{t,T} = 2\sum_{i=1}^{T-t}\left(\frac{F_{t+i,T}}{F_{t+i-1,T}} - 1 - \log\frac{F_{t+i,T}}{F_{t+i-1,T}}\right), \tag{3.4}$$

with the fair strike $\tilde{F}_{t,T}$. At first sight, $\widetilde{RV}_{t,T}$ may not look like a variance measure but it turns out to be the same as realized variance computed from simple returns minus cubed simple returns.[4] This new measure of realized variance has the useful property that it allows for a perfect replication of the variance contract for every price path and every partition. We summarize our findings in the following Proposition.

Proposition 3. *For any process $F_{t,T}$, the payoff $\widetilde{RV}_{t,T}$ can be perfectly replicated by a static position in*

$$\tilde{F}_{t,T} = \frac{2}{p_{t,T}}\left(\int_0^{F_{t,T}} \frac{P_{t,T}(K)}{K^2}dK + \int_{F_{t,T}}^{\infty}\frac{C_{t,T}(K)}{K^2}dK\right), \tag{3.5}$$

[4] A formal derivation can be found in Appendix 9.

and a dynamic position in the underlying, which at any time $s \in \mathcal{T}$ holds $2\left(\frac{1}{F_{s,T}} - \frac{1}{F_{t,T}}\right)$. *Hence,*

$$\tilde{F}_{t,T} = \mathbb{E}_t^{\mathbb{Q}_T}\left[\widetilde{RV}_{t,T}\right].$$

Proof: See Appendix 8.

Hence, it follows that the replicating strategy for the realized variance consists of two parts: First, a path-independent payoff from options and, second, a dynamic strategy in the underlying. Note that we have not made any assumption about $F_{t,T}$ or about the frequency with which we re-balance the portfolio. Proposition 3 looks almost identical to Proposition 2, and indeed, if $F_{t,T}$ is continuous, both $\widetilde{RV}_{t,T}$ and $\mathrm{RV}_{t,T}^{\log}$ converge to $\int_t^T \sigma_u^2 du$ as the partition goes to zero.

Note, however, that the results of Proposition 3 are valid under any partition $\mathcal{T}$ as long as re-balancing in the underlying takes place on the same dates, $t_0, t_1, \ldots, t_n$. It, hence, does not matter whether we sample $\widetilde{RV}_{t,T}$ from daily, weekly, or monthly data as long as we re-balance at the daily, weekly, or monthly frequency. The reason for this lies in the aggregation property of the realized variance estimator given in Equation (3.4), see Neuberger (2012). The aggregation property essentially tells us that for any real-valued function g, for any martingale process X_t, for any measure $\mathbb{M}$, and for any times $0 \leq s \leq t \leq u \leq T$,

$$E_t^{\mathbb{M}}\left[g(X_u - X_s) - g(X_u - X_t) - g(X_t - X_s)\right] = 0.$$

In other words, for any function g which satisfies this restriction and if X_t is the forward price $F_{t,T}$—which we know is a martingale under the $\mathbb{Q}_T$-measure—then the discretely-sampled payoff $\sum_{i=1}^n g(F_{t_i,T} - F_{t_{i-1},T})$ has the same market price as the path-independent time T payoff of $g(F_{T,T} - F_{0,T})$. Thus, the realized variance which we calculate from higher frequency future returns (daily) is an unbiased estimate of the lower frequency counterpart (monthly). A priori, we do not expect the effect of discrete sampling to be large for pricing (see Broadie and Jain (2008)), however, the aggregation property tells us that it is exactly zero for the measure of realized variance that we propose.

Next, we implement the generalized Treasury variance swap introduced in Proposition 3 using a large panel of Treasury options data. We first outline how to convert American option prices on futures into European option prices on forwards that are needed to calculate the model-free implied variance measure. We will also compare the payoffs of the GTVS (Proposition 3) with those of the LTVS (Proposition 2).

3.2 Data and Measurement of Variance Swaps

In this section, we briefly introduce the data used in our analysis. We use futures and options data to construct the bond variance swap payoffs. To put our results for the Treasury market in perspective, we also calculate returns to variance trading strategies using S&P500 futures and options. Before we can start, however, we need to account for two features of the data that mainly affect the calculation of the fair strike. First, the observed options are written on futures rather than on forwards as implied by our theoretical contracts in Propositions 2 and 3. Second, the futures options are American rather than European. Based on existing evidence for the Eurodollar futures market, we expect both effects to be small. For example, Flesaker (1993) and Cakici and Zhu (2001) show in the Eurodollar futures market that the effect of having futures as the underlying as opposed to forward contracts is very small especially for options with shorter maturities.

The data is available from October 1982, May 1985, and June 1990 to May 2012 for the 30 year, 10 year, and 5 year Treasury bond futures and options, respectively. Using a monthly frequency throughout the paper, we have at most 355, 325, and 264 observations available, respectively. To make our trading strategies comparable, we present our baseline results using a sample starting in June 1990.

3.2.1 Futures and Options Data

Treasury Futures and Options: To calculate implied and realized variance measures for Treasury bonds, we use futures and options data from the Chicago Mercantile Exchange (CME). For our benchmark results we use end-of-day price data for the 30-year Treasury bond futures, the 10- and 5-year Treasury notes futures, and end-of-day prices of options written on the underlying futures, respectively.[5]

Treasury futures are traded electronically as well as by open outcry. While the quality of electronic trading data is higher, the data only becomes available in August 2000. To maximize our time span, we use data from electronic as well as pit trading sessions.

The contract months for the Treasury futures are the first three (30y Treasury futures) or five (10y and 5y Treasury futures) consecutive contracts in the March, June, September, and December quarterly cycle. This means that at any given point in time, up to five contracts on the same underlying are traded. To get one time series, we roll the futures on the 28^{th} of the month preceding the contract month.

For options, the contract months are the first three consecutive months (two serial expirations and one quarterly expiration) plus the next two (30y futures) or four (10y and 5y futures) months in the March, June, September, and December quarterly cycle. Serials exercise into the first nearby quarterly futures contract, quarterlies exercise into futures contracts of the same delivery period. We roll our options data consistent with the procedure applied to the futures.[6]

S&P500 Index Futures and Options: In line with our approach for Treasuries, we calculate the implied and realized variance measures for the stock market using futures and options on the S&P500 index from CME. The sample period is from January 1983 to May 2012.[7]

3.2.2 Differences between Futures and Forwards, and the Effect on Option

Options on Treasury and S&P500 index futures are American. We need to adjust for these two features since Equation (3.5) is derived under the assumption that the options are

[5]We use settlement prices for both options and futures which do not suffer from stale trading or bid-ask spreads. CME calculates settlement prices simultaneously for all options based on their last bid and ask.

[6]Detailed information about the contract specifications of Treasury futures and options can be found on the CME website, www.cmegroup.com.

[7]We compare our results to the VIX and VXO measures that are calculated using options on the S&P500 cash index instead of S&P500 index futures. The VIX is the implied volatility calculated using a model-free approach, whereas the VXO is calculated using the Black and Scholes (1973) implied volatility. The VIX is available starting in January 1990 and the VXO is available since January 1986. Over the common sample period, the VIX and our implied volatility measure from index futures options using the same methodology have a correlation of over 99.4% and the root mean squared error is below 1%.

European and written on forwards. In order to obtain implied volatilities from the available option prices, we hence need a suitable model for the futures and American feature.

The approach we use closely follows Cakici and Zhu (2001) and is outlined in the Online Appendix. In line with the results in Cakici and Zhu (2001), we find the adjustment from American futures options to European forward options to be very small and never exceeding two percent of the implied volatility level for all tenors and option maturities we consider. For the 30y Treasury options the difference ranges between 0.60% and 0.91% of the price for short-term and long-term options, respectively. For the 10y Treasury options, the average differences range between 0.65% and 0.90%, and finally for 5y Treasury options they range between 0.56% and 0.77%, respectively.[8]

3.2.3 Construction of Implied Volatility

In this section we show how to empirically calculate the fair strike of the variance swap at time t with maturity T. In line with the equity literature, we report our benchmark results for a horizon $\tau = T - t$ of one-month. However, we also consider a term structure using horizons ranging from ten days to one year.

Before including an option in the calculation, we apply a set of filters to clean the data: (i) We eliminate all data where either the futures or option price, the strike, the maturity, or the open interest are equal to zero. (ii) We also delete data when option prices fail to pass the no arbitrage boundary conditions.[9] (iii) We eliminate deep in-the-money options, i.e., we eliminate calls if the strike is less than $0.94 \times F_{t,T}$ and puts if the strike is greater than $1.06 \times F_{t,T}$, where $F_{t,T}$ is the underlying futures price. Thus, note that we do not restrict ourselves to out-of-the-money options only. Using strike prices greater than $0.94 \times F_{t,T}$ for calls or less than $1.06 \times F_{t,T}$ for puts means we still include some in-the-money and near-the-money options. At each maturity, we then fit a spline for the available implied volatilities against their corresponding strike prices and from the fitted spline, we obtain a fine grid of implied volatilities. We then convert the grid of implied volatilities into European option

[8]To save space, we defer more detailed summary statistics to the Online Appendix.

[9]The boundary condition for the call options are $C_{t,T}(K) - \max\left(F_{t,T} - K\right)^{+} \geq 0$ and for the put options it is $P_{t,T}(K) - \max\left(K - F_{t,T}\right)^{+} \geq 0$.

prices and numerically evaluate Equation (3.5) with daily options data on 30y, 10y, and 5y Treasury futures, and the S&P500 futures adjusted for the forward/futures feature.

3.2.4 Construction of Realized Volatility

We now describe the construction of the realized variance measures used to calculate the payoffs at maturity T to a generic variance swap that pays the difference between realized variance and the strike $\tilde{F}_{t,T}$.

Proposition 3 implies that the fair strike of the GTVS is given by $\tilde{F}_{t,T}$ irrespective of the sampling partition and the underlying price process if we use the definition given in Equation (3.4) for the realized variance leg, $\widetilde{RV}_{t,T}$. As variance swaps generally use a daily sampling frequency, the benchmark results we report are based on daily data. In addition, we also calculate the standard measure of realized variance, $RV^{\log}_{t,T}$, using daily log returns.

Note that the realized variance measures $\widetilde{RV}_{t,T}$ and $RV^{\log}_{t,T}$ are only observed ex post at maturity T of the variance swap. For our benchmark horizon of one month and using a daily sampling frequency, $\widetilde{RV}_{t,T}$ and $RV^{\log}_{t,T}$ are therefore based on the futures prices of the previous 21 trading days. We use end-of-day prices measured at 14:00 CT in line with the end of pit trading hours at the CME. In addition to the one-month horizon, we calculate a term structure of realized variances ranging from ten days to one year (252 trading days) to match the term structure of implied variances.

3.2.5 Summary Statistics and Variance Risk Premiums

Table 3.1 reports summary statistics of implied volatilities as described in Equation (3.5) and Figure 3.1 provides a plot of the average implied volatilities. As it is custom in practice to report volatilities, we present the square root of the annualized variances expressed in percent.

We note that all term structures are downward sloping in the maturity dimension. For 30y (10y, 5y) Treasury options, average implied volatilities range from 11.8% (8.3%, 5.5%) for the ten-day options to 9.9% (6.7%, 4.5%) for options with one-year to maturity. Hence, along the tenor dimension, the term structure is upward sloping. The slope of the term

Table 3.1
Summary Statistics Treasury Implied Volatilities

The table reports means, standard deviations, minima and maxima of annualized implied volatilities for the three different tenors (30y, 10y, and 5y) and different maturities ranging from ten days to one year. Volatilities are the square root of the implied variances extracted from daily option prices using Equation (3.5). All numbers are annualized and expressed in percent. Data is monthly and runs from June 1990 to May 2012.

maturity	10d	20d	1m	3m	6m	9m	12m
	Panel A: 30y Treasury						
mean	11.83	11.21	10.98	10.47	10.35	10.06	9.88
std	3.57	2.85	2.65	2.41	2.11	2.09	2.13
min	4.21	6.01	6.60	5.78	6.07	5.82	5.41
max	31.18	26.20	24.15	21.12	20.30	20.02	19.88
	Panel B: 10y Treasury						
mean	8.27	7.65	7.43	6.90	6.87	6.77	6.72
std	2.56	1.77	1.63	1.54	1.66	1.86	1.98
min	0.31	4.32	4.26	2.99	3.11	2.91	2.66
max	17.08	14.76	14.33	13.08	12.52	13.69	14.24
	Panel C: 5y Treasury						
mean	5.54	4.99	4.79	4.49	4.49	4.48	4.47
std	1.94	1.29	1.16	1.22	1.36	1.45	1.50
min	0.74	2.34	2.59	2.06	1.77	1.46	1.28
max	11.88	9.90	9.69	9.40	9.33	9.30	9.29

structure of implied variances, which we define as the difference between the implied variance of a one-year option and a one-month option, is negative on average. The equity index option literature finds both upward and downward-sloping implied volatility term structures depending on the method used (see, e.g., Aït-Sahalia, Karaman, and Mancini (2015)). Using options on S&P500 futures, we find a slightly downward sloping term structure ranging between 20.1% and 19.4%.

In analogy to the calculation of the equity risk premium, we call the difference between the (ex post) realized and the implied variance in the variance swap the (unconditional) variance risk premium. In line with this definition, we use the terms ex post or realized

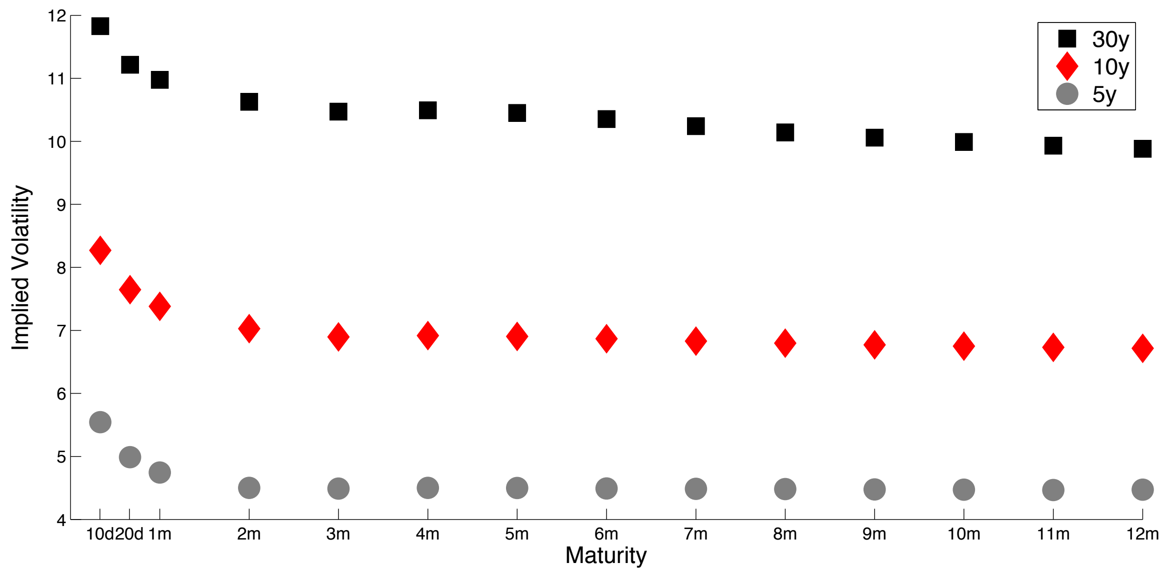

Figure 3.1. Term Structure of Implied Volatilities

The figure plots average implied volatilities for the three different tenors (30y, 10y, and 5y) and different maturities ranging from ten days to one year. Volatilities are the square root of the implied variances extracted from daily option prices using Equation (3.5). All numbers are annualized and expressed in percent. Data is monthly and runs from June 1990 to May 2012.

variance risk premiums to describe the time-varying differences between the realized and the implied variances.[10] The realized variance measure is calculated by summing up daily data.[11] We report average variance risk premiums in Table 3.2 together with the corresponding Sharpe ratios for different subsamples. Panel A reports variance risk premiums for the period June 1990 to May 2012 which is the time span for which we have available data for all tenors. There are several noteworthy observations. First, variance risk premiums are negative in line with expectations (investors are willing to pay a premium to be protected against volatility spikes) and existing research for the equity market (with the sole exception of 30y Treasury variance risk premiums for the one year horizon). Second, the downward sloping implied volatility term structures imply an upward sloping term structure of the

[10]Note that we explicitly take an *ex post* view by focusing on realized variance which is in line with the payoff traders face in a variance swap. Taking an *ex ante* view to study *conditional* variance risk premiums requires to form expectations about variance using, for example, an econometric model. See, e.g., Mueller, Sabtchevsky, Vedolin, and Whelan (2016) for an exploration of conditional variance risk premiums in the stock and bond markets.

[11]For different maturities we sum up daily data over various horizons. The sample averages are then taken over the whole sample. Regardless of the maturity, this is essentially the same as taking an unconditional average of all daily values.

ex post variance risk premiums. Given that the variance risk premiums are essentially all negative on average, this of course means that in absolute terms, the term structure of variance risk premiums is downward sloping. The same is true for the Sharpe ratios. These declining Sharpe ratios are consistent with findings of van Binsbergen and Koijen (2016), who document that Sharpe ratios in a range of markets (stocks, bonds, corporate bonds, index straddles, and housing) decline with maturity.[12] Third, the variance risk premium term structures look very similar across tenors. In absolute terms, the variance risk premium term structures are all downward sloping while the Sharpe ratios decline almost monotonically with the horizon. At the same time, average variance risk premiums increase with the tenor (again in absolute terms) and, for example, average 30y variance risk premiums are always larger than the corresponding variance risk premiums for the 10y and 5y tenors.

To check robustness of this pattern about the average shape of the term structure of variance risk premiums, we examine term structures for different subsamples. Figure 3.2 plots the average variance risk premiums and Sharpe ratios for a sample that excludes the financial crisis and ends in December 2007 (left panels) and a sample that starts in 2008 (right panels). We note that the variance risk premium surface looks very similar for both subsamples. However, the slope across tenors becomes much steeper during the post crisis period. For example, the difference between the ten-day variance risk premium on the 30y and 5y Treasury futures is -3.17 during the pre-crisis period, but increases to -6.45 after 2007. At the other end of the maturity spectrum, for one year options, we find that the difference between the variance risk premiums for 30y and 5y Treasury futures turns from being negative during the pre-crisis period to a positive value after the crisis, implying that while the 30y variance risk premium was more negative than the 5y variance risk premium in the early sample, this relationship has been reversed since 2008. A similar pattern emerges for Sharpe ratios which increase in absolute terms during the post-crisis periods, especially for shorter horizons while the Sharpe ratios for longer horizons are now slightly smaller.

[12] Applying a hedging based method for swaptions with different maturities, Duyvesteyn and de Zwart (2015) also find a downward sloping term structure of variance risk premiums in the swaptions market. Similarly, Aït-Sahalia, Karaman, and Mancini (2015) and Dew-Becker, Giglio, Le, and Rodriguez (2016) also find that the term structure of variance risk premiums is downward sloping for the S&P500.

Table 3.2
Term Structure of Variance Risk Premiums: Means and Sharpe Ratios

The table reports average variance risk premiums (mean) and average Sharpe ratios (SR) for the three different tenors (30y, 10y, and 5y) and different maturities ranging from ten days to one year. Variance risk premiums are computed by subtracting the implied variance as in Equation (3.5) from the ex-post realized variance as in Equation (3.4). They are monthly and expressed in squared percent. Sharpe ratios are calculated as the average variance risk premiums divided by the corresponding standard deviation of the monthly variance risk premiums. Panel A reports the results for the benchmarks sample period from June 1990 until May 2012 while Panel B shows the corresponding numbers for the maximally available data set that goes back to October 1982, May 1985, and June 1990 for the respective tenors, 30y, 10y, and 5y. Data is sampled monthly.

maturity		10d	20d	1m	3m	6m	9m	12m
Panel A: June 1990 - May 2012								
	30y	-4.90	-3.22	-2.41	-1.35	-0.88	-0.32	0.08
mean	10y	-3.15	-1.81	-1.33	-0.61	-0.49	-0.46	-0.40
	5y	-1.08	-0.61	-0.44	-0.16	-0.25	-0.27	-0.37
	30y	-0.81	-0.67	-0.51	-0.34	-0.23	-0.08	0.02
SR	10y	-0.94	-0.76	-0.69	-0.36	-0.25	-0.22	-0.18
	5y	-0.62	-0.48	-0.45	-0.19	-0.24	-0.25	-0.24
Panel B: all available data								
	30y	-6.82	-3.75	-2.84	-1.31	-0.56	0.18	0.66
mean	10y	-4.68	-2.46	-1.62	-0.52	-0.34	-0.26	-0.17
	5y	-1.08	-0.61	-0.44	-0.16	-0.25	-0.27	-0.37
	30y	-0.64	-0.50	-0.45	-0.26	-0.12	0.04	0.15
SR	10y	-0.67	-0.54	-0.60	-0.23	-0.14	-0.10	-0.07
	5y	-0.62	-0.48	-0.45	-0.19	-0.24	-0.25	-0.24

Finally, we examine to what extent the implied variances across tenors and option horizons are driven by a common set of factors. To this end, we perform a principal components analysis (PCA) on monthly changes in the implied variances for the three tenors and horizons ranging from ten days up to one year. We find that the first principal component (PC) explains around 70% of the overall variation while the second PC accounts for an additional 15%. Since the third PC captures 5% of the overall variation, the first three PCs are suffi-

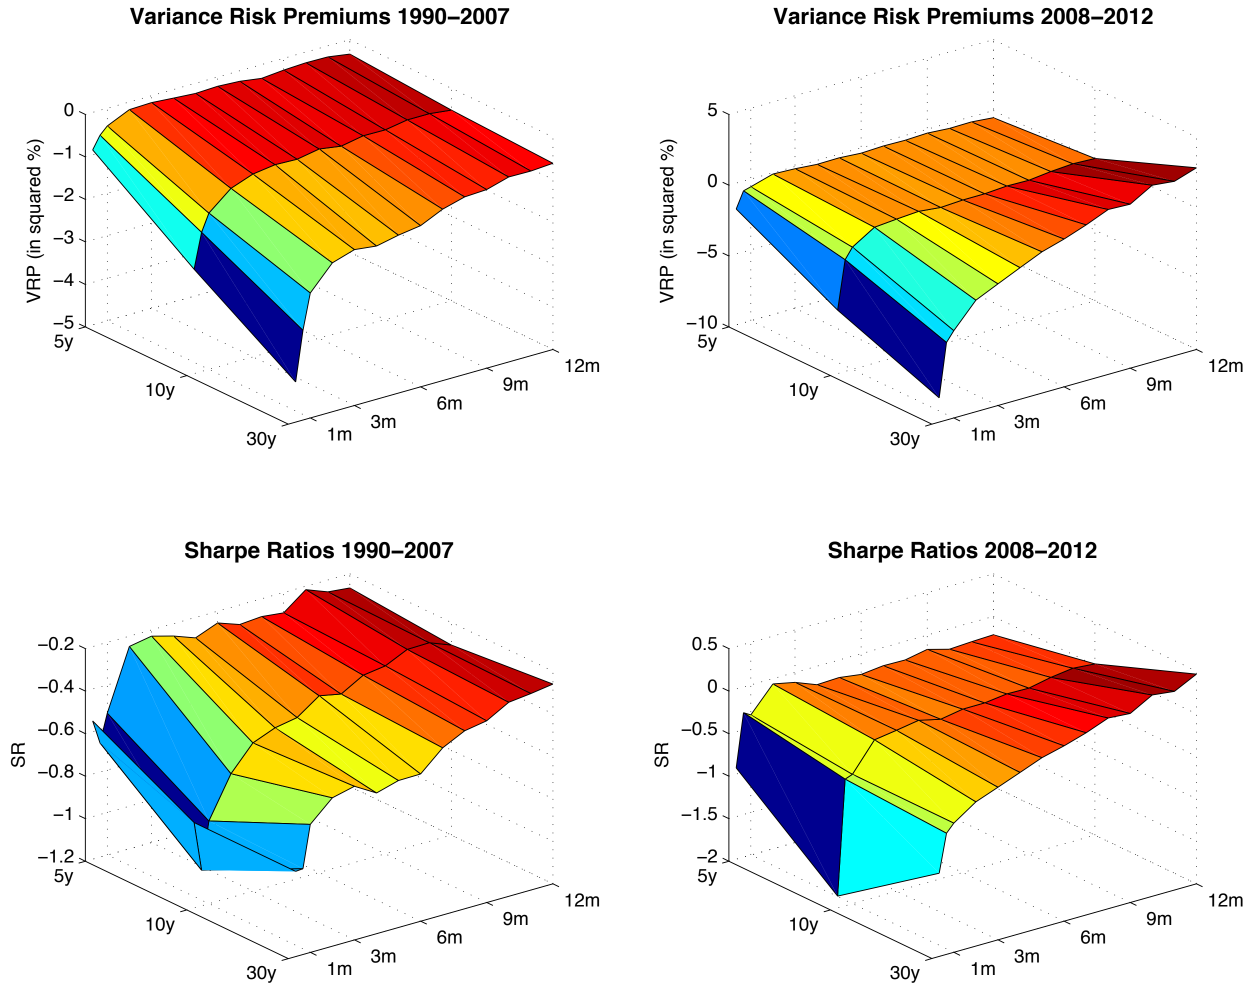

Figure 3.2. Term Structure of Variance Risk Premiums and Sharpe Ratios

The figure plots average variance risk premiums (upper panels) and average Sharpe ratios (lower panels) for the three different tenors (30y, 10y, and 5y) and different maturities ranging from ten days to one year. The left side panels display the results for the pre-crisis period from June 1990 until December 2007 while the right side panels show the corresponding numbers for the crisis and post-crisis period from January 2008 to May 2012. Variance risk premiums are computed by subtracting the implied variance as in Equation (3.5) from the ex-post realized variance as in Equation (3.4), and then expressed in squared percent after scaling them to monthly measures. Sharpe ratios are calculated as the average variance risk premiums divided by the corresponding standard deviation of the variance risk premiums. Data is monthly and runs from June 1990 to May 2012.

cient to explain almost 90% of the variation across tenors and for all horizons.[13] Figure 3.3

[13]This is in line with results for the swaptions implied volatility surface reported in Trolle and Schwartz (2014).

plots the factor loadings on the implied variance surface for the first two PCs. The first PC acts as a level factor, having a roughly uniform impact on implied variances for all tenors and option maturities. On the other hand, the second PC acts as a slope factor, having a more prominent impact on long-horizon variances and monotonically decreasing loadings across tenors. To summarize, only a small number of factors is needed to capture the variation in implied variances across tenors and option maturities, and the factors present the usual level and slope effects as in other asset classes (see, for example, Litterman and Scheinkman (1991) and Lustig, Roussanov, and Verdelhan (2011)).

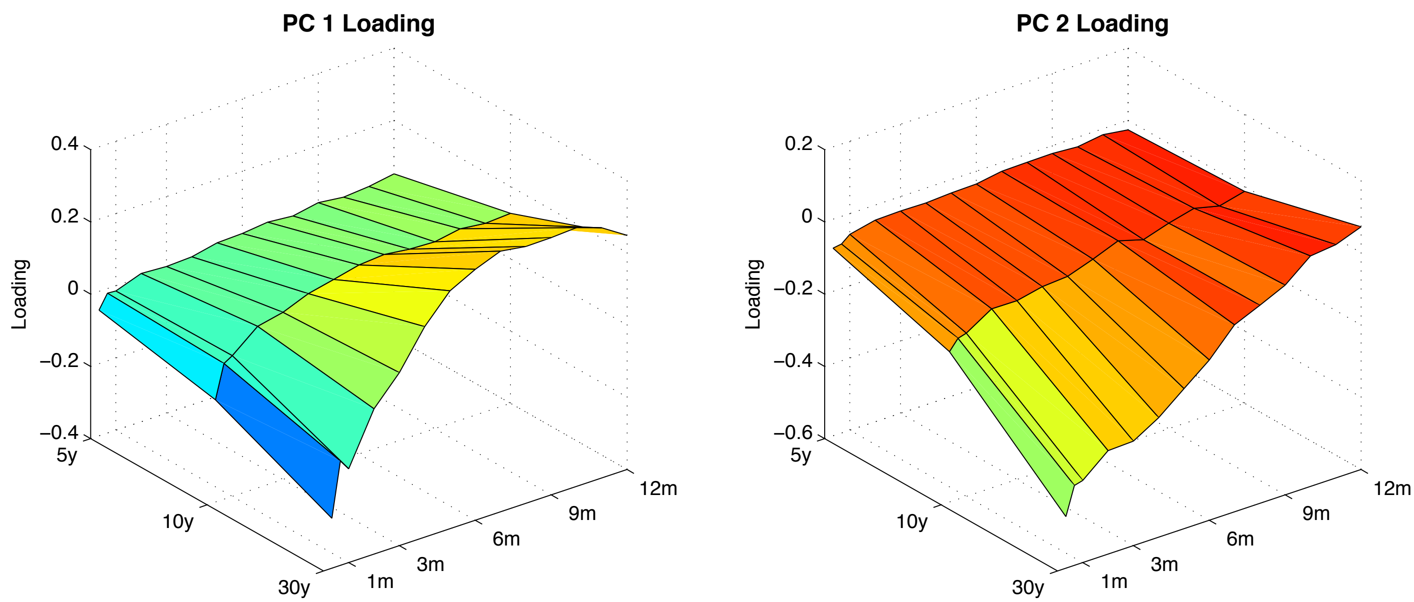

Figure 3.3. Factor Loadings on (Standardized) Implied Variances

The figure plots the factor loading of the first (left panel) and second (right panel) principal component of implied variances as in Equation (3.5) across the three tenors (5y, 10y and 30y) and for maturities ranging between ten days and one year. The loadings are constructed from the eigenvectors corresponding to the two largest eigenvalues of the correlation matrix of monthly changes in implied variances. Data is monthly and runs from June 1990 to May 2012.

3.3 Empirics

In this section, we evaluate different variance trading strategies using fixed income options. We first present summary statistics for the variance swap as well as straddles, which also provide exposure to variance risk. Our findings reveal that trading variance in fixed income markets is attractive even if we condition on other risk factors. We further benchmark our results for the fixed income markets against those realized by trading variance in the equity index market using options on S&P500 futures. As it is well established that variance trading in the equity market has been popular and profitable over the last decade, comparing the returns from the Treasury and equity market helps to put our novel results into perspective.

3.3.1 Trading Strategies

First, we study returns to a strategy based on our generalized Treasury variance swap which are calculated as follows

$$r_{t,T}^{GTVS} = \frac{\widetilde{RV}_{t,T}}{\mathbb{E}_t^{\mathbb{Q}_T}\left[\widetilde{RV}_{t,T}\right]} - 1, \tag{3.6}$$

where T is the maturity of the contract, $\widetilde{RV}_{t,T}$ is the realized variance as defined in Equation (3.4) and the denominator is given by Equation (3.5). This means that the return is the ex post variance risk premium scaled by the fair strike price. At the same time, this is the excess returns to a fully collateralized long position in the variance swap that posts $\mathbb{E}_t^{\mathbb{Q}_T}\left[\widetilde{RV}_{t,T}\right]$ dollars of collateral and receives $\widetilde{RV}_{t,T}$ at expiration plus interest on the collateral. Alternatively, one can label what we define as the realized variance risk premium the realized excess return on a unit position in the variance swap, whereas Equation (3.6) represents the scaled version thereof. However, we prefer the term "ex post variance risk premium" to highlight the fact that our main object of interest is the realized quantity of the ex ante variance risk premium that the literature generally focuses on.

We then compare these returns to those of two standard volatility trading strategies using straddles. First, we consider an at-the-money straddle, a classical position for getting exposure to volatility. Unfortunately, such a position loses sensitivity to volatility as the underlying moves away from the strike. To this end, in addition to an unhedged straddle, we also consider a delta-hedged position. Coval and Shumway (2001) and Santa-Clara and Saretto (2009) show that trading in straddles yields very attractive (annualized) Sharpe ratios above one for options on the S&P500 index futures.

To calculate the delta-hedged returns, we proceed as follows. Each month, we simultaneously purchase an at-the-money call and put option with 30 days to expiration (or the closest to 30 days). We track the path of the straddle until its expiration date and on a daily basis we go long or short the corresponding underlying future such that the whole position becomes delta-neutral at the end of each day. Denoting the delta of a straddle on a given date t as $\Delta_{S,t}$, the accumulated profit and loss from this hedging activity can be written as:

$$\sum_{i=t}^{T} -\Delta_{S,i-1} \left(F_{i,T} - F_{i-1,T}\right).$$

Hence, hold to expiration returns of the delta-hedged straddle strategy are defined as:

$$r_{t,T}^{\Delta S} = \frac{\left(K - F_{T,T}\right)^{+} + \left(F_{T,T} - K\right)^{+} - \sum_{i=t}^{T} \Delta_{S,i-1} \left(F_{i,T} - F_{i-1,T}\right)}{P_{t,T}\left(K\right) + C_{t,T}\left(K\right)} - 1. \tag{3.7}$$

Summary Statistics

Table 3.3 reports summary statistics of the trading strategies together with different performance measures, the Sharpe ratio and Jensen's alpha. Note that while Equation (3.6) is defined for an arbitrary maturity $T - t$ we focus on the one-month horizon in line with the sampling frequency for our benchmark results. Hence, in what follows we study one-month excess returns on one-month variance swaps only. To calculate the alpha, we employ two different market indices depending on whether we use Treasury or equity options, respectively. The market return for the options on the S&P500 futures is the value weighted excess return on all stocks in CRSP and for the bond options, we use the total return on the Barclays US Treasury bond index available from Datastream.

Panel A summarizes the annualized returns on the generalized Treasury variance swap, our main object of interest. We note that shorting a variance swap produces a monthly average return of around 20% for options on 30y, 10y, and 5y Treasury futures. All average returns are highly significantly different from zero as indicated by the t-statistics which range between 6.61 for the 5y and 12.71 for the 10y futures. The volatilities of the variance swap trading strategies are relatively small, leading to annualized Sharpe ratios ranging between 1.4 for the 5y and 2.7 for the 10y Treasury futures, respectively.[14] The associated alphas are only marginally smaller than the average returns ranging between between 18% and 27% and, hence, the market return does not explain the variance swap returns at all. Trolle

[14]Note that the tables report monthly Sharpe ratios.

Table 3.3
Summary Statistics Option Trading Strategies

This table reports monthly summary statistics for one-month returns on three different trading strategies described in Section 3.3.1: mean, standard deviation, maximum, skewness, kurtosis, Sharpe ratio (SR), and alpha. Panel A presents the summary statistics of monthly returns on the generalized Treasury variance swap defined in Equation (3.6) across three tenors (5y, 10y and 30y) and with a maturity $T - t = 1$ month. Panel B and Panel C report the summary statistics of monthly returns on un-hedged or delta-hedged at-the-money straddles with a maturity of one month, respectively. The delta-hedged straddle return is defined in Equation (3.7). The un-hedged return is the same but without the accumulated profit and loss from the hedging activity, i.e., $\sum_{i=t}^{T} -\Delta_{S,i-1}\left(F_{i,T} - F_{i-1,T}\right)$ in Equation (3.7). In addition, we also report the corresponding summary statistics for the trading strategies using options on S&P500 index futures. t-Statistics reported in parentheses are adjusted according to Newey and West (1987). Data is sampled monthly and runs from June 1990 to May 2012.

Panel A: Variance Swap

	mean	t-stat	std	max	skew	kurt	SR	alpha	t-stat
30y	-0.212	(-8.98)	0.382	2.167	2.128	7.805	-0.554	-0.202	(-8.54)
10y	-0.276	(-12.71)	0.353	2.123	1.893	8.104	-0.784	-0.270	(-11.14)
5y	-0.187	(-6.61)	0.460	2.313	2.191	7.857	-0.407	-0.181	(-5.39)
S&P500	-0.314	(-9.40)	0.542	4.182	3.711	21.634	-0.580	-0.279	(-7.33)

Panel B: Straddles not hedged

	mean	t-stat	std	max	skew	kurt	SR	alpha	t-stat
30y	0.007	(0.17)	0.710	2.409	0.912	3.723	0.011	0.005	(0.10)
10y	-0.019	(-0.42)	0.724	4.097	1.227	6.461	-0.026	-0.018	(-0.36)
5y	0.033	(0.70)	0.758	4.340	1.365	6.876	0.044	0.023	(0.45)
S&P500	-0.125	(-2.95)	0.686	2.895	1.379	5.607	-0.182	-0.150	(-3.22)

Panel C: Straddles hedged

	mean	t-stat	std	max	skew	kurt	SR	alpha	t-stat
30y	-0.034	(-2.47)	0.219	1.167	0.831	6.614	-0.153	-0.034	(-2.19)
10y	-0.070	(-4.91)	0.231	0.711	0.146	3.287	-0.304	-0.069	(-4.31)
5y	-0.045	(-2.99)	0.244	0.815	0.366	3.611	-0.185	-0.042	(-2.53)
S&P500	-0.075	(-4.11)	0.296	1.421	1.599	7.537	-0.254	-0.081	(-4.04)

and Schwartz (2014) study variance swaps in the swaptions market and report even higher average returns ranging between 44% and up to 66% per month. The associated volatilities are considerably higher than in the Treasury options market, so the Sharpe ratios they find are in line with ours. The returns to shorting a variance swap are negatively skewed and exhibit excess kurtosis, however, the values are comparable to the strategies based on straddles. Trading a variance swap on the S&P500 index futures is similarly attractive: The return is 31% per month, with an annualized Sharpe ratio of two. However, the reward comes with some additional risk as the strategy has much fatter tails with a kurtosis as high as 22.

Panel B reports the summary statistics for at-the-money (ATM) straddles without taking a position in the underlying futures. The average returns are much smaller than for the variance swap strategy. In fact, shorting an ATM straddle with options written on Treasury futures produces an average return close to, and not significantly different from, zero. Similarly, we also find the strategies' alpha to be insignificant. Moreover, the associated risk as proxied by the volatility is more than 50% higher compared to the variance swap strategies.

Panel C in Table 3.3 presents the results for the delta-hedged straddle strategies. By delta-hedging the straddle, the volatility of the strategy becomes considerably smaller. Average returns are highly significant and range between -3.4% (30y) and -7% (10y) per month. The corresponding annualized Sharpe ratios are 0.5 for the 30y and 1.1 for the 10y Treasury futures, respectively. The associated alphas are significant and also range between -3.4% (30y) and -7% (10y). Overall, we conclude that delta-hedged straddle strategies yield attractive returns but variance swaps still produce average returns and alphas that are on average larger by a factor of roughly four, both for Treasury as well as for S&P500 futures strategies.

To further investigate the variance swap strategy, we plot the time series of the GTVS returns for the 30y Treasury futures in Figure 3.4 (upper panel) together with the associated realized and implied volatility measures (lower panel). We note that most of the time, implied volatility exceeds the ex post realized one. Hence, on average, a strategy that is long realized and short implied variance produces a negative return. At the same time, there are some very distinct positive spikes which, interestingly, correspond to general periods of distress as can be gauged from the annotations in the upper panel. However, there is one single spike which is very specific to the bond market and it coincides with the large bond-market sell-off in July 2003 due to mortgage hedging activity (see Malkhozov, Mueller, Vedolin, and Venter (2016b)). To compare, we also plot the returns to the variance swap on the S&P500 futures in Figure 3.5. Different from the bond market, there is for example no

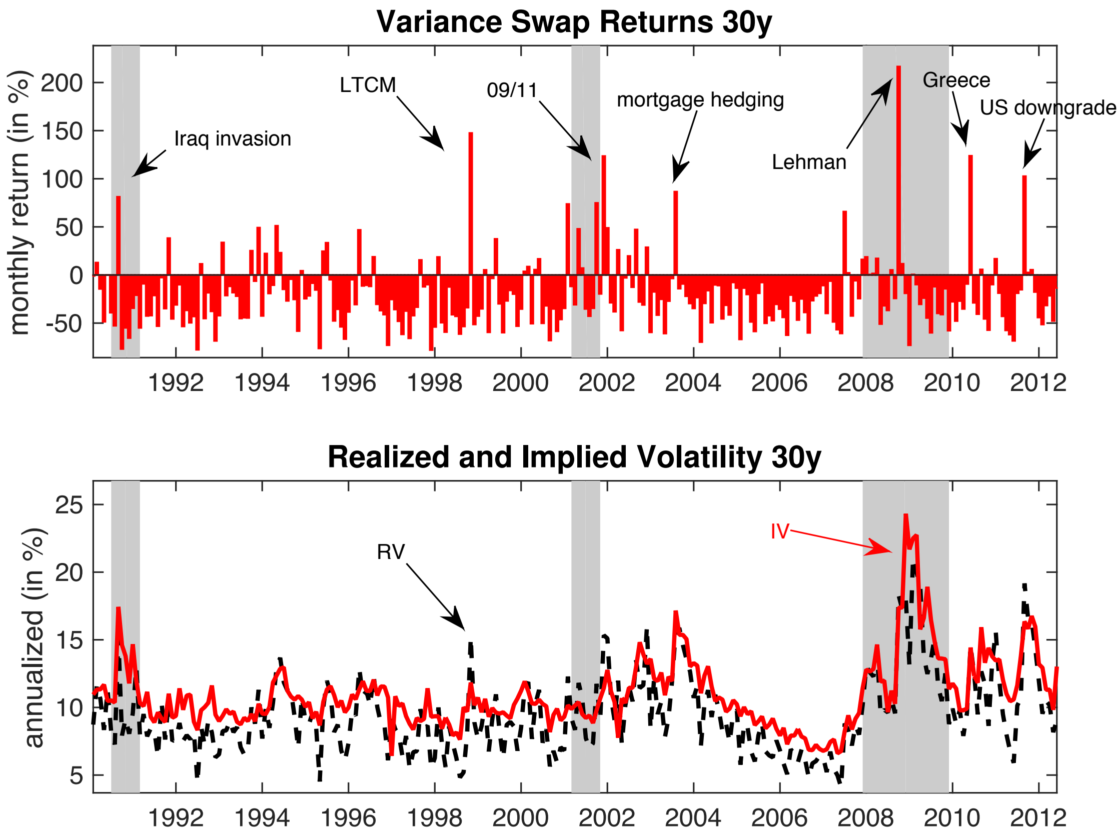

Figure 3.4. Generalized Treasury Variance Swap Returns 30y Treasury

The upper panel plots the monthly returns of the generalized Treasury variance swap (GTVS) with a 30y tenor and one month to maturity. The return is computed as the payoff of the one-month variance swap (implied variance minus ex-post realized variance) scaled by the fair strike of the variance swap (implied variance). The lower panel plots (annualized) realized and implied volatilities for the 30y Treasury futures. Gray bars indicate NBER recessions. Data is monthly and runs from January 1990 to May 2012.

spike in July 2003. On the other hand, there is one large positive spike in July 2002, when the S&P500 index lost 8% between June and July.

Statistical Significance

One might worry about statistical significance as it is well known that option trading strategies are non-Gaussian and the extant literature shows that performance measures such as the Sharpe ratio can change dramatically if returns do not follow a Normal distribution (see Lo (2002)).

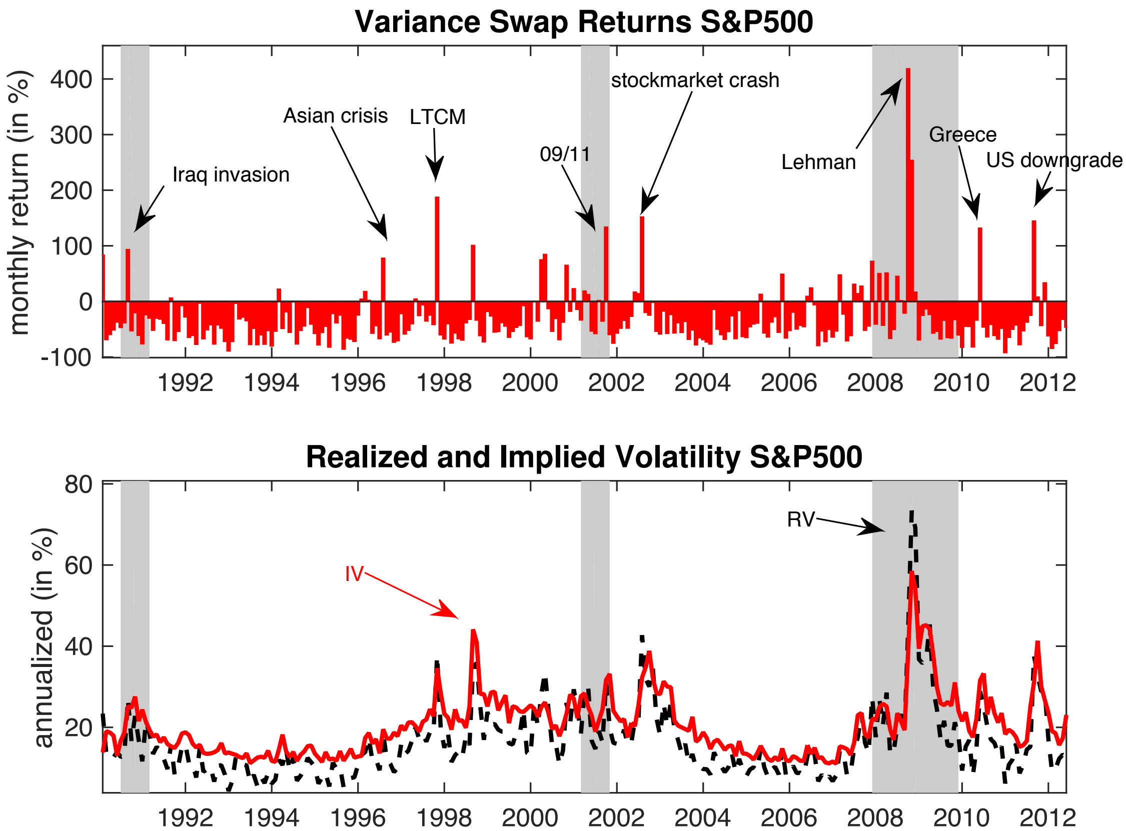

Figure 3.5. Variance Swap Returns S&P500

The upper panel plots the monthly returns of the variance swap on the S&P500 with a maturity of one month. The return is computed as the payoff of the one-month variance swap (implied variance minus ex-post realized variance) scaled by the fair strike of the variance swap (implied variance). The lower panel plots (annualized) realized and implied volatilities for the S&P500. Gray bars indicate NBER recessions. Data is monthly and runs from January 1990 to May 2012.

Thus, in the following, we use a studentized bootstrap to obtain confidence intervals on the mean, Sharpe ratio, and alpha of the trading strategies discussed earlier. Using a sample of 10,000 bootstrapped repetitions, we report 95% confidence intervals in Table 3.4.

In line with the previous results, average returns for the variance swap and delta-hedged straddle strategies are all significantly different from zero, as none of the confidence intervals includes the zero itself. The same applies to the Sharpe ratios as well as for the alphas.

Table 3.4
Bootstrapped Confidence Intervals

This table reports 95% bootstrapped confidence intervals for mean, Sharpe ratio (SR), and alpha for three trading strategies presented in Table 3.3. The empirical distribution of returns is obtained from 10,000 studentized bootstrap repetitions of our sample. Data is monthly and runs from June 1990 to May 2012.

	mean		SR		alpha	
Panel A: Variance Swap						
30y	-0.248	-0.177	-0.741	-0.415	-0.205	-0.199
10y	-0.311	-0.250	-1.050	-0.643	-0.272	-0.267
5y	-0.242	-0.138	-0.581	-0.254	-0.184	-0.177
S&P500	-0.376	-0.252	-0.851	-0.373	-0.282	-0.276
Panel B: Straddles not hedged						
	mean		SR		alpha	
30y	-0.097	0.078	-0.128	0.091	-0.056	0.042
10y	-0.099	0.060	-0.163	0.047	-0.041	0.048
5y	-0.051	0.116	-0.066	0.121	-0.032	0.280
S&P500	-0.167	-0.048	-0.316	-0.078	-0.200	-0.104
Panel C: Straddles hedged						
	mean		SR		alpha	
30y	-0.065	-0.014	-0.287	-0.045	-0.042	-0.021
10y	-0.091	-0.045	-0.432	-0.178	-0.088	-0.070
5y	-0.063	-0.014	-0.276	-0.056	-0.068	-0.027
S&P500	-0.099	-0.037	-0.415	-0.117	-0.095	-0.072

Risk-Adjusted Returns

In this section, we explore how and whether the returns of these trading strategies are related to market risk. To this end, we regress the strategy returns onto the returns of different equity and bond portfolios. In particular, we control for the market, size (SMB), book-to-market (HML), and momentum (MOM) portfolios. We also include two liquidity factors for bond and equity markets: The liquidity factor extracted from bonds used in Malkhozov, Mueller, Vedolin, and Venter (2016a) and the Pástor and Stambaugh (2003) equity liquidity factor. In order to make the two measures comparable, we multiply the latter by minus one to get an illiquidity measure. Table 3.5 presents the regression results

Table 3.5
Risk-Adjusted Returns Trading Strategies

This table reports Newey and West (1987) adjusted t-statistics of OLS coefficients from regressing the returns of three different trading strategies (the one-month generalized Treasury variance swap, un-hedged straddle, and delta-hedged straddle) on the market excess return (MRKT), size (SMB), book-to-market (HML), momentum factors (MOM), and an illiquidity factor.

$$r_t^i = \alpha + \beta_1 r_t^{\text{MRKT}} + \beta_2 r_t^{\text{SMB}} + \beta_3 r_t^{\text{HML}} + \beta_4 r_t^{\text{MOM}} + \beta_5 r_t^{\text{ILLIQ}} + \epsilon_t$$

For the bond option trading strategies we use the Barclays US Treasury bond index as a proxy for the market return. For the S&P500 futures strategies we use the excess return on the value-weighted return of all CRSP firms. The illiquidity measure for the bond market is taken from Malkhozov, Mueller, Vedolin, and Venter (2016b) and for the equity market we use the negative of the Pástor and Stambaugh (2003) liquidity factor such that a high value again measures illiquidity. The first column reports Jensen's alpha together with its t-statistic in parentheses. Data is monthly and runs from June 1990 to May 2012.

Panel A: Variance Swap								
	alpha	t-stat	MRKT	SMB	HML	MOM	ILLIQ	Adj. R^2
30y	-0.261	(-4.44)	-2.317	-3.061	1.924	-2.572	1.562	10.37%
10y	-0.418	(-8.21)	-1.997	-2.195	1.839	-2.483	3.211	10.48%
5y	-0.336	(-3.80)	-0.370	-2.106	2.466	-2.082	1.697	4.79%
S&P500	-0.321	(-9.65)	0.146	-1.133	-1.376	-1.653	4.249	31.72%

Panel B: Straddles not hedged								
	alpha	t-stat	MRKT	SMB	HML	MOM	ILLIQ	Adj. R^2
30y	-0.031	(-0.36)	0.264	-0.660	0.329	-0.985	0.591	0.72%
10y	-0.082	(-0.93)	0.261	-0.907	1.277	-1.076	0.903	1.61%
5y	-0.025	(-0.27)	0.895	-0.522	1.554	-1.018	0.564	1.56%
S&P500	-0.230	(-4.07)	0.111	-0.313	1.304	-1.126	0.440	2.76%

Panel C: Straddles hedged								
	alpha	t-stat	MRKT	SMB	HML	MOM	ILLIQ	Adj. R^2
30y	-0.071	(-2.58)	0.180	-0.238	0.888	-0.185	1.623	1.58%
10y	-0.125	(-4.48)	0.581	-1.191	1.922	-1.040	2.482	5.14%
5y	-0.083	(-2.83)	0.566	-0.245	2.410	-1.067	1.513	3.82%
S&P500	-0.176	(-4.67)	1.091	-1.095	0.410	-0.334	2.628	5.28%

for each strategy. We report the alpha together with its associated t-statistic while for the other regressors we only report (Newey and West (1987) adjusted) t-statistics.

Panel A reports the results for variance swap returns. We first note the high significance of the momentum factor (t-statistics all above two for the Treasury variance swap returns) and the borderline significance of the equity momentum factor for variance swaps on the S&P500. The alpha of the strategy is still negative and highly significant, which indicates that while these factors explain some of the variation (adjusted R^2 range between 4% to 10%), the majority is left unexplained.

The alphas for the un-hedged straddles (Panel B) are not statistically significant except for the S&P500. On the other hand, once the straddles are delta-hedged (Panel C) we find alphas to be highly significant. Naturally, due to the delta-hedging, the market return is insignificant. Moreover, we find the momentum portfolio to have no correlation with the strategy returns.

Overall, we reconfirm that trading variance in the fixed income market produces high average returns and attractive Sharpe ratios. The associated alphas are large and statistically significant even when controlling for standard risk factors.

Transaction Costs

It is well known that transaction costs lower the returns of option strategies and that the impact is increasing with decreasing moneyness, i.e., it is worst for deep out-of-the-money options (see, for example, George and Longstaff (1993) and Santa-Clara and Saretto (2009)). In the following, we explore the impact of bid-ask spreads onto the profitability of the trading strategies. To this end, we recalculate the summary statistics in Table 3.3 by assuming that we always buy at the ask price and sell at the bid price. The results are summarized in Table 3.6.

Not surprisingly, we find that average returns drop but shorting variance is still attractive: Annualized Sharpe ratios vary between 1.1 and 2.4 for the variance swaps. Moreover, average returns and alphas remain statistically significant both for the variance swaps and the delta-hedged straddles. We conclude that while bid-ask spreads lower average returns by up to one third, trading variance swaps is still very profitable.

Table 3.6
Summary Statistics Option Trading Strategies Bid-Ask Spread Adjusted

This table reports monthly summary statistics for one-month returns on three different trading strategies described in Section 3.3.1, taking into account bid and ask spreads: mean, standard deviation, Sharpe ratio (SR), and alpha. First, we report returns on the generalized Treasury variance swap defined in Equation (3.6) across three tenors (5y, 10y and 30y) and with a maturity $T - t = 1$ month. Second, we report the summary statistics of monthly returns on un-hedged or delta-hedged at-the-money straddles with a maturity of one month, respectively. The delta-hedged straddle return is defined in Equation (3.7). The un-hedged return is the same but without the accumulated profit and loss from the hedging activity, i.e., $\sum_{i=t}^{T} -\Delta_{S,i-1}(F_{i,T} - F_{i-1,T})$ in Equation (3.7). t-Statistics reported in parentheses are adjusted according to Newey and West (1987). Data is sampled monthly and runs from June 1990 to May 2012.

	Variance Swap			Straddles not hedged			Straddles hedged		
	30y	10y	5y	30y	10y	5y	30y	10y	5y
mean	-0.138	-0.254	-0.173	-0.001	-0.004	0.020	-0.021	-0.062	-0.035
t-stat	(-5.32)	(-11.42)	(-5.95)	(-0.01)	(-0.21)	(0.41)	(-2.11)	(-4.24)	(-2.57)
std	0.421	0.362	0.474	1.625	0.309	0.793	0.162	0.238	0.221
SR	-0.327	-0.703	-0.364	-0.001	-0.013	0.025	-0.130	-0.261	-0.158
alpha	-0.127	-0.245	-0.166	-0.002	-0.005	0.014	-0.027	-0.046	-0.038
t-stat	(-5.00)	(-10.69)	(-5.17)	(-0.00)	(-0.01)	(0.21)	(-1.96)	(-3.15)	(-3.26)

Margins

The strategies we consider are implemented using options traded on the CME. In practice, this means that investors are subject to margin constraints which are likely going to influence a strategy's profitability. For example, if a strategy of shorting variance leads to large losses, an investor could be forced to close down the position if she does not have unlimited funds. In the following, we follow a similar procedure as in Santa-Clara and Saretto (2009) and study the impact of margin requirements on the trading strategies discussed earlier.

In practice, variance swap strikes are quoted in terms of volatility (expressed in percent), not variance (see, for example, Allen, Einchomb, and Granger (2006)). The payoff to a variance swap is the difference between the ex post realized variance and the squared strike price multiplied by the so-called variance notional, which represents the profit or loss per point difference between realized and implied variance. Since market participants often think in terms of volatility, the vega notional is often used instead of the variance notional.

The vega notional shows the profit or loss from a 1% change in volatility and it is calculated as $N_{vega} = N_{var} \times 2K$, where K is the strike of the variance swap expressed in terms of volatility.[15]

Margin requirements in general depend on the type of strategy employed. Variance swaps are margined in a similar way to options. In the following, we assume that the required margin is nine times the vega notional. For our sample period, the assumed margin is sufficient to withstand a daily adverse move in volatility in excess of three standard deviations based on monthly data. This means a trading portfolio that can be re-balanced daily will never be wiped out in a single day.[16] We further assume that during the life of the swap, the variation margin as well as the initial margin are set in the same fashion. For example, the variation margin for a short position with notional value of one dollar at time t is set as

$$PV_t\left(T\right) \times \left[\left\{\frac{t}{T} \times \widetilde{RV}_{0,t} + \frac{T-t}{T} \times \left(\text{Implied Vol}\left(t,T\right) + 9\right)^2\right\} - K^2\right], \tag{3.8}$$

where K is the fair strike at initiation and $PV_t\left(T\right)$ is the t-present value of one dollar at time T. The minimum required margin is then the maximum of the initial margin and the variation margin at each point in time t.[17]

Margins influence our strategies along two dimensions: i) they limit the number of contracts that an investor can write (execution) and ii) they may force the investor to close down positions and take losses (profitability). To evaluate these effects, we assume in line with Santa-Clara and Saretto (2009) a zero-cost strategy. In particular, at the beginning of each month the investor borrows one dollar and allocates that amount to a risk-free rate account which she can use to cover the margin requirements. Then, the investor takes a short position in a variance swap contract for a notional amount which is equivalent to a fraction

[15]Thus, a vega notional of one dollar means a trader with a short position in a variance swap with a strike of 10% will lose one dollar if the volatility increases to 11%.

[16]Allen, Einchomb, and Granger (2006) provide an example of a term sheet with a collateral requirement equal to three times the vega notional, which is not sufficient to prevent a complete loss in a single trading day. In practice, payoffs to variance swaps are often capped. As a result, a three notional vega margin may be enough. However, this further complicates hedging and pricing (see also Martin (2013)). To abstract from the pricing issues, we thus impose a higher margin.

[17]As the month progresses, Equation (3.8) requires the calculation of the variance swap rate with a decreasing time to maturity. To ensure that our results are not biased by interpolation, we enter into a position when the options have exactly one month to maturity instead of sampling at the end of the month.

of that one dollar. This quantity is referred to as the target notional, and the corresponding vega amount of the contract is referred to as the target vega notional. The initial margin requirement is then approximately equal to nine times the target vega notional.

During the month, we assume that the investor cannot borrow additional capital. In other words, margin calls are met by liquidating the investment in the risk-free account. When the risk-free account is no longer sufficient to meet the margin call, then the position is liquidated at the swap value. The investor is then allowed to open a new position such that the new margin does not exceed 90% of the available wealth. The maturity date of the new contract remains the same as before, but a new strike is defined for the value of the swap to be zero. Moreover, the (vega) notional of the new contract is adjusted accordingly. Hence, unless no re-balancing occurs during the month, the effective (vega) notional will differ from the target (vega) notional. At the end of the month, the variance swap position is closed and the proceeds are added to the risk-free account. Together with the interest earned on the risk-free account (which in general is negligible) this allows to calculate both the P/L of the strategy as well as the return in month t compared to the initial position of one dollar.

Given the relatively high margin requirements, the maximum target vega notional possible is 10 cents on the dollar or 10%. Table 3.7 reports results using target vega notionals ranging between 1% and 10%, meaning that the initial margin ranges from just under 10% to 90% of the one dollar borrowed at the beginning of each month. For comparison, we also calculate un-margined returns based on the same zero-cost strategy described above but without applying the variational margin. This means that the initial position is never forcibly closed out and remains the same until the end of the month. Note that the returns in Table 3.7 are calculated for a short position in a variance swap and they are, thus, positive on average.[18] We report bootstrapped confidence intervals for the quantities of interest.

First, the results in Table 3.7 show that margined returns are increasing in the target vega notional. This is not surprising as shorting variance is profitable and a higher target vega notional is tantamount to increasing the exposure to variance for a given amount of

[18]Moreover, note that the returns in Table 3.7 are defined slightly differently than those presented in Table 3.3 and, hence, they cannot be directly compared.

Table 3.7
Treasury Variance Swaps and Margin Requirements

This table analyzes the impact of margin requirements on the returns to writing one-month Treasury variance swaps. The table reports the effective vega notional, monthly average and monthly Sharpe ratio (SR) of margined returns at different levels of target vega notional in addition to the corresponding un-margined returns. The proportion of months with forced rescaling (out of 264) is reported in the third column of the table. 95% confidence intervals (in brackets) are obtained using a bootstrap with 10,000 draws. Data is monthly and runs from June 1990 to May 2012.

	margined				un-margined	
target vega	effective vega notional	rescaled month (fraction)	mean	SR	mean	SR
Panel A: 30y Treasury						
0.010	0.010	0.00	0.01	0.33	0.01	0.33
	[0.010,0.010]	[0.00,0.00]	[0.00,0.01]	[0.18,0.50]	[0.00,0.01]	[0.19,0.50]
0.050	0.050	0.01	0.03	0.32	0.03	0.33
	[0.050,0.050]	[0.00,0.02]	[0.02,0.05]	[0.18,0.50]	[0.02,0.05]	[0.19,0.50]
0.075	0.068	0.52	0.05	0.29	0.05	0.33
	[0.067,0.069]	[0.46,0.58]	[0.03,0.07]	[0.15,0.47]	[0.03,0.07]	[0.19,0.50]
0.100	0.067	1.00	0.05	0.30	0.07	0.33
	[0.066,0.068]	[1.00,1.00]	[0.03,0.06]	[0.15,0.47]	[0.04,0.09]	[0.19,0.50]
Panel B: 10y Treasury						
0.010	0.010	0.00	0.01	0.53	0.01	0.53
	[0.010,0.010]	[0.00,0.00]	[0.01,0.01]	[0.37,0.72]	[0.01,0.01]	[0.37,0.72]
0.050	0.050	0.00	0.04	0.52	0.04	0.53
	[0.050,0.050]	[0.00,0.03]	[0.03,0.04]	[0.37,0.71]	[0.03,0.04]	[0.37,0.72]
0.075	0.060	0.93	0.04	0.49	0.05	0.53
	[0.059,0.061]	[0.90,0.96]	[0.03,0.05]	[0.34,0.68]	[0.04,0.07]	[0.37,0.72]
0.100	0.060	1.00	0.04	0.50	0.07	0.53
	[0.059,0.060]	[1.00,1.00]	[0.03,0.05]	[0.34,0.68]	[0.06,0.09]	[0.37,0.72]
Panel C: 5y Treasury						
0.010	0.010	0.00	0.00	0.43	0.00	0.43
	[0.010,0.010]	[0.00,0.00]	[0.00,0.01]	[0.28,0.62]	[0.00,0.01]	[0.28,0.62]
0.050	0.048	0.25	0.02	0.44	0.02	0.43
	[0.047,0.048]	[0.20,0.31]	[0.01,0.03]	[0.29,0.63]	[0.01,0.03]	[0.28,0.62]
0.075	0.050	1.00	0.02	0.43	0.03	0.43
	[0.049,0.051]	[1.00,1.00]	[0.02,0.03]	[0.28,0.62]	[0.02,0.04]	[0.28,0.62]
0.100	0.050	1.00	0.02	0.43	0.04	0.43
	[0.049,0.051]	[1.00,1.00]	[0.02,0.03]	[0.28,0.62]	[0.03,0.05]	[0.28,0.62]

capital. At the same time, a higher vega notional also leads to a higher probability of forced rescaling if the volatility increases too much and a margin call cannot be met with the available capital. For example a 10% vega notional leads to a forced rescaling in every month of our sample and for all underlying tenors, while a 1% target vega notional never leads to a forced rescaling.

The relationship between the target vega notional and the volatility of our strategy is less straightforward. Absent rescaling, the volatility of the strategy is increasing in the vega notional. However, once the margin requirements take effect, forced rescaling often leads to a reduction in the volatility of the trading strategy as positions are often scaled down during periods of turmoil. This has an overall positive effect on Sharpe ratios that is at times, however, offset by a reduction in the mean return. Hence, Sharpe ratios are not necessarily monotonic in the target vega notional. However, they remain on average very sizable and range between 1 and 1.8 (annualized) depending on the underlying and the target vega notional.

As mentioned above, the assumed margin requirements are rather restrictive. This ensures that the investor cannot loose the total available capital in a daily move and before rescaling is allowed, which makes our strategy viable in the long-run. Applying for example the same setup to variance swaps on S&P500 futures results in the investor losing all the capital before being able to rescale a total of six times during the sample period. Hence, the strategy with variance swaps on equity is much riskier compared to using Treasury variance swaps and thus, in practice, margins would have to be higher in order to ensure that the portfolio can withstand even periods of severe turmoil in financial markets.[19]

Figure 3.6 plots in log scale the accumulation of wealth for different target vega notional short positions in variance swaps on 30y Treasury futures. The starting wealth is one dollar and it is assumed that the full amount is reinvested each month. As benchmark, we plot the total return index from investing one dollar in the S&P500 stock index. The figure confirms the results from Table 3.7, namely that that shorting variance in the Treasury market remains a profitable strategy when taking margins into account. Measured over 22 years, the Treasury variance swap strategies with vega notionals ranging between 5% and

[19] As mentioned earlier, one can alternatively impose a cap on the payoff as is regularly done in practice.

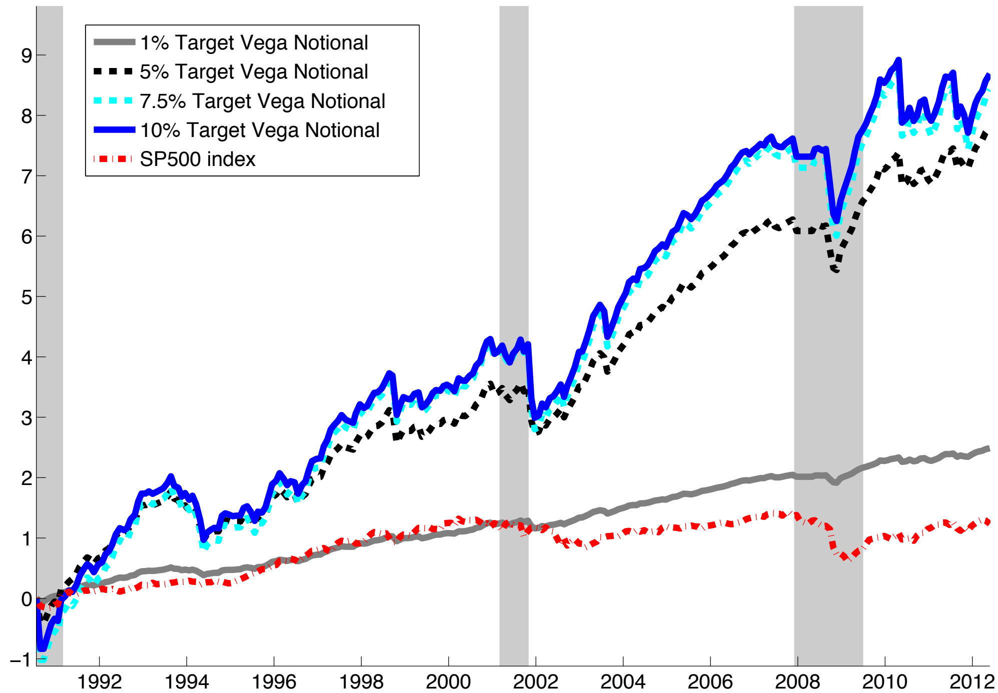

Figure 3.6. Margined Variance Swap Returns for 30y Treasury

This figure plots the wealth evolution in log scale for portfolios with an initial value of one USD. Each month the investor takes a short position in a 30y generalized Treasury variance swaps with a given target vega notional (ranging from 1% to 10%) and a maturity of one month. The required margin is nine times the vega notional. If losses within the month exceed the margin, the investor has to rescale the position. For comparison purposes, we also plot the evolution (also in log scale) of an investment in the S&P500 market index. Gray bars indicate NBER recessions. Data is monthly and runs from June 1990 to May 2012.

10% deliver on average around 40% return per year, compared with just above 6% for the S&P500 index. Moreover, the results also highlight that with proper margin requirements, the strategy remains very attractive even in a setting where investors do not have unlimited access to capital and may need to close out or scale down their positions during market turmoil.

3.3.2 The Impact of Realized Variance

In this section we gauge to what extent it makes a difference whether we use realized variance as defined in Equation (3.4) or whether we follow common practice and use daily squared log

Table 3.8
Summary Statistics Variance Swap Different RV

This table reports monthly summary statistics, mean, standard deviation, kurtosis, skewness, and the Sharpe ratio (SR) for one-month variance swap returns that are based on different measures of realized variance. GTVS corresponds to the variance swap payoff using realized variance as defined in Equation (3.4) that allows for perfect replication. LTVS corresponds to a variance swap that uses realized variance calculated using daily squared log returns. Data is monthly and runs from June 1990 to May 2012.

	30y Treasury		10y Treasury		5y Treasury	
	GTVS	LTVS	GTVS	LTVS	GTVS	LTVS
mean	-0.212	-0.212	-0.276	-0.276	-0.187	-0.187
t-stat	(-8.98)	(-8.98)	(-12.71)	(-12.71)	(-6.61)	(-6.61)
std	0.383	0.383	0.353	0.353	0.460	0.460
kurtosis	10.635	10.824	10.929	11.163	10.686	10.896
skewness	2.116	2.131	1.882	1.901	2.178	2.197
SR	-0.553	-0.552	-0.782	-0.782	-0.407	-0.406

returns. The results are reported in Table 3.8. The summary statistics reveal that *on average* it does not matter whether one uses one or the other measure of realized variance. This is intuitive, as for the data and sample period we study, the number and size of positive and negative jumps is roughly the same. The distinction between the two approaches becomes more evident, however, once we consider the time series of the differences between a variance swap that is defined using daily squared log returns and of our GTVS that uses realized variance as defined in Equation (3.4). The time series of the differences is depicted in Figure 3.7.

We note that there are both positive and negative differences in line with the notion that there are positive and negative jumps. The largest positive spikes correspond to the same distinct spikes that we observe in Figure 3.4. The three largest spikes are: July 2003 which corresponds to the month with the largest mortgage refinancing activity, August 1990 which was when Iraq invaded Kuwait, and September 2008 right after the Lehman default. Positive spikes correspond to negative jumps. In Appendix 9 we derive the difference between the realized variance measure that we use for our variance swap contract and realized variance calculated using daily squared log returns. It is straightforward to show

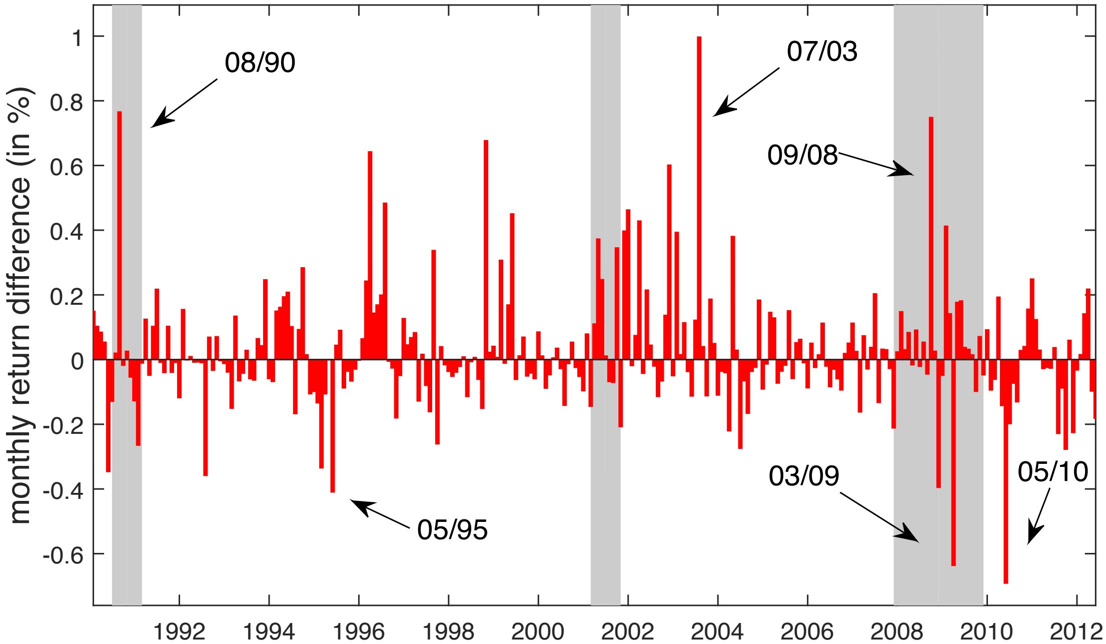

Figure 3.7. Difference in 30y Variance Swap Returns between $RV^{\log}$ and $\widetilde{RV}$

This figure plots the differences in monthly returns on 30y variance swap with a maturity of one month, but from two different measures of realized variance, log squared returns ($RV_T^{\log}$) and realized variance $\widetilde{RV}_T$ defined in Equation (3.4). Each return is computed as the payoff of the one-month variance swap (implied variance minus $RV_T^{\log}$ or $\widetilde{RV}_T$ respectively) scaled by the fair strike of the variance swap (implied variance). Gray bars indicate NBER recessions. Data is monthly and runs from January 1990 to May 2012.

that $\widetilde{RV}_{t,T} = RV_{t,T}^{\log} -$ cubed returns. Hence, large negative jumps render the log realized variance measure larger than $\widetilde{RV}_{t,T}$. Since the payoff is the realized leg minus the fair strike, the payoff becomes larger (smaller) compared to using $\widetilde{RV}_{t,T}$ in the presence of negative (positive) jumps.

3.3.3 Treasury Implied Volatility and Economic Activity

The VIX is often referred to as a fear gauge and, hence, it seems natural to ask whether the VIX has any predictive power for future economic activity.[20] In the following, we study whether Treasury implied variance and the slope of the implied variance term structure

[20] For example Bekaert and Hoerova (2014) present empirical evidence that the VIX^2 is an excellent predictor of economic activity at the monthly, quarterly, and annual frequency.

has any predictive power for economic activity as captured by the Chicago Fed National Activity Index (CFNAI).[21] We also check how implied volatility and the slope of implied variances is related to a measure of economic stress, namely the St. Louis Fed Stress Index (STLFSI).[22]

Even though we have implied variance available for different tenors, we focus on the longest maturity, i.e., the 30y Treasury futures options. To simplify notation, we call implied volatility from options on 30y Treasury futures, TIV (Treasury Implied Volatility, the implied variance is then TIV^2). We denote $\text{slope}^{\text{TIV}}$ the slope of the term structure of implied variances which is defined as the difference between the implied variance from a one-year and one-month option.

Figure 3.8 plots $\text{slope}^{\text{TIV}}$ together with the CFNAI, and the STLFSI multiplied by minus one. The co-movement between the three time series is strikingly high, especially during the recent crisis period. Indeed, the unconditional correlation between the slope and the CFNAI is 68% and it is 76% with the STLFSI. To test the relationship more formally, we run predictive regressions from the economic activity/stress index onto TIV^2, $\text{slope}^{\text{TIV}}$, VIX^2, the slope of the implied variance on the S&P500 ($\text{slope}^{\text{S\&P500}}$), and the slope of the term structure at different horizons, n, ranging from zero to twelve months:

$$\text{CFNAI}_{t+n}/\text{STLFSI}_{t+n} = \beta_n^{\text{TIV}}\text{TIV}_t^2 + \beta_n^{\text{slope TIV}}\text{slope}_t^{\text{TIV}} + \ldots + \epsilon_{t+n}.$$

The results are reported in Tables 3.9 and 3.10.[23] TIV^2 is an excellent predictor of future economic activity up to eight months: For the contemporaneous regressions, we find that for any one standard deviation shock in the TIV^2, there is a 0.3 standard deviation shock to economic activity. t-statistics range from -6.45 to -3.01 for the eight month horizon. The R^2s drop fast with the horizon: For contemporaneous regressions, the R^2 is 29% but it drops by half, i.e., to 14% after four months. Similarly to the level, the slope of the TIV has

[21] The CFNAI is available on the Chicago Fed web page as a weighted average of 85 existing monthly indicators of national economic activity including (i) production and income, (ii) employment, unemployment, and hours, (iii) personal consumption and housing, and (iv) sales, orders, and inventories. A positive index corresponds to growth above trend and a negative index corresponds to growth below trend.

[22] The STLFSI is available from FRED and is a principal component from 18 different time-series including (i) interest rates, (ii) yield spreads, and (iii) other indicators which include variables like the VIX. A positive (negative) index corresponds to a higher (lower) degree of stress compared to the trend.

[23] Variables are standardized, meaning, we de-mean and divide by the standard deviation.

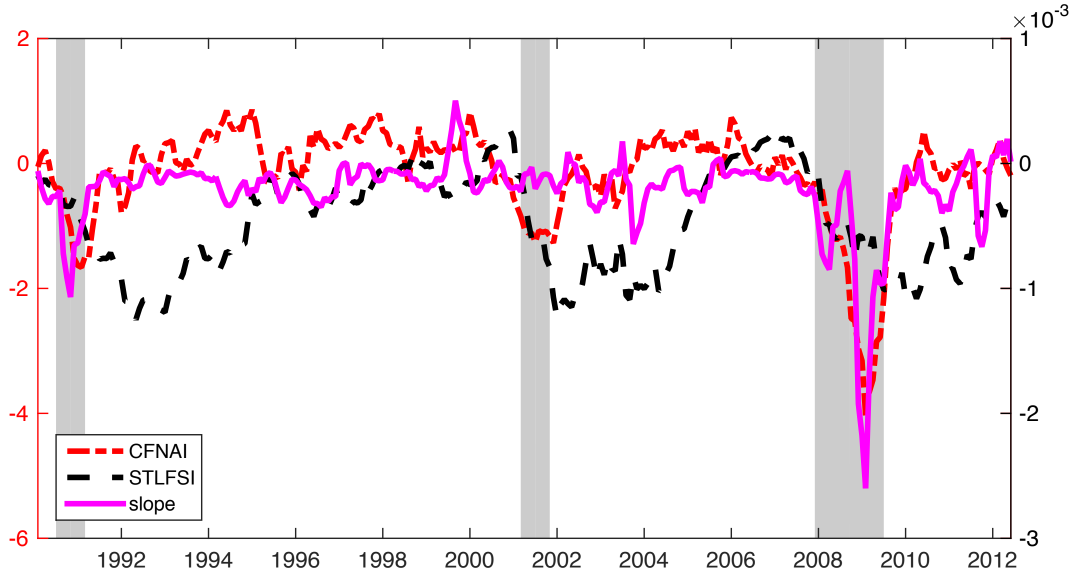

Figure 3.8. Slope of 30y Implied Variance, Chicago Fed National Activity Index (CFNAI), & St. Louis Fed Stress Index (SLFSI)

This figure plots the slope of the implied variance term structure which is defined as the difference between the one-year and one-month implied variance together with the Chicago Fed National Activity Index and the (negative of the) St. Louis Fed Stress Index. The variables are de-meaned and standardized. Gray bars indicate NBER recessions. Data is monthly and runs from January 1990 to May 2012.

strong predictive power up to eight months with R^2s monotonically decreasing from 20% for contemporaneous regressions to virtually zero for horizons longer than eight months. When we include both the level and the slope in the regression, the significance of both factors is not affected.

Adding VIX^2 or the slope of the VIX term structure does not alter the results much: The predictive power of the slope of the TIV is quantitatively the same, however, some of the predictive power of TIV^2 is subsumed by VIX^2. Adding the slope of the Treasury yield term structure itself does not change the results either. The slope of the term structure is known to be a long-term predictor of the business cycle rather than a predictor of short-term fluctuations. Hence, at horizons up to eight months, we observe basically no predictive power from the slope of the term structure, whereas at longer horizons, where the other right-hand side variables lose any power, the slope of the term structure becomes significant.

Table 3.9
Predictive Regressions Economic Activity

This table reports estimated coefficients from predictive regressions from the Chicago Fed National Activity Index (CFNAI) onto the Treasury implied volatility index squared (TIV^2), the slope of the TIV, and other factors:

$$\text{CFNAI}_{t+n} = \beta_n^{\text{TIV}}\text{TIV}_t^2 + \beta_n^{\text{slope TIV}}\text{slope}_t^{\text{TIV}} + \ldots + \epsilon_{t+n}, \quad n = 0, \ldots, 12.$$

Variables are standardized, meaning we de-mean and divide by the standard deviation. t-Statistics presented in parentheses are calculated using Newey and West (1987). Regressions are run contemporaneously and for forecast horizons up to twelve months. Data is monthly and runs from June 1990 to May 2012.

horizon	contemp.	3m	6m	9m	12m
TIV^2	-0.302	-0.226	-0.121	-0.051	-0.038
	(-6.45)	(-5.24)	(-3.30)	(-2.51)	(-1.34)
Adj. R^2	28.89%	16.20%	4.64%	0.81%	0.47%
slope TIV	0.622	0.552	0.327	0.152	0.121
	(3.11)	(4.49)	(3.61)	(1.72)	(0.99)
Adj. R^2	19.71%	15.46%	5.32%	1.16%	0.74%
TIV^2	-0.237	-0.154	-0.074	-0.026	-0.018
	(-6.05)	(-4.21)	(-1.80)	(-0.69)	(-0.38)
slope TIV	0.330	0.360	0.232	0.118	0.098
	(3.74)	(3.47)	(2.23)	(0.91)	(0.59)
Adj. R^2	32.84%	20.83%	6.28%	0.94%	0.43%
TIV^2	-0.172	-0.022	0.019	0.004	0.010
	(-3.43)	(-0.38)	(0.31)	(0.07)	(0.13)
slope TIV	0.339	0.375	0.261	0.183	0.179
	(3.01)	(3.06)	(2.01)	(1.12)	(0.84)
VIX^2	-0.064	-0.131	-0.097	-0.046	-0.046
	(-1.22)	(-2.08)	(-1.58)	(-0.79)	(-0.72)
slope VIX	-0.037	-0.073	-0.101	-0.173	-0.210
	(-0.32)	(-0.67)	(-0.71)	(-0.95)	(-1.09)
Adj. R^2	33.52%	25.20%	7.66%	1.23%	1.37%
TIV^2	-0.188	-0.015	0.039	0.038	0.054
	(-3.82)	(-0.26)	(0.69)	(0.63)	(0.71)
slope TIV	0.309	0.388	0.298	0.243	0.261
	(2.90)	(3.11)	(2.24)	(1.41)	(1.19)
VIX^2	-0.053	-0.136	-0.112	-0.072	-0.081
	(-1.03)	(-2.16)	(-1.81)	(-1.17)	(-1.14)
slope VIX	-0.014	-0.084	-0.135	-0.232	-0.288
	(-0.13)	(-0.78)	(-0.94)	(-1.23)	(-1.39)
slope yields	-0.099	0.044	0.124	0.204	0.268
	(-1.73)	(0.72)	(1.65)	(2.18)	(2.38)
Adj. R^2	34.34%	25.12%	8.96%	5.40%	9.14%

Table 3.10
Predictive Regressions Stress Index

This table reports estimated coefficients from predictive regressions from the St. Louis Fed Stress Index (STLFSI) onto the Treasury implied volatility index squared (TIV^2) and the slope of the TIV:

$$\text{STLFSI}_{t+n} = \beta_n^{\text{TIV}} \text{TIV}_t^2 + \beta_n^{\text{slope TIV}} \text{slope}_t^{\text{TIV}} + \epsilon_{t+n}, \quad n = 0, \ldots, 12.$$

Variables are standardized, meaning we de-mean and divide by the standard deviation. t-Statistics presented in parentheses are calculated using Newey and West (1987). Regressions are run contemporaneously and for forecast horizons up to twelve months. Data is monthly and runs from January 1994 to May 2012.

horizon	contemp.	3m	6m	9m	12m
TIV^2	0.276	0.192	0.088	0.025	0.006
	(2.98)	(2.43)	(1.53)	(0.71)	(0.27)
Adj. R^2	28.74%	13.99%	2.93%	0.24%	0.02%
slope TIV	-0.791	-0.600	-0.356	-0.285	-0.220
	(-4.14)	(-4.12)	(-3.69)	(-2.21)	(-1.47)
Adj. R^2	35.37%	20.14%	6.95%	4.43%	2.55%
TIV^2	0.160	0.098	0.022	-0.048	-0.055
	(2.03)	(1.42)	(0.41)	(-0.98)	(-1.06)
slope TIV	-0.576	-0.466	-0.325	-0.354	-0.299
	(-3.89)	(-3.22)	(-2.47)	(-1.98)	(-1.65)
Adj. R^2	42.18%	22.44%	6.65%	4.60%	2.86%

We run the same exercise but now use the stress index as the left-hand side variable.[24] The results are reported in Table 3.10. The results depict a similar picture as for the activity index but with opposite signs: An increase in TIV^2 implies an increase in the stress index. Estimated coefficients are significant up to four months. The estimated slope coefficients of $\mathrm{slope}^{\mathrm{TIV}}$ are negative and highly significant up to a horizon of eight months. The effects are not only statistically but also economically significant: For any one standard deviation change in the slope, there is almost a 0.8 standard deviation change in the stress index.

We conclude that both TIV^2 and the slope of the TIV are excellent predictors of future economic activity or stress. The predictive power is only at the short horizon and dies out after eight months. One might now obviously suspect that our results are mainly driven by the large co-movement between all series during the summer of 2008. We therefore run robustness checks and study the predictive power using a sample which ends in August 2008. Indeed, the TIV^2 has no predictive power for CFNAI if we end the sample in 2008, however, the slope of the TIV is still significant but only at short horizons. For the stress index, we still find that both TIV^2 and the slope are highly significant. To save space, these results are deferred to the Online Appendix.

3.4 Conclusions

This paper studies the returns of variance trading strategies in the Treasury options market. We first derive theoretically how to replicate variance swaps. The fair strike can easily be constructed using a continuum of put and call options. The way we formalize the realized leg is critical: Instead of using squared log returns, we use simple returns which allows for a perfect replication of the variance swap payout even in the presence of jumps and regardless of the sampling partition.

Using a large panel data set of Treasury options with different tenors, we juxtapose the variance swap with delta-hedged ATM straddle strategies—an alternative strategy well known to be very profitable in the equity markets. Trading variance through straddles is attractive but the profitability of these strategies is dwarfed by the variance swaps: The

[24]We do not run a regression which includes the slope from yields, VIX^2, or the slope of the VIX as the stress index contains both the slope of the term structure and the VIX itself.

average excess return on a variance swap is about 20% per month regardless of the tenor. Moreover, these returns come at reasonable risk as the annualized Sharpe ratio is above two. We then examine how and whether these returns can be explained by standard risk factors such as the market, book-to-market, size, momentum, or liquidity factors. We find that none of these factors have significant explanatory power for the returns of variance swaps. The alpha of the strategy is large and highly significant. Using realistic assumptions on margins, we verify that variance trading in fixed income markets yields substantial profits even when investors cannot rely on an unlimited supply of capital to implement their strategies.

We also study the term structure of implied variances and variance risk premiums. We document that variance risk premiums are negative and their term structure is downward sloping in absolute terms. At longer maturities, the premiums are smaller in magnitude. The term structure of implied variances is also downward sloping on average while over time, it is highly time-varying and strongly pro-cyclical. In particular, the slightly negative slope becomes extremely negative during crisis periods. Using this observation, we find that both the level of implied volatility and the slope are excellent predictors of both economic activity and stress, especially at short horizons.

Appendix 1 Proof of Proposition 1

Lemma 1. *Suppose that $F_t(T,\mathbb{T})$ is a diffusion process*

$$\frac{dF_t(T,\mathbb{T})}{F_t(T,\mathbb{T})} = v_t' dB_t^{\mathbb{Q}_T}$$

Then,

$$E_t^{\mathbb{Q}_T}\left[\int_t^T v_s v_s' ds\right] = \frac{2}{P_{t,T}}\left[\int_0^{F_t(T,\mathbb{T})} \frac{Put_t(K,T,\mathbb{T})}{K^2} dK + \int_{F_t(T,\mathbb{T})}^{\infty} \frac{Call_t(K,T,\mathbb{T})}{K^2} dK\right]$$

Proof. See Appendix 8. □

Lemma 2. *Suppose that the numeraire of $\mathbb{Q}_T$ forward measure, $P_{t,T}$, has the following dynamics under $\mathbb{Q}$*

$$dP_{t,T} = P_{t,T}\left(r_t dt + \sigma_{t,T}' dB_t^{\mathbb{Q}}\right) \tag{A-1}$$

Then,

$$dB_t^{\mathbb{Q}_T} = -\sigma_{t,T} dt + dB_t^{\mathbb{Q}}$$

Proof. See for example Björk (2009). □

Proof of Proposition 1. Under the two assumptions that (i) the state follows affine diffusion and (ii) the short rate is an affine function of the state, Duffie and Kan (1996) show that $P_{t,T}$ is exponentially affine in the state.

$$P_{t,T} = \exp\left[\bar{A}\left(\Theta^{\mathbb{Q}}, t, T\right) + \bar{B}\left(\Theta^{\mathbb{Q}}, t, T\right)' Z_t\right] \tag{A-2}$$

Since no-arbitrage condition implies $F_t = P_{t,\mathbb{T}}/P_{t,T}$, the forward price is also exponentially affine in the state

$$F_t(T,\mathbb{T}) = P_{t,\mathbb{T}}/P_{t,T} = \exp\left[\psi_0 + \psi_1' Z_t\right]$$

where ψ_0 and ψ_1 are functions of $\{\Theta^{\mathbb{Q}}, T, \mathbb{T}\}$. Applying Ito's lemma to F_t under $\mathbb{Q}_T$, we have the following diffusion process

$$\frac{dF_t(T,\mathbb{T})}{F_t(T,\mathbb{T})} \equiv v_t' dB_t^{\mathbb{Q}_T} = \psi_1' \Sigma_{Z,t} dB_t^{\mathbb{Q}_T}$$

The expected quadratic variation of the forward up to time T is

$$E_t^{\mathbb{Q}_T}\left[\int_t^T v_s' v_s ds\right] = \int_t^T \psi_1' \Sigma_{Z0} \psi_1 ds + \sum_{i=1}^{m} \int_t^T \psi_1' \Sigma_{Zi} \psi_1 E_t^{\mathbb{Q}_T}[V_{is}] ds \tag{A-3}$$

The proof is completed by Lemma 1 once the right-hand side of equation (A-3) is an affine function of V_t only. Intuitively, the Gaussian factor X_t cannot affect $E_t^{\mathbb{Q}_T}[V_{is}]$, otherwise the expected value of V_s under $\mathbb{Q}_T$ could have a negative value. As a last step, apply Ito's lemma to equation (A-2). Then, the diffusion term of it, the $\sigma_{t,T}$ term in case of equation (A-1), is given by

$$\sigma_{t,T} = \Sigma_{Z,t} \bar{B}\left(\Theta^{\mathbb{Q}}, t, T\right)$$

As a consequence of Lemma 2, the dynamics of the state under $\mathbb{Q}_T$ can be written as

$$dZ_t = d\left(X_t', V_t'\right)' = \left[\left(\mu_{X,t}^{\mathbb{Q}\prime}, \mu_{V,t}^{\mathbb{Q}\prime}\right)' + \Sigma_{Z,t}\Sigma_{Z,t}'\bar{B}\left(\Theta^{\mathbb{Q}}, t, T\right)\right] dt + \Sigma_{Z,t} dB^{\mathbb{Q}_T}$$

Since $\mu_{V,t}^{\mathbb{Q}}$ is assumed to be affine in V_t only (i.e. as in Duffie, Filipović, and Schachermayer (2003), K_{1VX} in equation (1.1) is set to be zero for admissibility) and $\Sigma_{Z,t}\Sigma_{Z,t}'\bar{B}\left(\Theta^{\mathbb{Q}}, t, T\right)$ is a linear function of V_t solely, $\mu_{V,t}^{\mathbb{Q}_T}$ is also an affine function of V_t. $\square$

Appendix 2 The bond VIX^{2}s in a long-run risk framework

I initially solve the model of Bansal and Shaliastovich (2013) in the continuous-time framework, and then demonstrate the linear mapping between the bond VIXs and the two macroeconomic uncertainties in the model. The dynamics of consumption C_t, inflation π_t ,and their long-run risks as well as quantity of risk are specified as

$$\begin{aligned}
\frac{dC_t}{C_t} &= \left(\mu_c + X_{ct}\right) dt + \sigma_c dZ_{1t} \\
\frac{d\pi_t}{\pi_t} &= \left(\mu_\pi + X_{\pi t}\right) dt + \sigma_\pi dZ_{2t} \\
dX_{ct} &= \left(-\rho_c X_{ct} - \rho_{c\pi} X_{\pi t}\right) dt + \sqrt{V_{ct}} dW_{1t} \\
dX_{\pi t} &= -\rho_\pi X_{\pi t} dt + \sqrt{V_{\pi t}} dW_{2t} \\
V_{ct} &= \kappa_c \left(\bar{V}_c - V_{ct}\right) dt + w_c \sqrt{V_{ct}} dB_{1t} \\
V_{\pi t} &= \kappa_\pi \left(\bar{V}_\pi - V_{\pi t}\right) dt + w_\pi \sqrt{V_{\pi t}} dB_{2t}
\end{aligned}$$

where $Z_{1t}, Z_{2t}, W_{1t}, W_{2t}, B_{1t}$ and B_{2t} are independent Brownian motions. Following Duffie and Epstein (1992), the representative agent's objective is

$$J_t = \max_{\{C_s\}} E_t \left[\int_t^T f\left(C_s, J_s\right) ds\right]$$

where the normalized aggregator $f\left(C, J\right)$ is given by

$$f\left(C, J\right) = \beta \left(\frac{1-\gamma}{1-1/\psi}\right) J \left[\left(\frac{C}{\left(\left(1-\gamma\right) J\right)^{1/(1-\gamma)}}\right)^{1-1/\psi} - 1\right] \tag{A-4}$$

with β the rate of time preference, γ the relative risk version, and ψ the elasticity of intertemporal substitution. Conjecture J as

$$J\left(W_t, Z_t\right) = \exp\left(a_0 + a_1 X_{ct} + a_2 X_{\pi t} + a_3 V_{ct} + a_4 V_{\pi t}\right) \frac{W_t^{1-\gamma}}{1-\gamma}. \tag{A-5}$$

Since the envelop condition, $f_C = J_W$, can be written as

$$C = J_W^{-\psi} \left[\left(1-\gamma\right) J\right]^{\frac{1-\gamma\psi}{1-\gamma}} \beta^\psi, \tag{A-6}$$

substituting equation (A-5) into equation (A-6) enables us to express the log consumption-wealth ratio in terms of the state variables:

$$\log\left(\frac{C_t}{W_t}\right) = \psi\log\beta + \frac{1-\psi}{1-\gamma}\left(a_0 + a_1X_{ct} + a_2X_{\pi t} + a_3V_{ct} + a_4V_{\pi t}\right)$$

In addition, substituting equation (A-6) into equation (A-4) gives

$$f = \left(\frac{1-\gamma}{1-1/\psi}\right)\left(\frac{C_t}{W_t} - \beta\right)J$$

Applying log-linear approximation from Campbell (1993) to the consumption-wealth ratio

$$\frac{C_t}{W_t} \approx g_1 + g_1\log g_1 + g_1\log\left(\frac{C_t}{W_t}\right) \tag{A-7}$$

where g_1 is the long-term mean of the consumption-wealth ratio. Then

$$\begin{aligned} f &\approx \left(\frac{1-\gamma}{1-1/\psi}\right)\left[g_1 + g_1\log g_1 + g_1\log\left(\frac{C_t}{W_t}\right) - \beta\right]J \\ &= \left(\frac{1-\gamma}{1-1/\psi}\right)\left[\xi + \frac{g_1(1-\psi)}{1-\gamma}\left(a_0 + a_1X_{ct} + a_2X_{\pi t} + a_3V_{ct} + a_4V_{\pi t}\right)\right]J \end{aligned}$$

where $\xi = g_1 + g_1\log g_1 + g_1\psi\log\beta - \beta$. As shown in Duffie and Epstein (1992), the state price process is identified as $\zeta_t = \exp\left[\int_0^t f_J\left(C_s, J_s\right)ds\right] f_C\left(C_t, J_t\right)$ and the corresponding nominal pricing kernel is defined as $\tilde{\zeta}_t = \frac{\zeta_t}{\pi_t}$. Applying Ito's lemma for the nominal pricing kernel results in

$$\begin{aligned} \frac{d\tilde{\zeta}_t}{\tilde{\zeta}_t} &= -\left(r_0 + r_1X_{ct} + r_2X_{\pi t} + r_3V_{ct} + r_4V_{\pi t}\right)dt \\ &\quad -\lambda_{1t}dZ_{1t} - \lambda_{2t}dZ_{2t} - \lambda_{3t}dW_{1t} - \lambda_{4t}dW_{2t} - \lambda_{5t}dB_{1t} - \lambda_{6t}dB_{2t} \end{aligned}$$

where

$$\begin{aligned} \lambda_{1t} &= \gamma\sigma_c, & \lambda_{2t} &= \sigma_\pi \\ \lambda_{3t} &= -\frac{1-\gamma\psi}{1-\gamma}a_1\sqrt{V_{ct}}, & \lambda_{4t} &= -\frac{1-\gamma\psi}{1-\gamma}a_2\sqrt{V_{\pi t}} \\ \lambda_{5t} &= -\frac{1-\gamma\psi}{1-\gamma}a_3w_c\sqrt{V_{ct}}, & \lambda_{6t} &= -\frac{1-\gamma\psi}{1-\gamma}a_4w_\pi\sqrt{V_{\pi t}} \end{aligned}$$

and

$$
\begin{aligned}
r_0 &= -\xi_1 - \frac{1-\psi}{1-\gamma}\left(a_1\mu_\pi + a_3\kappa_c\bar{V}_c + a_4\kappa_\pi\bar{V}_\pi\right) + \gamma\mu_c - \frac{1}{2}\gamma\left(\gamma-1\right)\sigma_c^2 + \mu_\pi - \sigma_\pi^2 \\
r_1 &= \gamma + a_1\left(\rho_c + g_1\right)\frac{1-\gamma\psi}{1-\gamma} \\
r_2 &= 1 + \left(a_1\rho_{c\pi} + a_2\left(\rho_\pi + g_1\right)\right)\frac{1-\gamma\psi}{1-\gamma} \\
r_3 &= a_3\left(\kappa_c + g_1\right)\frac{1-\gamma\psi}{1-\gamma} - \frac{1}{2}\left(a_1^2 + a_3^2 w_c^2\right)\left(\frac{1-\gamma\psi}{1-\gamma}\right)^2 \\
r_4 &= a_4\left(\kappa_\pi + g_1\right)\frac{1-\gamma\psi}{1-\gamma} - \frac{1}{2}\left(a_2^2 + a_4^2 w_\pi^2\right)\left(\frac{1-\gamma\psi}{1-\gamma}\right)^2
\end{aligned}
$$

Because of the Girsanov theorem, the price of risk dynamics implies that the state variable also follows affine diffusion under $\mathbb{Q}$, and the short rate is affine in the state as can be seen.

Appendix 3 The bond VIX^{2}s and Gaussian quadratic term structure models

In Gaussian quadratic term structure models, the state variable X_t is assumed to follow the Ornstein-Uhlenbeck process under the risk-neutral measure $\mathbb{Q}$:

$$dX_t = \left[K_{0X}^{\mathbb{Q}} + K_{1X}^{\mathbb{Q}} X_t\right] dt + \sqrt{\Omega}\, dB_t^{\mathbb{Q}}$$

where $\sqrt{\Omega}$ represents the Cholesky decomposition of a positive definite matrix Ω. The short rate is a quadratic function of the Gaussian affine-diffusion state

$$r_t = \Psi_0 + \Psi_1 X_t + X_t' \Psi_2 X_t$$

Then, bond prices are represented as

$$P_{t,T} = \exp\left[\tilde{A}\left(\Theta^{\mathbb{Q}}, t, T\right) + X_t'\tilde{B}\left(\Theta^{\mathbb{Q}}, t, T\right) + X_t'\tilde{C}\left(\Theta^{\mathbb{Q}}, t, T\right) X_t\right]$$

with a symmetric matrix $\tilde{C}$ (see for example Ahn, Dittmar, and Gallant (2002)). Applying Ito's lemma, the dynamics of the forward are

$$\frac{dF_t\left(T, \mathbb{T}\right)}{F_t\left(T, \mathbb{T}\right)} \equiv v_t' dB_t^{\mathbb{Q}_T} = \left[\xi_0 + 2\xi_1 X_t\right]' \sqrt{\Omega}\, dB_t^{\mathbb{Q}_T}$$

where

$$
\begin{aligned}
\xi_0 &= \tilde{B}\left(\Theta^{\mathbb{Q}}, t, \mathbb{T}\right) - \tilde{B}\left(\Theta^{\mathbb{Q}}, t, T\right) \\
\xi_1 &= \tilde{C}\left(\Theta^{\mathbb{Q}}, t, \mathbb{T}\right) - \tilde{C}\left(\Theta^{\mathbb{Q}}, t, T\right)
\end{aligned}
$$

Then, the expected quadratic variation can be written as

$$E_t^{\mathbb{Q}_T}\left[\int_t^T v_s' v_s\, ds\right] = E_t^{\mathbb{Q}_T}\left[\int_t^T \left(\xi_0 + 2\xi_1 X_s\right)' \Omega \left(\xi_0 + 2\xi_1 X_s\right) ds\right]$$

Suppose that the i-th element of X_t has no impact on the volatilities in yields. This implies that i-th row and column of $\tilde{C}\left(\Theta^{\mathbb{Q}}, t, T\right)$ is zero for all $T > 0$. Since the instantaneous volatility of $P_{t,T}$ is

$$\sigma_{t,T} = \sqrt{\Omega}'\left[\tilde{B}\left(\Theta^{\mathbb{Q}}, t, T\right) + 2\tilde{C}\left(\Theta^{\mathbb{Q}}, t, T\right)X_t\right],$$

the dynamics of the state under $\mathbb{Q}_T$ is

$$dX_t = \left[K_{0X}^{\mathbb{Q}_T} + K_{1X}^{\mathbb{Q}_T}X_t\right]dt + \sqrt{\Omega}dB_t^{\mathbb{Q}_T}$$

where

$$\begin{aligned}
K_{0X}^{\mathbb{Q}_T} &= K_{0X}^{\mathbb{Q}} + \Omega\tilde{B}\left(\Theta^{\mathbb{Q}}, t, T\right) \\
K_{1X}^{\mathbb{Q}_T} &= K_{1X}^{\mathbb{Q}} + 2\Omega\tilde{C}\left(\Theta^{\mathbb{Q}}, t, T\right)
\end{aligned}$$

Appendix 4 Discrete-time term structure model with stochastic volatility

Appendix 4.1 Zero-coupon bonds' loading on pricing factors

Denote the price of zero-coupon bond with maturity of n as $P_t^{(n)}$. Then, we can show that $\ln P_t^{(n)} = -\bar{A}_n - \bar{B}_{V,n}V_t - \bar{B}_{X,n}X_t$ with loadings given by

$$\begin{aligned}
\bar{A}_n &= \delta_0 + \bar{A}_{n-1} + K_{0X}^{\mathbb{Q}\prime}\bar{B}_{X,n} + c^{\mathbb{Q}}\nu^{\mathbb{Q}\prime}\bar{B}_{V,n-1} - \frac{1}{2}\alpha_{n-1} \\
\bar{B}_{V,n} &= \delta_V + K_{1XV}^{\mathbb{Q}\prime}\bar{B}_{X,n-1} + \rho^{\mathbb{Q}\prime}\bar{B}_{V,n-1} - \frac{1}{2}\beta_{n-1} \qquad \text{(A-8)} \\
\bar{B}_{X,n} &= \delta_X + K_{1X}^{\mathbb{Q}\prime}\bar{B}_{X,n-1}
\end{aligned}$$

where

$$\begin{aligned}
\alpha_n &= \bar{B}_{X,n}'\Sigma_{0X}\Sigma_{0X}'\bar{B}_{X,n} - 2v_V^{\mathbb{Q}\prime}\left[\log\mathcal{J}_n - c_V^{\mathbb{Q}\prime}\bar{B}_{XV,n}\right] \qquad \text{(A-9)} \\
\beta_n &= -2\mathcal{G}_n + \left(\bar{B}_{X,n}'\Sigma_{1X}\Sigma_{1X}'\bar{B}_{X,n}, \cdots, \bar{B}_{X,n}'\Sigma_{mX}\Sigma_{mX}'\bar{B}_{X,n}\right)' \qquad \text{(A-10)}
\end{aligned}$$

and

$$\begin{aligned}
\mathcal{G}_{n-1} &= \rho^{\mathbb{Q}\prime}c^{\mathbb{Q}-1\prime}\left(\left[\operatorname{diag}\left(\mathcal{J}_{n-1}\right)\right]^{-1} - I_m\right)c^{\mathbb{Q}\prime}\bar{B}_{XV,n-1} \\
\mathcal{J}_{n-1} &= \iota_m + c^{\mathbb{Q}\prime}\bar{B}_{XV,n-1} \\
\bar{B}_{XV,n-1} &= \Sigma_{XV}'\bar{B}_{X,n-1} + \bar{B}_{V,n-1}
\end{aligned}$$

The initial condition of the difference equation is $\bar{A}_0 = \bar{B}_{V,0} = \bar{B}_{X,n} \equiv 0$.

In addition, denoting the yield on a zero-coupon bond of maturity n as $y_{n,t}$, then, we have

$$y_{n,t} = A_n + B_nZ_t$$

where $A_n = \frac{1}{n}\bar{A}_n$, $B_{V,n} = \frac{1}{n}\bar{B}_{V,n}$ and $B_{V,n} = \frac{1}{n}\bar{B}_{X,n}$. Stacked yields, y_t, can be represented as

$$y_t = A + B_VV_t + B_XX_t$$

where A, B_V and B_X are corresponding stacked A_n, $B_{V,n}$ and $B_{X,n}$. Furthermore it can be written as a function of the observable factor $\mathcal{Z}_t$ as in Section 1.5. With $W \in \mathbb{R}^{(N-m)\times J}$ being the weighting matrix to construct $\mathcal{P}_t$,

$$y_t = A + B_V V_t + B_X X_t = \mathcal{A} + \mathcal{B}_{\mathcal{P}}\mathcal{P}_t + \mathcal{B}_{\mathcal{V}}\mathcal{V}_t$$

where

$$\begin{aligned} \mathcal{B}_{\mathcal{P}} &= B_X \left(W B_X\right)^{-1} \\ \mathcal{B}_{\mathcal{V}} &= \left(I - \mathcal{B}_{\mathcal{P}} W\right) B_V \beta^{-1} \\ \mathcal{A} &= \left(I - \mathcal{B}_{\mathcal{P}} W\right) A - \mathcal{B}_{\mathcal{V}}\alpha \end{aligned} \tag{A-11}$$

Appendix 4.2 $\mathcal{V}_t$'s loading on the latent variance factor

As shown in Appendix 8,

$$\begin{aligned} \mathcal{V}_t &\equiv \frac{2}{P_{t,T}}\left[\int_0^{F_t(T,\mathbb{T})} \frac{\mathrm{Put}_t\left(K,T,\mathbb{T}\right)}{K^2}dK + \int_{F_t(T,\mathbb{T})}^{\infty} \frac{\mathrm{Call}_t\left(K,T,\mathbb{T}\right)}{K^2}dK\right] \\ &= 2\left[\ln E_t^{\mathbb{Q}_T}\left(F_T\left(T,\mathbb{T}\right)\right) - E_t^{\mathbb{Q}_T}\left(\ln F_T\left(T,\mathbb{T}\right)\right)\right] \end{aligned}$$

In the case of one-month TIVs or TYVIX, T is equal to $t+1$. Hence, the calculation for α and β should be taken under $\mathbb{Q}_{t+1}$ forward measure (i.e. the risk neutral measure). I use the notation $\mathbb{Q}$ instead of $\mathbb{Q}_{t+1}$ according to the convention in the literature, and denote the time to maturity of the underlying bond on the expiration date of the forward as $n = \mathbb{T} - (t+1)$ for notational simplicity. Then,

$$\begin{aligned} \ln E_t^{\mathbb{Q}}\left[F_{t+1}\left(t+1,\mathbb{T}\right)\right] &= \ln F_t\left(t+1,\mathbb{T}\right) = \ln P_t^{(n+1)} - \ln P_t^{(1)} = -\left(A_{n+1} - A_1\right) - \left(B'_{n+1} - B'_1\right) Z_t \\ E_t^{\mathbb{Q}}\left[\ln F_{t+1}\left(t+1,\mathbb{T}\right)\right] &= E_t^{\mathbb{Q}}\left(\ln P_{t+1}^{(n)}\right) = -A_n - B'_n E_t^{\mathbb{Q}}\left(Z_{t+1}\right) \\ &= -A_n - B'_{X,n}\left(K_{0X}^{\mathbb{Q}} + K_{1X}^{\mathbb{Q}} X_t + K_{1XV}^{\mathbb{Q}} V_t\right) - B'_{V,n}\left(c^{\mathbb{Q}}\nu^{\mathbb{Q}} + \rho^{\mathbb{Q}} V_t\right) \end{aligned}$$

Using the difference equations in (A-8), we have

$$\begin{aligned} \mathcal{V}_t^{(n)} &= 2\ln E_t^{\mathbb{Q}}\left[F_{t+1}\left(t+1,\mathbb{T}\right)\right] - 2E_t^{\mathbb{Q}}\left[\ln F_{t+1}\left(t+1,\mathbb{T}\right)\right] \\ &= \alpha_n + \beta'_n V_t \end{aligned}$$

where α_n and β_n are given in equations (A-9) and (A-10).

Appendix 5 Test for the USV effect

Denote R as the number of priced factors. For $J \geq R$, a $(J \times 1)$ vector of yields can be written as

$$y_t = A + B_X X_t + B_V V_t = \mathcal{A} + \mathcal{B}_{\mathcal{P}}\mathcal{P}_t + \mathcal{B}_{\mathcal{V}}\mathcal{V}_t$$

where $\mathcal{P}_t\left(\equiv W y_t\right)$ is the first $(R-m)$ principal components of y_t, $\mathcal{V}_t$ is the observable measure of variance risk, and

$$\mathcal{B}_{\mathcal{V}} = \left(I - B_X\left(W B_X\right)^{-1} W\right) B_V \beta^{-1}$$

as in equation (A-11). Suppose that, without loss of generality, the first volatility factor V_{1t} is an unspanned volatility. Then, there exists Φ such that $B_{V_1} = B_X\Phi$ (see, for example, Lemma 2 in Joslin (2015)), which implies

$$\begin{aligned}\mathcal{B}_{\mathcal{V}} &= \left(I - B_X\left(WB_X\right)^{-1}W\right)\left[\begin{array}{cc} B_X\Phi, & B_{V_{2:m}}\end{array}\right]\beta^{-1} \\ &= \left[\begin{array}{cc} 0, & \left(I - B_X\left(WB_X\right)^{-1}W\right)B_{V_{2:m}}\end{array}\right]\beta^{-1}\end{aligned}$$

Hence, $\mathcal{B}_{\mathcal{Z}} \equiv \left[\begin{array}{cc}\mathcal{B}_{\mathcal{P}}, & \mathcal{B}_{\mathcal{V}}\end{array}\right]$ cannot be a full rank matrix in the presence of USV.

Appendix 6 Returns on generalized mimicking portfolios

As in Joslin, Priebsch, and Singleton (2014), we have

$$\mathbb{E}_t^{\mathbb{P}}\left[P_{t+1}^{(n-1)}/P_t^{(n)}\right] = \exp\left[k_t^{\mathbb{P}}\left(\mathcal{Z}_{t+1}; -B_{n-1}\right) - k_t^{\mathbb{Q}}\left(\mathcal{Z}_{t+1}; -B_{n-1}\right) + r_t\right]$$

where $k_t\left(\mathcal{Z}_{tt+1}; u\right)$ denotes the conditional cumulant generating function of $\mathcal{Z}_{t+1}$ at time t, and B_n is the corresponding factor loading. In addition, the Laplace transform of the mixture of Gaussian and multivariate non-central gamma distribution is given by

$$\begin{aligned}\mathbb{E}\left[\exp\left(u'\mathcal{Z}_{t+1}\right)\right] &= \exp\left(u_{\mathcal{P}}'\left(K_{0\mathcal{P}} + K_{1\mathcal{P}}\mathcal{P}_t + K_{1\mathcal{P}\mathcal{V}}\mathcal{V}_t - \Sigma_{\mathcal{P}\mathcal{V}}\left[\mu + c\nu + \rho\left(\mathcal{V}_t - \mu\right)\right]\right) + \frac{1}{2}u_{\mathcal{P}}'\Sigma_{\mathcal{P},t}\Sigma_{\mathcal{P},t}'u_{\mathcal{P}}\right) \\ &\quad\times\exp\left(u_{\mathcal{P}\mathcal{V}}'\mu - \sum_{i=1}^{m} v_i\log\left(1 - e_i'c'u_{\mathcal{P}\mathcal{V}}\right) + \sum_{i=1}^{m}\frac{e_i'c'u_{\mathcal{P}\mathcal{V}}}{1 - e_i'c'u_{\mathcal{P}\mathcal{V}}}e_i'c^{-1}\rho\left(\mathcal{V}_t - \mu\right)\right)\end{aligned}$$

where $u_{\mathcal{P}\mathcal{V}} = \Sigma_{\mathcal{P}\mathcal{V}}'u_{\mathcal{P}} + u_{\mathcal{V}}$ and e_i is a zero vector except its i-th element is 1. μ denotes the lower bound of $\mathcal{V}_t$, and is equal to α_n in Appendix 4.2 when the lower bound on latent V_t is set to be zero. Under the normalization scheme of $c^{\mathbb{P}} = c^{\mathbb{Q}}$ and $\Sigma_{XV}^{\mathbb{P}} = \Sigma_{XV}^{\mathbb{Q}}$, we have

$$\begin{aligned}& k_t^{\mathbb{P}}\left(\mathcal{Z}_{t+1}; -B_{n-1}\right) - k_t^{\mathbb{Q}}\left(\mathcal{Z}_{t+1}; -B_{n-1}\right) \\ =\; & B_{\mathcal{P},n-1}'\left[\left(K_{0\mathcal{P}}^{\mathbb{Q}} - K_{0\mathcal{P}}^{\mathbb{P}}\right) + \left(K_{1\mathcal{P}}^{\mathbb{Q}} - K_{1\mathcal{P}}^{\mathbb{P}}\right)\mathcal{P}_t + \left(K_{1\mathcal{P}\mathcal{V}}^{\mathbb{Q}} - K_{1\mathcal{P}\mathcal{V}}^{\mathbb{P}}\right)\mathcal{V}_t\right] \\ & -B_{\mathcal{P},n-1}'\Sigma_{\mathcal{P}\mathcal{V}}\left[c\left(\nu^{\mathbb{Q}} - \nu^{\mathbb{P}}\right) + \left(\rho^{\mathbb{Q}} - \rho^{\mathbb{P}}\right)\left(\mathcal{V}_t - \mu\right)\right] \\ & +\sum_{i=1}^{m}\left(v_i^{\mathbb{Q}} - v_i^{\mathbb{P}}\right)\log\left(1 + \mathcal{A}_i\right) + \sum_{i=1}^{m}\frac{\mathcal{A}_i}{1 + \mathcal{A}_i}e_i'c^{-1}\left(\rho^{\mathbb{Q}} - \rho^{\mathbb{P}}\right)\left(\mathcal{V}_t - \mu\right)\end{aligned}$$

where the constant $\mathcal{A}_i$ is given by

$$\mathcal{A}_i = e_i'c'\left(\Sigma_{\mathcal{P}\mathcal{V}}'B_{\mathcal{P},n-1} + B_{\mathcal{V},n-1}\right)$$

Since $\frac{1}{1+x} \approx x$ and $\log(1+x) \approx x$ for small x, the above can be approximated by

$$\begin{aligned}
& B'_{\mathcal{P},n-1}\left[\left(K^{\mathbb{Q}}_{0\mathcal{P}} - K^{\mathbb{P}}_{0\mathcal{P}}\right) + \left(K^{\mathbb{Q}}_{1\mathcal{P}} - K^{\mathbb{P}}_{1\mathcal{P}}\right)\mathcal{P}_t\right] + B'_{\mathcal{P},n-1}\left(K^{\mathbb{Q}}_{1\mathcal{P}\mathcal{V}} - K^{\mathbb{P}}_{1\mathcal{P}\mathcal{V}}\right)\mathcal{V}_t \\
& + \left[B'_{\mathcal{P},n-1}\Sigma_{\mathcal{P}\mathcal{V}} - \sum_{i=1}^{m}\mathcal{A}_i e'_i c^{-1}\right]\left(\rho^{\mathbb{P}} - \rho^{\mathbb{Q}}\right)(\mathcal{V}_t - \mu) \\
& + B'_{\mathcal{P},n-1}\Sigma_{\mathcal{P}\mathcal{V}} c\left(\nu^{\mathbb{P}} - \nu^{\mathbb{Q}}\right) - \sum_{i=1}^{m}\mathcal{A}_i\left(v^{\mathbb{P}}_i - v^{\mathbb{Q}}_i\right)
\end{aligned}$$

When $m = 1$ and $\nu^{\mathbb{Q}}_i = \nu^{\mathbb{P}}_i$, this can be further simplified as

$$B'_{\mathcal{P},n-1}\left[\left(K^{\mathbb{Q}}_{0\mathcal{P}} - K^{\mathbb{P}}_{0\mathcal{P}}\right) + \left(K^{\mathbb{Q}}_{1\mathcal{P}} - K^{\mathbb{P}}_{1\mathcal{P}}\right)\mathcal{P}_t + \left(K^{\mathbb{Q}}_{1\mathcal{P}\mathcal{V}} - K^{\mathbb{P}}_{1\mathcal{P}\mathcal{V}}\right)\mathcal{V}_t\right] + B_{\mathcal{V},n-1}\left(\rho^{\mathbb{Q}} - \rho^{\mathbb{P}}\right)(\mathcal{V}_t - \mu)$$

Now consider a linear combination of yields $y^a_t = \sum_{i=1}^N a_i y^{n_i}_t$ where n_i denotes the maturity of the i-th yield. To construct a mimicking portfolio of it, we need to find $\{w_i\}_{i=1}^N$ such that

$$\frac{dP^w_t}{dy^a_t} = \sum_{i=1}^{N}\frac{dP^w_t}{dy^{n_i}_t}\frac{dy^{n_i}_t}{dy^a_t} = -\sum_{i=1}^{N} w_i n_i P^{n_i}_t \frac{1}{a_i} = 1$$

which will hold for weights

$$w_i = -\frac{a_i}{N n_i P^{n_i}_t}$$

Consider the one-period excess return on portfolio P^w_t:

$$\frac{\sum_i w_i\left(P^{n_i-1}_{t+1} - e^{r_t}P^{n_i}_t\right)}{\left|\sum_i w_i P^{n_i}_t\right|} = \frac{-\sum_i a_i/n_i\left(P^{n_i-1}_{t+1}/P^{n_i}_t - e^{r_t}\right)}{\left|\sum_i w_i P^{n_i}_t\right|}$$

Using $e^x \approx 1 + x$, we have

$$\begin{aligned}
\mathbb{E}^{\mathbb{P}}_t\left[P^{n_i-1}_{t+1}/P^{n_i}_t\right] &= \exp\left[k^{\mathbb{P}}_t(\mathcal{Z}_{t+1}; -B_{n-1}) - k^{\mathbb{Q}}_t(\mathcal{Z}_{t+1}; -B_{n-1}) + r_t\right] \\
&\approx 1 + k^{\mathbb{P}}_t(\mathcal{Z}_{t+1}; -B_{n-1}) - k^{\mathbb{Q}}_t(\mathcal{Z}_{t+1}; -B_{n-1}) + r_t
\end{aligned}$$

which implies that

$$\frac{-\sum_i a_i/n_i E^{\mathbb{P}}_t\left[P^{n_i-1}_{t+1}/P^{n_i}_t - e^{r_t}\right]}{\left|\sum_i w_i P^{n_i}_t\right|} = \frac{-\sum_i a_i/n_i\left[k^{\mathbb{P}}_t(\mathcal{Z}_{t+1}; -B_{n-1}) - k^{\mathbb{Q}}_t(\mathcal{Z}_{t+1}; -B_{n-1})\right]}{\left|\sum_i w_i P^{n_i}_t\right|}$$

Hence, the expected excess return on portfolio P^w_t, to a first-order approximation, is given by

$$\frac{\sum_i a_i/n_i B_{n_i-1}\left[\Lambda_0 + \Lambda_1 \mathcal{Z}_t\right]}{\left|\sum_i a_i/n_i\right|} \tag{A-12}$$

where

$$\begin{aligned}\Lambda_0 &= \begin{bmatrix} K_{0\mathcal{P}}^{\mathbb{P}} - K_{0\mathcal{P}}^{\mathbb{Q}} \\ -\left(\rho^{\mathbb{P}} - \rho^{\mathbb{Q}}\right)\mu \end{bmatrix} \\ \Lambda_1 &= \begin{bmatrix} \left(K_{1\mathcal{P}\mathcal{P}}^{\mathbb{P}} - K_{1\mathcal{P}\mathcal{P}}^{\mathbb{Q}}\right) & K_{1\mathcal{P}\mathcal{M}}^{\mathbb{P}} & \left(K_{1\mathcal{P}\mathcal{V}}^{\mathbb{P}} - K_{1\mathcal{P}\mathcal{V}}^{\mathbb{Q}}\right) \\ 0 & 0 & \rho^{\mathbb{P}} - \rho^{\mathbb{Q}} \end{bmatrix}\end{aligned}$$

and

$$B_{n_i-1} = \left(B'_{\mathcal{P},n-1}, B'_{\mathcal{V},n-1}\right)'$$

Since the first $(N-m)$ elements of $\mathcal{Z}_t$ correspond to the first $(N-m)$ principal components of yields

$$PCj_t = \sum_{i=1}^{N} l_i^j y_t^{n_i} = \sum_{i=1}^{N} l_i^j \left(A_{n_i}/n_i + B_{n_i}/n_i \mathcal{Z}_t\right)$$

it follows that $\sum_{i=1}^{N} l_i^j B_{n_i}/n_i$ is the selection vector for the j-th element (e.g. $(1,0,0)$ for $j=1$) for $j \leq N-m$. Replacing a_i of equation (A-12) with l_i^j and approximating B_{n_i-1} with B_{n_i}, we have

$$\frac{\sum_i l_i^j B_{n_i-1}/n_i \left[\Lambda_0 + \Lambda_1 \mathcal{Z}_t\right]}{\left|\sum_i l_i^j/n_i\right|} = \frac{\overbrace{\sum_i l_i^j B_{n_i}/n_i}^{\text{selction vector}} \left[\Lambda_0 + \Lambda_1 \mathcal{Z}_t\right]}{\left|\sum_i l_i^j/n_i\right|}$$

which implies that $xPCj$ is given by the j-th row of $\Lambda_0 + \Lambda_1 \mathcal{Z}_t$ scaled by $\left|\sum_i l_i^j/n_i\right|$ for $j \leq N-m$.

Appendix 7 Estimates of $\mathcal{M}_C$

Under the risk-neutral measure $\mathbb{Q}$, the latent state variable is drift-normalized as in Joslin (2015) or Creal and Wu (2015), for econometric identification. Then, the short rate equation for $\mathcal{M}_C$ is[1]

$$r_t = r_\infty^{\mathbb{Q}} + \iota' X_t + \delta_V V_t \tag{A-13}$$

where ι denotes a vector of ones. δ_V can take ± 1 or 0, and each possible value induces different local maxima: see for example Creal and Wu (2015). The conditional mean of the normalized state $Z_t \equiv (X_t', V_t)'$ is

$$\begin{bmatrix} E_t^{\mathbb{Q}}(X_{t+1}) \\ E_t^{\mathbb{Q}}(V_{t+1}) \end{bmatrix} = \begin{bmatrix} 0 \\ K_{0V}^{\mathbb{Q}} \end{bmatrix} + \begin{bmatrix} diag\left(\lambda^{\mathbb{Q}}\right) & 0 \\ 0 & \rho^{\mathbb{Q}} \end{bmatrix} \begin{bmatrix} X_t \\ V_t \end{bmatrix} \tag{A-14}$$

where $K_{0V}^{\mathbb{Q}}$ is the product of the scale and shape parameters of V_t.

Table A-1 presents the estimates of the $\mathbb{Q}$ parameters for $\mathcal{M}_C$ and $\mathcal{M}_J$. The persistency of Gaussian factors, measured by $\lambda^{\mathbb{Q}}$, is similar across the two models. For $\mathcal{M}_J$, δ_V in equation (A-13) should be zero and V_t does not affect the conditional variance of X_t. Then, $r_\infty^{\mathbb{Q}}$ can be interpreted as the long-run $\mathbb{Q}$ mean of the short rate, since the long-run mean of X_t is set to zero under $\mathbb{Q}$

[1] For ease of explanation, $K_{1X}^{\mathbb{Q}}$ is assumed to have real distinct eigenvalues - this is overidentifying. For details, see Joslin, Singleton, and Zhu (2011).

as in equation (A-14). For $\mathcal{M}_C$, the long-run $\mathbb{Q}$ mean of the short rate is $r_\infty^{\mathbb{Q}} + \delta_V K_{0V}^{\mathbb{Q}} / \left(1 - \rho^{\mathbb{Q}}\right)$ rather than $r_\infty^{\mathbb{Q}}$, which induces the difference between the $r_\infty^{\mathbb{Q}}$s of $\mathcal{M}_C$ and $\mathcal{M}_J$. In addition, the likelihood of $\mathcal{M}_C$ is maximized with $\delta_V = 1$ among the three possible values of δ_V. This is in line with $r_\infty^{\mathbb{Q}}$ of $\mathcal{M}_C$ being slightly less than $r_\infty^{\mathbb{Q}}$ of $\mathcal{M}_J$ for the two models to have similar levels in the long-run $\mathbb{Q}$ mean of the short rate.

For each maturity n, Table A-2 reports the fraction $\text{var}(B_n \iota_i \iota_i' Z_t) / \text{var}(B_n Z_t)$ where ι_i denotes a vector of zero with i-th element being one. The table thus represents the relative contribution of each latent factor toward the yields curve movement. Note that the exercises in Section 1.4 analyze the marginal impact of $\mathcal{V}_t$ after controlling the yield curve factors. Because the yield curve factors themselves, $\mathcal{P}_t$, are linear functions of latent factors X_t and V_t, the sole impact of $\mathcal{V}_t$ on the shape of the yield cannot be assessed in this setting. Within a fully-fledged ADTSM, the impact of V_t on the cross-section of yields can be completely isolated from the impact of latent Gaussian factor X_t. The table performs this exercise, and shows that the shape of the yields curve is largely unexplained by V_t or $\mathcal{V}_t$. Note that the interpretations of $\mathcal{V}_t$ and V_t are freely interchangeable in the context, since one is an invariant transformation of the other: one can freely scale up or down V_t, then yield loadings on V_t are adjusted accordingly so that its impact on the yield curve still remains the same.

Table A-1
Persistence Parameters

This table reports the estimates of persistence parameters for the each model $\mathcal{M}_C$ and $\mathcal{M}_J$. Standard errors are given in parentheses.

Model	$r_\infty^{\mathbb{Q}}$	$\lambda_1^{\mathbb{Q}}$	$\lambda_2^{\mathbb{Q}}$	$\lambda_3^{\mathbb{Q}}$	$\rho^{\mathbb{Q}}$
$\mathcal{M}_C$	0.098	0.996	0.963	0.903	0.949
	(0.012)	(0.000)	(0.003)	(0.010)	(0.036)
$\mathcal{M}_J$	0.108	0.996	0.963	0.906	
	(0.003)	(0.000)	(0.002)	(0.008)	

Table A-2
Yield Curve Decomposition

This table reports the fraction $\text{var}(B_n \iota_i \iota_i' Z_t) / \text{var}(B_n Z_t)$ where ι_i denotes a vector of zero with i-th element being one.

$\iota_i' Z_t$	n						
	6 mon	1 yr	2 yr	3 yr	5 yr	7 yr	10 yr
X_1 ($\equiv \iota_1' Z_t$)	1.20	1.22	1.28	1.37	1.47	1.49	1.45
X_2 ($\equiv \iota_2' Z_t$)	3.60	3.02	2.24	1.73	1.07	0.69	0.40
X_3 ($\equiv \iota_3' Z_t$)	0.64	0.39	0.18	0.10	0.05	0.03	0.01
V ($\equiv \iota_4' Z_t$)	0.03	0.03	0.02	0.01	0.01	0.00	0.00

Appendix 8 Proofs of Proposition 2 and Proposition 3

Proof of Proposition 2. Following Carr and Madan (1998), we assume that there exist a function $g(F_{T,T})$ and that it is twice differentiable. It then follows that

$$\begin{aligned} g(F_{T,T}) &= g(x) + g'(x)\,(F_{T,T} - x) + \int_0^x g''(K)\,(K - F_{T,T})^+\,dK \\ &\quad + \int_x^\infty g''(K)\,(F_{T,T} - K)^+\,dK \end{aligned}$$

for any $x \geq 0$. If $x = F_{t,T}$, then

$$\begin{aligned} g(F_{T,T}) &= g(F_{t,T}) + g'(F_{t,T})\,(F_{T,T} - F_{t,T}) + \int_0^{F_{t,T}} g''(K)\,(K - F_{T,T})^+\,dK \\ &\quad + \int_{F_{t,T}}^\infty g''(K)\,(F_{T,T} - K)^+\,dK \end{aligned}$$

Now assumr $g(F) = \log F$. Then, we get

$$\begin{aligned} \log F_{T,T} &= \log F_{t,T} + \frac{1}{F_{t,T}}\,(F_{T,T} - F_{t,T}) \\ &\quad - \left(\int_0^{F_{t,T}} \frac{(K - F_{T,T})^+}{K^2} dK + \int_{F_{t,T}}^\infty \frac{(F_{T,T} - K)^+}{K^2} dK \right). \end{aligned} \tag{A-15}$$

Given that $d\left(\log F_{t,T}\right) = \frac{dF_{t,T}}{F_{t,T}} - \frac{1}{2}\sigma_t^2 dt$ due to Ito's lemma, the quadratic variation of $F_{t,T}$ can be written as

$$\begin{aligned} \int_t^T \sigma_u^2 du &= -2\log\frac{F_{T,T}}{F_{t,T}} + \int_t^T \frac{dF_{u,T}}{F_{u,T}} \\ &= 2\left(\frac{F_{T,T} - F_{t,T}}{F_{t,T}} - \log\frac{F_{T,T}}{F_{t,T}}\right) + 2\int_t^T \left(\frac{1}{F_{u,T}} - \frac{1}{F_{t,T}}\right) dF_{u,T} \\ &= 2\left(\int_0^{F_{t,T}} \frac{(K - F_{T,T})^+}{K^2} dK + \int_{F_{t,T}}^\infty \frac{(F_{T,T} - K)^+}{K^2} dK \right) \\ &\quad + 2\int_t^T \left(\frac{1}{F_{u,T}} - \frac{1}{F_{t,T}}\right) dF_{u,T}. \end{aligned}$$

where the last equality follows from Equation (A-15). Since $F_{t,T}$ is a martingale under the $\mathbb{Q}_T$ measure, the $\mathbb{Q}_T$ expectation of the dynamic strategy is zero and hence this implies that the market price of the realized variance is equal to $\tilde{F}_{t,T}$, i.e.

$$\tilde{F}_{t,T} = \mathbb{E}_t^{\mathbb{Q}_T}\left[\int_t^T \sigma_u^2 du\right].$$

In words, replicating the expression for the variance $\mathrm{RV}_{t,T}^{\log}$ involves

1. a path-independent payoff in out-of-the-money options
2. a dynamic trading strategy which is re-balanced to hold $2\left(\frac{1}{F_{t+i-1,T}} - \frac{1}{F_{t,T}}\right)$ of the underlying.

$\square$

Proof of Proposition 3.

$$
\begin{aligned}
\widetilde{RV}_{t,T} &= -2\log\frac{F_{T,T}}{F_{t,T}} + 2\sum_{i=1}^{T-t}\frac{F_{t+i,T}-F_{t+i-1,T}}{F_{t+i-1,T}} \\
&= 2\left(\frac{F_{T,T}-F_{t,T}}{F_{t,T}} - \log\frac{F_{T,T}}{F_{t,T}}\right) \\
&\quad +2\sum_{i=1}^{T-t}\left(\frac{1}{F_{t+i-1,T}}-\frac{1}{F_{t,T}}\right)\left(F_{t+i,T}-F_{t+i-1,T}\right) \\
&= 2\left(\int_0^{F_{t,T}}\frac{(K-F_{T,T})^+}{K^2}dK + \int_{F_{t,T}}^{\infty}\frac{(F_{T,T}-K)^+}{K^2}dK\right) \\
&\quad +2\sum_{i=1}^{T-t}\left(\frac{1}{F_{t+i-1,T}}-\frac{1}{F_{t,T}}\right)\left(F_{t+i,T}-F_{t+i-1,T}\right)
\end{aligned}
$$

where the last equality follows from the proof of Proposition 2. $\square$

Appendix 9 Squared Returns and New Variance Measure

Relationship between squared returns and new variance measure. The calculations follow Carr and Lee (2009). We start with a Taylor expansion of $2\log F_{t+i,T}$ around $F_{t,T}$,

$$
\begin{aligned}
2\log F_{t+i,T} &= 2\log F_{t,T} + \frac{2}{F_{t,T}}\left(F_{t+i,T}-F_{t,T}\right) - \left(\frac{F_{t+i,T}-F_{t,T}}{F_{t,T}}\right)^2 + \\
&\quad \frac{2}{3}\left(\frac{F_{t+i,T}-F_{t,T}}{F_{t,T}}\right)^3 + O\left(r_{t+i,T}^4\right) \\
&= 2\log F_{t,T} + 2r_{t+i,T} - r_{t+i,T}^2 + \frac{2}{3}r_{t+i,T}^3 + O\left(r_{t+i,T}^4\right),
\end{aligned}
$$

where $r_{t+i,T} := F_{t+i,T}/F_{t,T} - 1$. Rearranging the right and left hand sides of the above equation yields

$$
r_{t+i,T}^2 = 2r_{t+i,T} - 2\left(\log F_{t+i,T} - \log F_{t,T}\right) + \frac{2}{3}r_{t+i,T}^3 + O\left(r_{t+i,T}^4\right).
$$

Summing over different i's we get

$$
\begin{aligned}
RV_{t,T} &= \widetilde{RV}_{t,T} + \frac{2}{3}\sum_{i=1}^{n} r_{t+i,T}^3 + \sum_{i=1}^{n} O\left(r_{t+i,T}^4\right) \\
&= \widetilde{RV}_{t,T} + O\left(\sum_{i=1}^{n} r_{t+i,T}^3\right),
\end{aligned}
$$

where $RV_{t,T} = \sum_{i=1}^{n} r_{t+i,T}^2$ and $\widetilde{RV}_{t,T} = 2\sum_{i=1}^{n}\left[r_{t+i,T} - \log\left(1+r_{t+i,T}\right)\right]$. $\square$

Bibliography

Adrian, T., and H. S. Shin (2010): "Liquidity and Leverage," *Journal of Financial Intermediation*, 19, 418–437.

Ahn, D.-H., R. F. Dittmar, and A. R. Gallant (2002): "Quadratic Term Structure Models: Theory and Evidence," *Review of Financial Studies*, 15(1), 243–288.

Aït-Sahalia, Y., M. Karaman, and L. Mancini (2015): "The Term Structure of Variance Swaps and Risk Premia," Working Paper, Princeton University.

Aït-Sahalia, Y., and A. W. Lo (2000): "Nonparametric risk management and implied risk aversion," *Journal of Econometrics*, 94(1–2), 9–51.

Allen, P., S. Einchomb, and N. Granger (2006): "Variance Swaps," Discussion paper, JPMorgan.

Almeida, C., J. Graveline, and S. Joslin (2011): "Do Interest Rate Options Contain Information About Excess Bond Returns?," *Journal of Econometrics*, 164, 35–44.

Amengual, D., and D. Xiu (2014): "Resolution of Policy Uncertainty and Sudden Declines in Volatility," Working paper, CEMFI.

Andersen, T. G., and L. Benzoni (2010): "Do Bonds Span Volatility Risk in the U.S. Treasury Market? A Specification Test for Affine Term Structure Models," *Journal of Finance*, 65, 603–653.

Andries, M., T. Eisenbach, M. Schmalz, and Y. Wang (2015): "Price vs. Quantity in the Term Structure of Variance Risk Premia," Working Paper, University of Toulouse.

Ang, A., R. J. Hodrick, Y. Xing, and X. Zhang (2006): "The Cross-Section of Volatility and Expected Returns," *Journal of Finance*, 61, 259 – 299.

Backus, D., M. Chernov, and S. Zin (2014): "Sources of Entropy in Representative Agent Models," *The Journal of Finance*, 69(1), 51–99.

Bansal, R., and I. Shaliastovich (2013): "A Long-Run Risks Explanation of Predictability Puzzles in Bond and Currency Markets," *Review of Financial Studies*, 26, 1–33.

Bansal, R., and A. Yaron (2004): "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, 59, 1481–1509.

Bao, J., J. Pan, and J. Wang (2011): "The Illiquidity of Corporate Bonds," *Journal of Finance*, 66, 911–946.

Bauer, M. D., and G. D. Rudebusch (2016): "Resolving the Spanning Puzzle in Macro-Finance Term Structure Models," *forthcoming, Review of Finance.*

Bean, C. (2005): "Commentary: separating the business cycle from other economic fluctuations," in *Proceedings-Economic Policy Symposium-Jackson Hole*, no. Aug, pp. 181–186. Federal Reserve Bank of Kansas City.

Bekaert, G., and M. Hoerova (2014): "The VIX, the variance premium and stock market volatility," *Journal of Econometrics*, 183, 181–192.

Bekaert, G., M. Hoerova, and M. L. Duca (2013): "Risk, Uncertainty and Monetary Policy," *Journal of Monetary Economics*, 60, 771–788.

Bekaert, G., and J. Liu (2004): "Conditioning Information and Variance Bounds on Pricing Kernels," *Review of Financial Studies*, 17(2), 339–378.

Bibkov, R., and M. Chernov (2009): "Unspanned Stochastic Volatility in Affine Models: Evidence from EurodollarFutures and Options," *Management Science*, 55, 1292–1305.

——— (2011): "Yield Curve and Volatility: Lessons from Eurodollar Futures and Options," *Journal of Financial Econometrics*, 9, 66–105.

Björk, T. (2009): *Arbitrage Theory in Continuous Time*. Oxford University Press, 3rd edn.

Black, F. (1976): "The Pricing of Commodity Contracts," *Journal of Financial Economics*, 3, 167–179.

Black, F., and M. Scholes (1973): "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81, 637–659.

Bollerslev, T. G., G. Tauchen, and H. Zhou (2009): "Expected Stock Returns and Variance Risk Premia," *Review of Financial Studies*, 22, 4463–4492.

Bondarenko, O. (2014): "Variance Trading and Market Price of Variance Risk," *Journal of Econometrics*, 180, 81–97.

Breeden, D., and R. H. Litzenberger (1978): "Prices of State-Contingent Claims Implicit in Option Prices," *Journal of Business*, 51, 621–651.

Breeden, D. T. (1979): "An intertemporal asset pricing model with stochastic consumption and investment opportunities," *Journal of Financial Economics*, 7(3), 265 – 296.

Broadie, M., and A. Jain (2008): "The Effect of Jumps and Discrete Sampling on Volatility and Variance Swaps," *International Journal of Theoretical and Applied Finance*, 11, 761–797.

Brunnermeier, M., S. Nagel, and L. H. Pedersen (2009): "Carry Trades and Currency Crashes," *NBER Macroeconomics Annual 2008*, 23, 313–347.

Cakici, N., and J. Zhu (2001): "Pricing Eurodollar Futures Options with the Heath-Jarrow-Morton Model," *Journal of Futures Markets*, 21, 655–680.

Cambpell, J. Y., and R. J. Shiller (1991): "Yield Spreads and Interest Rate Movements: A Bird's Eye View," *Review of Economic Studies*, 58, 495–514.

Campbell, J. Y. (1993): "Intertemporal Asset Pricing without Consumption Data," *The American Economic Review*, 83(3), pp. 487–512.

Campbell, J. Y., and J. H. Cochrane (1999): "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107, 205–251.

Carr, P., and R. Lee (2009): "Volatility Derivatives," *Annual Review of Financial Economics*, 1, 1–21.

Carr, P., and D. Madan (1998): "Towards a Theory of Volatility Trading," in *Volatility: New Estimation Techniques for Pricing Derivatives*, ed. by R. Jarrow, pp. 417–427. RISK Publications.

Carr, P., and L. Wu (2009): "Variance Risk Premiums," *Review of Financial Studies*, 22(3), 1311–1341.

Chabi-Yo, F. (2008): "Conditioning Information and Variance Bounds on Pricing Kernels with Higher- Order Moments: Theory and Evidence," *Review of Financial Studies*, 21(1), 181–231.

Chen, R.-R., and L. Scott (1993): "Maximum Likelihood Estimation for a Multifactor Equilibrium Model of the Term Structure of Interest Rates," *The Journal of Fixed Income*, 3, 14–31.

Cheridito, P., D. Filipović, and R. L. Kimmel (2007): "Market price of risk specifications for affine models: Theory and evidence," *Journal of Financial Economics*, 83(1), 123 – 170.

Chernov, M., and P. Mueller (2012): "The Term Structure of Inflation Expectations," *Journal of Financial Economics*, 106, 367–394.

Choi, H., P. Mueller, and A. Vedolin (2016): "Bond Variance Risk Premiums," *Working Paper, London School of Economics.*

Cieslak, A., and P. Povala (2016): "Information in the Term Structure of Yield Curve Volatility," *Journal of Finance*, 71, 1393–1434.

Cochrane, J. H., and M. Piazzesi (2005): "Bond Risk Premia," *American Economic Review*, 95, 138–160.

——— (2008): "Decomposing the Yield Curve," Working Paper, Chicago Booth.

Collin-Dufresne, P., and R. S. Goldstein (2002): "Do Bonds Span the Fixed-Income Markets?," *Journal of Finance*, 57, 1685—1730.

Collin-Dufresne, P., R. S. Goldstein, and C. S. Jones (2009): "Can interest rate volatility be extracted from the cross section of bond yields?," *Journal of Financial Economics*, 94, 47–66.

Coval, J., and T. Shumway (2001): "Expected Option Returns," *Journal of Finance*, 56, p. 983–1009.

Creal, D. D., and J. C. Wu (2015): "Estimation of Affine Term Structure Models with Spanned or Unspanned Stochastic Volatility," *Journal of Econometrics*, 185(1), 60–81.

Dai, Q., and K. Singleton (2002): "Expectation Puzzles, Time-Varying Risk Premia, and Affine Models of the Term Structure," *Journal of Financial Economics*, 63, 415–441.

DAI, Q., AND K. J. SINGLETON (2000): "Specification Analysis of Affine Term Structure Models," *The Journal of Finance*, 55(5), 1943–1978.

DEMETERFI, K., E. DERMAN, M. KAMAL, AND J. ZOU (1999): "A Guide to Volatility and Variance Swaps," *Journal of Derivatives*, 6, 9–32.

DEW-BECKER, I., S. GIGLIO, A. LE, AND M. RODRIGUEZ (2016): "The Price of Variance Risk," *forthcoming, Journal of Financial Economics*, Working Paper, UNC Chapel Hill.

DRECHSLER, I., AND A. YARON (2011): "What's Vol Got to Do With It," *Review of Financial Studies*, 24, 1–45.

DUARTE, J. (2004): "Evaluating an Alternative Risk Preference in Affine Term Structure Models," *The Review of Financial Studies*, 17(2), pp. 379–404.

DUARTE, J., F. LONGSTAFF, AND F. YU (2007): "Risk and Return in Fixed-Income Arbitrage: Nickels in Front of a Steamroller?," *Review of Financial Studies*, 20, 769–811.

DUFFEE, G. R. (2010): "Sharpe Ratios in Term Structure Models," Working Paper, Johns Hopkins University.

——— (2011a): "Forecasting with the Term Structure: the Role of No-arbitrage Restrictions," Working Paper, Johns Hopkins University.

——— (2011b): "Information in (and not in) the Term Structure," *Review of Financial Studies*, 24(9), 2895–2934.

DUFFIE, D., AND L. G. EPSTEIN (1992): "Asset Pricing with Stochastic Differential Utility," *Review of Financial Studies*, 5(3), 411–436.

DUFFIE, D., D. FILIPOVIĆ, AND W. SCHACHERMAYER (2003): "Affine Processes and Applications in Finance," *The Annals of Applied Probability*, 13(3), 984–1053.

DUFFIE, D., AND R. KAN (1996): "A yield-factor model of interest rates," *Mathematical finance*, 6, 379–406.

DUPIRE, B. (1994): "Pricing with a Smile," *Risk*, 7, 18–20.

DUYVESTEYN, J., AND G. DE ZWART (2015): "Riding the swaption curve," *Journal of Banking & Finance*, 59, 57–75.

EGLOFF, D., M. LEIPPOLD, AND L. WU (2010): "The Term Structure of Variance Swap Rates and Optimal Variance Swap Investments," *Journal of Financial and Quantitative Analysis*, 45, 1279–1310.

FAMA, E. F., AND R. R. BLISS (1987): "The Information in Long-Maturity Forward Rates," *American Economic Review*, 77, 680–692.

FAMA, E. F., AND J. MACBETH (1973): "Risk, Return, and Equilibrium: Empirical Tests," *Journal of Political Economy*, 81, 607–636.

FERSON, W. E., AND A. F. SIEGEL (2003): "Stochastic Discount Factor Bounds with Conditioning Information," *Review of Financial Studies*, 16(2), 567–595.

FILIPOVIĆ, D., E. GOURIER, AND L. MANCINI (2015): "Quadratic Variance Swap Models," *forthcoming, Journal of Financial Economics*.

FLESAKER, B. (1993): "Arbitrage Free Pricing of Interest Rate Futures and Forward Contracts," *Journal of Futures Markets*, 13, 77–91.

George, T. J., and F. A. Longstaff (1993): “Bid-Ask Spreads and Trading Activity in the S&P100 Index Option Market,” *Journal of Financial and Quantitative Analysis*, 28, 381–397.

Ghosh, A., C. Julliard, and A. P. Taylor (2016a): “An Information Based One-Factor Asset Pricing Model,” Working Paper, Carnegie Mellon.

——— (2016b): “What is the Consumption-CAPM Missing? An Information-Theoretic Framework for the Analysis of Asset Pricing Models,” *forthcoming, Review of Financial Studies*.

Gourieroux, C., and J. Jasiak (2006): “Autoregressive gamma processes,” *Journal of Forecasting*, 25(2), 129–152.

Gürkaynak, R., B. Sack, and J. Wright (2007): “The U.S. Treasury Yield Curve: 1961 to the Present,” *Journal of Monetary Economics*, 54, 2291–2304.

Hamilton, J. D., and J. C. Wu (2012): “Identification and Estimation of Gaussian Affine Term Structure Models,” *Journal of Econometrics*, 168(2), 315 – 331.

Hansen, L., and R. Jagannathan (1991): “Implications of Security Market Data for Models of Dynamic Economies,” *Journal of Political Economy*, 99, 225–262.

Jagannathan, R., A. Kaplin, and S. Sun (2003): “An evaluation of multi-factor CIR models using LIBOR, swap rates, and cap and swaption prices,” *Journal of Econometrics*, 116, 113 – 146.

Joslin, S. (2014): “Pricing and Hedging in Fixed Income Markets,” Working Paper, USC.

——— (2015): “Can Unspanned Stochastic Volatility Models Explain the Cross Section of Bond Volatilities?,” Working Paper, USC.

Joslin, S., and A. Le (2014): “Interest Rate Volatility and No-Arbitrage Affine Term Structure Models,” Working Paper, USC.

Joslin, S., A. Le, and K. J. Singleton (2013): “Why Gaussian Macro-Finance Term Structure Models Are (Nearly) Unconstrained Factor-VARs,” *Journal of Financial Economics*, 109, 604 – 622.

Joslin, S., M. Priebsch, and K. J. Singleton (2014): “Risk Premiums in Dynamic Term Structure Models with Unspanned Macro Risks,” *Journal of Finance*, 69, 1197–1233.

Joslin, S., K. J. Singleton, and H. Zhu (2011): “A New Perspective on Gaussian Dynamic Term Structure Models,” *Review of Financial Studies*.

Kelly, B., L. Pástor, and P. Veronesi (2014): “The Price of Political Uncertainty: Theory and Evidence From the Option Market,” Working Paper, Chicago Booth.

Kim, D. H. (2007): “Spanned stochastic volatility in bond markets: a reexamination of the relative pricing between bonds and bond options,” Unpublished working paper 239. BIS, Basel, Switzerland.

Kleibergen, F., and R. Paap (2006): “Generalized reduced rank tests using the singular value decomposition,” *Journal of Econometrics*, 133(1), 97 – 126.

Le, A., and K. J. Singleton (2013): "The Structure of Risks in Equilibrium Affine Models of Bond Yields," Working Paper, UNC.

Le, A., K. J. Singleton, and Q. Dai (2010): "Discrete-Time Affine Term Structure Models with Generalized Market Prices of Risk," *Review of Financial Studies*, 23(5), 2184–2227.

Lewellen, J., S. Nagel, and J. Shanken (2010): "A Skeptical Appraisal of Asset Pricing Tests," *Journal of Financial Economics*, 96, 175–194.

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

Lo, A. W. (2002): "The Statistics of Sharpe Ratios," *Financial Analyst Journal*, 58, 36–52.

Lucas, R. E. (1978): "Asset Prices in an Exchange Economy," *Econometrica*, 46, 1429–1446.

Ludvigson, S. C., and S. Ng (2009): "Macro Factors in Bond Risk Premia," *Review of Financial Studies*, 22, 5027–5067.

Lustig, H., N. Roussanov, and A. Verdelhan (2011): "Common Risk Factors in Currency Markets," *Review of Financial Studies*, 24, 3731–3777.

Malkhozov, A., P. Mueller, A. Vedolin, and G. Venter (2016a): "International Illiquidity," Working Paper, London School of Economics.

——— (2016b): "Mortgage Risk and the Yield Curve," *Review of Financial Studies*, 29(5), 1220–1253.

Martin, I. (2013): "Simple Variance Swaps," Working Paper, Stanford University.

——— (2015): "What is the Expected Return on the Market?," Working Paper, London School of Economics.

Mele, A., and Y. Obayashi (2013): "Volatility Indexes and Contracts for Government Bonds and Time Deposits," Working Paper, Applied Academics LLC.

Menzly, L., T. Santos, and P. Veronesi (2004): "Understanding Predictability," *Journal of Political Economy*, 112, 1–47.

Merener, N. (2012): "Swap Rate Variance Swaps," *Quantitative Finance*, 12, 249–261.

Mueller, P., P. Sabtchevsky, A. Vedolin, and P. Whelan (2016): "Variance Risk Premia on Stocks and Bonds," Working Paper, London School of Economics.

Nagel, S. (2012): "Evaporating liquidity," *Review of Financial Studies*, 25, 2005–2039.

Neuberger, A. (1994): "Hedging Volatility - the Case for a New Contract," *Journal of Portfolio Management*, 20(2), 74–80.

——— (2012): "Realized Skewness," *Review of Financial Studies*, 25, 3423–3455.

Newey, W. K., and K. D. West (1987): "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55(3), 703–708.

Orphanides, A. (2003): "Historical monetary policy analysis and the Taylor rule," *Journal of Monetary Economics*, 50(5), 983 – 1022.

Pástor, L., and R. F. Stambaugh (2003): "Liquidity Risk and Expected Stock Returns," *Journal of Political Economy*, 111, p. 642 – 685.

Pearson, N. D., and T.-S. Sun (1994): "Exploiting the Conditional Density in Estimating the Term Structure: An Application to the Cox, Ingersoll, and Ross Model," *Journal of Finance*, 49(4), 1279–1304.

Piazzesi, M. (2010): "Affine Term Structure Models," in *Handbook of Financial Econometrics: Tools and Techniques*, ed. by Y. Aït-Sahalia, and L. Hansen, vol. 1 of *Handbooks in Finance*, pp. 691 – 766. North-Holland, San Diego.

Piazzesi, M., M. Schneider, and S. Tuzel (2007): "Housing, consumption and asset pricing," *Journal of Financial Economics*, 83(3), 531 – 569.

Rosenberg, J., and R. Engle (2002): "Empirical pricing kernels," *Journal of Financial Economics*, 64(3), 341–372.

Ross, S. (2015): "The Recovery Theorem," *The Journal of Finance*, 70(2), 615–648.

Rubinstein, M. (1976): "The Valuation of Uncertain Income Streams and the Pricing of Options," *The Bell Journal of Economics*, 7(2), 407–425.

Rudebusch, G. D. (2006): "Monetary Policy Inertia: Fact or Fiction?," *International Journal of Central Banking*, 2(4), 88–135.

Santa-Clara, P., and A. Saretto (2009): "Option Strategies: Good Deals and Margin Calls," *Journal of Financial Markets*, 12, 391–417.

Schneider, P. (2015): "Generalized Risk Premia," *Journal of Financial Economics*, 116, 487–504.

Schneider, P., and F. Trojani (2015): "Fear Trading," Working Paper, University of Lugano.

Shanken, J. (1985): "Multivariate tests of the zero-beta CAPM," *Journal of Financial Economics*, 14(3), 327 – 348.

——— (1992): "On the Estimation of Beta Pricing Models," *Review of Financial Studies*, 5, 1–34.

Taylor, J. B. (1993): "Discretion versus policy rules in practice," *Carnegie-Rochester Conference Series on Public Policy*, 39(1), 195–214.

——— (1999): "A Historical Analysis of Monetary Policy Rules," in *Monetary Policy Rules*, NBER Chapters, pp. 319–348. National Bureau of Economic Research, Inc.

Trolle, A. B. (2009): "The Price of Interest Rate Variance Risk And Optimal Investments in Interest Rate Derivatives," Working Paper, EPFL.

Trolle, A. B., and E. S. Schwartz (2009): "A General Stochastic Volatility Model for the Pricing of Interest Rate Derivatives," *Review of Financial Studies*, 22, 2007–2057.

——— (2014): "The Swaption Cube," *Review of Financial Studies*, 27, 2307–2353.

Tsai, J. (2016): "Rare Disasters and the Term Structure of Interest Rates," Working Paper, University of Oxford.

van Binsbergen, J. H., and R. S. Koijen (2016): "The Term Structure of Returns: Facts and Theory," *forthcoming, Journal of Financial Economics*.

WACHTER, J. (2006): "A Consumption-Based Model of the Term Structure of Interest Rates," *Journal of Financial Economics*, 79, 365–399.

WACHTER, J. A. (2013): "Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility?," *The Journal of Finance*, 68(3), 987–1035.

WU, L. (2010): "Variance Dynamics: Joint Evidence from Options and High-Frequency Returns," *Journal of Econometrics*, 160, 280–287.

ZHOU, G., AND Y. ZHU (2015): "Macroeconomic Volatilities and the Long-run Risks of Asset Prices," *Management Science*, 61, 413 – 430.