Van Tassel, Peter

Working Paper

Relative pricing and risk premia in equity volatility markets

Staff Report, No. 867

Provided in Cooperation with:
Federal Reserve Bank of New York

Suggested Citation: Van Tassel, Peter (2018) : Relative pricing and risk premia in equity volatility markets, Staff Report, No. 867, Federal Reserve Bank of New York, New York, NY

This Version is available at:
https://hdl.handle.net/10419/210719

Federal Reserve Bank of New York
Staff Reports

Relative Pricing and Risk Premia in Equity Volatility Markets

Peter Van Tassel

Staff Report No. 867
September 2018

Relative Pricing and Risk Premia in Equity Volatility Markets

Peter Van Tassel

Federal Reserve Bank of New York Staff Reports, no. 867
September 2018
JEL classification: C58, G12, G13

Abstract

This paper provides empirical evidence that volatility markets are integrated through the time-varying term structure of variance risk premia. These risk premia predict the returns from selling volatility for different horizons, maturities, and products, including variance swaps, straddles, and VIX futures. In addition, the paper derives a closed-form relationship between the prices of variance swaps and VIX futures. While tightly linked, VIX futures exhibit deviations of varying significance from the no-arbitrage prices and bounds implied by the variance swap market. The paper examines these pricing errors and their relationship to VIX futures' return predictability.

Key words: variance swaps, term structure, variance risk premium, VIX futures, options, return predictability

---

Van Tassel: Federal Reserve Bank of New York (email: peter.vantassel@ny.frb.org). The author thanks Tobias Adrian, Torben Andersen, Tim Bollerslev, Richard Crump, Robert Engle, Cam Harvey, David Lucca, George Tauchen, and Erik Vogt for helpful comments and conversations, as well as seminar participants at Duke University, the Federal Reserve Bank of New York, and NYU Stern. The views expressed in this paper are those of the author and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

To view the author's disclosure statement, visit
https://www.newyorkfed.org/research/staff_reports/sr867.html.

1 Introduction

In efficient financial markets, the relative prices and risk premia of assets with closely related payoffs are tightly linked. By adhering to no-arbitrage restrictions and forecasting returns, relative prices and risk premia can provide evidence that markets are integrated by a common stochastic discount factor (Gromb and Vayanos 2010). In contrast, deviations of relative prices from no-arbitrage relationships and differences in return predictability for similar assets can indicate market inefficiencies and segmentation (Merton 1987, Shleifer and Vishny 1997). Documenting the behavior of closely related assets is thus important for understanding how financial markets function and for testing different asset pricing theories.

Equity volatility markets provide an ideal setting to study the relative pricing and return predictability of closely related assets. Since the financial crisis, rapid growth in the trading of S&P 500 index options and VIX futures has led to the development of separate derivatives markets where investors and firms can manage their volatility and stock market risk. As of 2016, the average open interest in VIX futures was over 414 thousand contracts per day, a more than 10-fold increase over the past decade equal to approximately $414 million of gains and losses for each one-point change in the VIX. In comparison, the 2016 average daily open interest for S&P 500 index options was $2.38 billion of Black-Scholes vega, more than five times the VIX futures open interest.

When managing volatility risk, investors can now trade in either of these large and liquid exchange-traded-markets. In practice, volatility traders often use separate models for valuing and hedging different derivatives. This risk management approach can make it difficult to determine whether relative valuations and risk exposures are accurate, as different models may not be consistent with each other (Longstaff et al. 2001). In theory, however, arbitrage pricing places restrictions on the relative valuation of different derivatives. For example, within equity volatility markets, variance swaps can be valued from a portfolio of options by using a model-free formula that holds under certain assumptions (Carr and Wu 2009). This relationship forms the basis for the VIX index. Similarly, VIX futures can be valued from variance swaps and VIX options, as well as bounded by variance and volatility swaps (Carr and Wu 2005). How accurate are these no-arbitrage relationships in practice? When investors buy a volatility hedge or sell volatility to earn the variance risk premium, should they trade in the index options, variance swap, or VIX futures market?

This paper examines these questions from the perspective of a dynamic term-structure model that provides closed form prices for variance swaps and VIX futures. The model is estimated with synthetic variance swap rates that are computed from index option prices and realized variance data. As such, the model prices for VIX futures can be interpreted as

the fair value or no-arbitrage price implied by variance swaps. In addition to relative pricing, the model decomposes variance swap rates into distinct measures of financial stability that are of direct interest to investors and policymakers: realized variance forecasts and variance term premia. The realized variance forecasts measure the expected quantity of stock market volatility over different horizons. The variance term premia measure the expected holding period return from receiving fixed in variance swaps over different horizons.[1] In addition to tracking investor risk aversion, these measures can be used for risk management and portfolio choice decisions.

The paper tests the hypothesis that equity volatility markets are integrated by examining the model's accuracy in pricing VIX futures and the model's ability to predict the returns from selling volatility across different markets and products, despite being estimated with only realized variance and variance swap rate data. The paper finds mixed empirical results. On one hand, there is significant evidence of market efficiency and integration across volatility markets. Synthetic variance swap rates constructed from index option prices closely track over-the-counter variance swap quotes. VIX futures prices implied by variance swaps and the no-arbitrage model closely track observed futures prices. Model expected returns significantly forecast the returns from selling volatility through variance swaps, index option straddles, and VIX futures. On the other hand, there is also evidence of inefficiency and segmentation. While the VIX futures pricing errors are small on average, their size varies significantly over time and tends to increase during periods of financial distress. The pricing errors also predict VIX futures returns, which suggests that VIX futures are mispriced at times relative to the fair value implied by variance swaps and the model. A pseudo out-of-sample trading strategy in VIX futures based on the model expected returns and pricing errors earns an annualized Sharpe ratio of 1.80 from 2005 to 2016 with minimal stock market exposure.

In comparison to the literature, the model in this paper obtains a closed form relationship between the prices of variance swaps and VIX futures by modeling the logarithm of realized variance. Existing affine and quadratic models deliver closed form solutions for the prices of variance swaps but not VIX futures (Egloff et al. 2010, Filipović et al. 2016, Eraker and Wu 2017). Beyond pricing VIX futures, modeling the logarithm of realized variance is also advantageous because it guarantees non-negative variance swap rates and realized variance forecasts, unlike affine models. This restriction is important in low volatility environments because negative variance swap rates are arbitrage opportunities, similar to zero lower bound violations in fixed income settings. More broadly, this approach builds on Andersen et al.

[1]This paper uses the terms realized variance forecasts and variance risk premia interchangeably with the terms volatility forecasts and volatility risk premia. The former terms are technically what the no-arbitrage model estimates. Similarly, the paper uses the terms risk premia and term premia to refer to expected holding period returns.

(2003) and Andersen et al. (2007) who forecast volatility using the logarithm of realized variance.

The model prices realized variance exactly by including the logarithm of realized variance in the state vector as the observable payoff to the floating leg of a variance swap. This approach makes estimation fast and tractable as it avoids the need to filter latent stochastic volatility factors. A detailed investigation of the model's in-sample and out-of-sample pricing errors finds that a three-factor logarithmic model performs well relative to competing models of different sizes and to linear models. The three factors are the logarithm of realized variance and the first two principal components from the logarithm of variance swap rates. This setup is similar to fixed income models that use the short rate, level, and slope of the yield curve as pricing factors, with the difference that the variance swap factors are in logs, not levels.

In comparison to existing variance swap models, this paper is similar to Aït-Sahalia et al. (2015) and Dew-Becker et al. (2017), who estimate three-factor affine models with two cross-sectional factors from variance swap rates and one time-series factor that is either realized variance or the stock market index. In contrast, Egloff et al. (2010) and Giglio and Kelly (2017) estimate two-factor affine models by assuming that realized variance is spanned by variance swap rates. Empirically, realized variance is only partially spanned by variance swap rates.[2] This observation can motivate including realized variance directly in the model, as in this paper. Beyond model size and setup, the analysis also highlights the outperformance of the logarithmic model to affine models at out-of-sample return prediction. For nearly all combinations of forecast horizons and model sizes, the preferred logarithmic model outperforms reduced form linear return forecasts, often by as much as 10% to 20%.

Similar to previous studies of the variance risk premium, the estimated variance term premia tend to increase during periods of financial distress and decrease during expansions (Bollerslev et al. 2009, Drechsler 2013). This business cycle variation drives the return predictability of the model. Across the variance swap curve, the paper finds that long-end variance swap rates are primarily driven by risk premia whereas short-end variance swap rates are driven by both the quantity and price of volatility risk. Decomposing the variance term premia further, the paper finds that each of the pricing factors contributes significantly to the time-variation in the estimated risk premia with differential effects that change over time. These results reflect the nonlinear nature of the model, providing an alternative perspective

[2]The spanning assumption is common in fixed income settings where short rates are well explained by yield curve principal components. For example, the first three principal components of standardized 1, 2, 3, 5, 7, and 10-year zero-coupon yields from Gürkaynak et al. (2007) span three-month UST bill rates with an $R^2_{adj} = 99.1\%$ from 1980 to 2016. In comparison, the first two principal components of standardized variance swap rates in annualized variance (volatility) units only span realized variance with an $R^2_{adj} = 79.1\%$ (65.5%) from 1996 to 2016 using the synthetic variance swap rates and realized variance estimates from this paper.

to the linear return predictability regressions analyzed in the literature (Van Tassel and Vogt 2016, Johnson 2017).

Beyond forecasting returns, the variance term premia also reflects how investors' pricing of risk has changed over time. Prior to the financial crisis, the term-structure of risk premia was relatively flat on average and the slope sometimes switched signs. After the crisis, long-dated risk premia increased relative to short-dated risk premia and have remained persistently high. These results build on the prior literature which has documented a downward sloping term-structure of unconditional Sharpe ratios for variance swap and straddle returns (Andries et al. 2015, Dew-Becker et al. 2017). Going beyond these unconditional results, the analysis in this paper illustrates how the price of risk for realized variance varies over time and over horizon. According to the model estimates, the price of risk for bearing realized variance shocks over longer horizons has differentially increased in the post-crisis period.

The remainder of the paper proceeds as follows. Section 2 describes the data and presents a model for pricing variance swaps and VIX futures. Section 3 discusses model estimation and presents the variance risk premia estimates. Section 4 reports the return predictability results and examines the relative pricing of variance swaps and VIX futures. Section 5 concludes. The Appendix includes additional details and robustness checks.

2 The Time-Varying Price of Volatility Risk

2.1 Variance Swaps

Variance swaps are over-the-counter derivatives that allow investors to hedge and speculate on volatility over different horizons. The only cashflow occurs at maturity and is equal to the difference between the fixed variance swap rate and the floating amount of realized variance that the underlying asset exhibits over the life of the swap. The fixed rate is priced to make the swap costless to enter at the time of trade. Variance swaps can be interpreted as a form of volatility insurance, with the fixed rate and maturity representing the insurance premium and length of coverage. By trading variance swaps, investors give rise to a term structure of market implied volatility that embeds information about volatility expectations and risk premia over different horizons.

This paper constructs a detailed variance swap dataset that includes daily data for variance swaps written on the S&P 500 Index from 1996 to 2016 on a monthly grid from one-month to two years. The sample is obtained by combining synthetic variance swap rates from 2000 to 2016 with over-the-counter variance swap quotes from 1996 to 2000. The synthetic rates are computed from index option prices using OptionMetrics data. This exploits the

well known no-arbitrage relationship between option prices and variance swap rates (Carr and Wu 2009), leveraging the long time series and rich quantity of available index option data. The over-the-counter quotes are obtained from a hedge fund for the early years in the sample when long maturity index options are less liquid. The Appendix contains a detailed description of the synthetic variance swap rate construction. Overall, the synthetic rates closely align with the hedge fund quotes both in their time series variation and levels. In addition, the synthetic rates closely match over-the-counter variance swap quotes from Markit Totem as well as the volatility indexes from the CBOE.

2.2 Realized Variance

The floating leg of a variance swap pays the realized variance of the underlying asset from the trade date until the maturity of the swap. To make this feasible in practice, variance swaps need to specify a definition for computing realized variance. Contracts can differ on this dimension. For example, variance swaps must specify whether to use log or simple returns, whether to demean returns or not, how to annualize estimates using different day count conventions, etc. From a theoretical perspective, it is desirable to choose a definition that produces an accurate estimate of the quadratic variation of the underlying asset. This follows from the no-arbitrage replication argument for pricing variance swaps, which relies on computing the risk-neutral expectation of an asset's quadratic variation through an application of Itô's lemma.

Based on these observations, I define the realized variance payoff for my empirical application using the two-scale realized variance estimator from Zhang et al. (2005). I compute the two-scale estimator using one-minute high frequency data for the S&P 500 Index from Thomson Reuters Tick History (TRTH). This choice reflects the trade-offs in using high frequency data to estimate realized variance. On one hand, sampling more finely allows for more accurate volatility estimation (Merton 1980). On the other hand, sampling too finely can magnify microstructure noise such as bid-ask spread and price discreteness which can severely bias estimation. The two-scale estimator balances these trade-offs by averaging realized variance estimates from a sparse sampling frequency across subsamples on a finer grid. For the application in this paper, I compute first stage realized variance estimates that are equal to the sum of squared five-minute intraday log returns plus the squared overnight log return for each day in the sample. The choice of a five-minute intraday sampling frequency is common in the empirical literature and motivated by Liu et al. (2015). I then average the first stage realized variance estimates across one-minute subsamples to reduce sampling variability, resulting in a second stage daily estimate of realized variance. The monthly payoff

to a variance swap is defined as the sum of the second stage daily realized variance estimates every 21 business days. For a more detailed description of the estimation approach and an outline of the steps used for cleaning the high frequency data, see the Appendix.

2.3 Pricing Variance Swaps and VIX Futures

Variance swaps are modeled as the expected value of future realized variance under the risk-neutral measure $\mathbb{Q}$ from the trade date until the maturity of the swap,

$$VS_{t,n} = E_t^{\mathbb{Q}} \left[\sum_{i=1}^{n} RV_{t+i} \right]. \tag{1}$$

Time is discrete with each period representing one month.[3] To model variance swap dynamics, I assume the systematic risk in the economy can be summarized by a $K \times 1$ vector of state variables X_t that follows a stationary vector autoregression under the physical measure $\mathbb{P}$,

$$X_{t+1} = \mu + \Phi X_t + v_{t+1}, \tag{2}$$

with shocks v_{t+1} that are conditionally Normal $v_{t+1}|\mathcal{F}_t \overset{\mathbb{P}}{\sim} N(0, \Sigma_v)$. This specification can be motivated by the intertemporal capital asset pricing model (ICAPM) of Merton (1973) or the arbitrage pricing theory (APT) of Ross (1976). I set the state vector equal to,

$$X_t' = [\ln RV_t \; Y_t']. \tag{3}$$

The first element is the logarithm of realized variance $\ln RV_t$ which spans variance payoffs. The subsequent variables Y_t can be any financial or macroeconomic variables that help to price the cross section of variance swap rates or explain the time series variation of variance swap returns.

To model risk premia and derive variance swap rates, I assume the stochastic discount factor is equal to,

$$M_{t+1} = e^{-r_t - \frac{1}{2}\lambda_t'\lambda_t - \lambda_t'\Sigma_v^{-1/2} v_{t+1}}, \tag{4}$$

with an affine price of risk,

$$\lambda_t = \Sigma_v^{-1/2} \left(\Lambda_0 + \Lambda_1 X_t \right). \tag{5}$$

[3]The Appendix presents an analogous model in continuous time. While variance swap pricing remains tractable in continuous time, closed form solutions for VIX futures are only available in discrete time. This distinction and the ease of estimation for the discrete time model are advantages for the approach adopted in the paper.

This links the physical and risk-neutral dynamics through the relationships $\mu^{\mathbb{Q}} = \mu - \Lambda_0$ and $\Phi^{\mathbb{Q}} = \Phi - \Lambda_1$ with the state vector under the risk-neutral measure $\mathbb{Q}$ following,

$$X_{t+1} = \mu^{\mathbb{Q}} + \Phi^{\mathbb{Q}} X_t + v_{t+1}^{\mathbb{Q}}, \tag{6}$$

with shocks that are conditionally normal $v_{t+1}^{\mathbb{Q}} | \mathcal{F}_t \overset{\mathbb{Q}}{\sim} N(0, \Sigma_v)$.

In deriving variance swap rates, it is convenient to first obtain prices for variance swap forwards. Variance swap forwards are defined as,

$$F_{t,n} = E_t^{\mathbb{Q}} \left[RV_{t+n} \right], \tag{7}$$

where the zero-month forward rate is equal to the current value of realized variance. Variance swap forwards decompose the variance swap curve into one-month swap rates with forward starting dates,

$$VS_{t,n} = \sum_{i=1}^{n} F_{t,i}, \tag{8}$$

similar to the relationship between forward rates and yields in fixed income.

The excess return from receiving fixed in variance swap forwards is equal to,

$$Rx_{t+1,n} = F_{t,n} - F_{t+1,n-1}. \tag{9}$$

This trade corresponds to receiving fixed in an n month variance swap forward at time t and paying fixed in an $n-1$ month variance swap forward at time $t+1$. Since this trade costs zero dollars, it is equivalent to the risk-neutral pricing equation,

$$E_t^{\mathbb{Q}} \left[F_{t,n} - F_{t+1,n-1} \right] = 0. \tag{10}$$

Put differently, the risk-neutral expected value from trading variance swaps is zero. Variance swap forwards are a martingale under the risk-neutral measure.[4]

To derive variance swap rates, I guess and verify that variance swap forwards are exponential affine in the state vector,

$$F_{t,n} = e^{A_n + B_n' X_t}. \tag{11}$$

I set the initial condition to $A_0 = 0$ and $B_0 = [1 \, \vec{0}]$ so that the model prices realized variance exactly. This restriction reduces the number of parameters to estimate resulting in a more parsimonious model. The risk-neutral pricing equation for the one-month variance swap rate

[4]This follows from the common assumption in the variance swap literature that interest rates are deterministic or independent from realized variance. For example, see Carr and Wu (2009), Egloff et al. (2010), Aït-Sahalia et al. (2015), Filipović et al. (2016), and Dew-Becker et al. (2017) .

is thus equal to,

$$\begin{aligned} E_t^{\mathbb{Q}}\left[Rx_{t+1,1}\right] &= E_t^{\mathbb{Q}}\left[F_{t,1} - F_{t+1,0}\right] \\ &= E_t^{\mathbb{Q}}\left[VS_{t,1} - RV_{t+1}\right] \\ &= e^{A_1 + B_1' X_t} - e^{A_0 + B_0'\left(\mu^{\mathbb{Q}} + \Phi^{\mathbb{Q}} X_t\right) + \frac{1}{2} B_0' \Sigma_v B_0} \\ &= 0. \end{aligned} \tag{12}$$

Since this equation must hold state by state, matching coefficients determines A_1 and B_1. For longer maturities, plugging the guess into the risk-neutral pricing equation produces the following system of well-known recursive equations,

$$\begin{aligned} A_n &= A_{n-1} + B_{n-1}'\mu^{\mathbb{Q}} + \tfrac{1}{2} B_{n-1}' \Sigma_v B_{n-1} \\ B_n' &= B_{n-1}' \Phi^{\mathbb{Q}}. \end{aligned} \tag{13}$$

These recursions coupled with the initial condition determine variance swap forward rates.

Variance swap rates are equal to the sum of variance swap forward rates,

$$VS_{t,n} = \sum_{i=1}^{n} e^{A_i + B_i' X_t}. \tag{14}$$

The adjustment $\sqrt{12/n \cdot VS_{t,n}}$ expresses variance swap rates in annualized volatility units. To compute realized variance forecasts and the variance term premia in the model, variance swap rates can be decomposed as,

$$VS_{t,n} = \underbrace{E_t^{\mathbb{P}}\left[\sum_{i=1}^{n} RV_{t+i}\right]}_{RVF_{t,n} \equiv \text{Realized Variance Forecast}} + \underbrace{\left(E_t^{\mathbb{Q}}\left[\sum_{i=1}^{n} RV_{t+i}\right] - E_t^{\mathbb{P}}\left[\sum_{i=1}^{n} RV_{t+i}\right]\right)}_{VTP_{t,n} \equiv \text{Variance Term Premium}}. \tag{15}$$

The variance term premia are equal to the expected holding period return from receiving fixed in variance swaps over an n-month horizon. I compute variance term premia by subtracting the realized variance forecasts form variance swap rates,

$$VTP_{t,n} = \textstyle\sum_{i=1}^{n} e^{A_i + B_i' X_t} - \sum_{i=1}^{n} e^{A_i^P + (B_i^P)' X_t}. \tag{16}$$

The realized variance forecasts are obtained by replacing $\mu^{\mathbb{Q}}$ and $\Phi^{\mathbb{Q}}$ with μ and Φ in the recursions above to compute the coefficients A_n^P and B_n^P. This shuts down the prices of risk, allowing for forecasts under the physical as opposed to the risk-neutral measure.

Beyond pricing variance swaps, the model also admits closed form prices for VIX futures.

This is an advantage of modeling the logarithm of realized variance. The exponential affine price for variance swap forwards naturally absorbs the convexity adjustment. To see this, define the VIX as,

$$VIX_t \equiv \sqrt{E_t^{\mathbb{Q}}\left[RV_{t+1}\right]} = \sqrt{VS_{t,1}}. \tag{17}$$

It follows that the price of the n-month VIX futures contract is,

$$\begin{aligned} Fut_{t,n} &= E_t^{\mathbb{Q}}\left[VIX_{t+n}\right] \\ &= E_t^{\mathbb{Q}}\left[\sqrt{E_{t+n}^{\mathbb{Q}}\left[RV_{t+n+1}\right]}\right] \\ &= E_t^{\mathbb{Q}}\left[\sqrt{e^{A_1+B_1'X_{t+n}}}\right] \\ &= E_t^{\mathbb{Q}}\left[e^{\frac{1}{2}A_1+\frac{1}{2}B_1'X_{t+n}}\right] \\ &= e^{A_n^F+(B_n^F)'X_t}. \end{aligned} \tag{18}$$

The coefficients A_n^F and B_n^F for pricing VIX futures follow the same recursions as A_n and B_n for pricing variance swaps with an adjusted initial condition $A_0^F = \frac{1}{2}A_1$ and $B_0^F = \frac{1}{2}B_1$. To express VIX futures prices in annualized volatility units, simply multiply the formula above by $\kappa = 100 \cdot \sqrt{12}$.[5] Finally, in addition to pricing VIX futures, the model also provides volatility swap rates, option prices for VIX futures, and bounds on VIX futures prices which are included in the Appendix.

3 Variance Term Premia Estimation

3.1 Model Estimation

I estimate the model using daily observations of variance swap rates from 1996 to 2016 for $\tau = \{1, 3, 6, 9, 12, 18, 24\}$ month maturities. Adjusting the notation slightly to allow for daily data, the model can be summarized by the following system of equations,

$$\begin{aligned} X_{t+h} &= \mu + \Phi X_t + v_{t+h}, & v_{t+h}|\mathcal{F}_t \sim N(0, \Sigma_v) \\ Y_{t,n} &= g_n(X_t, \mu^{\mathbb{Q}}, \Phi^{\mathbb{Q}}, \Sigma_v) + e_{t,n}, & E[e_{t,n}|X_t] = 0. \end{aligned} \tag{19}$$

The state vector X_t follows a monthly vector autoregression with overlapping observations and a horizon of $h = 21$ trading days. Variance swap rates $Y_{t,n}$ are observed with measurement errors $e_{t,n}$ that are mean zero conditioned on the state vector. The model prices expressed in annualized volatility units are,

[5] In practice the VIX is defined as $VIX_t = 100 \cdot \sqrt{12 \cdot E_t^{\mathbb{Q}}\left[RV_{t+1}\right]}$.

$$g_n(X_t, \mu^{\mathbb{Q}}, \Phi^{\mathbb{Q}}, \Sigma_v) = \sqrt{\frac{12}{n}\sum_{i=1}^{n} e^{A_i + B_i' X_t}}. \tag{20}$$

The parameters to be estimated are $\Theta = (\mu, \Phi, \mu^{\mathbb{Q}}, \Phi^{\mathbb{Q}}, L_v)$ where L_v is the Cholesky decomposition of $\Sigma_v = L_v L_v'$.

Estimating the model with overlapping daily data has several advantages relative to month-end data. Daily data increases the sample size to allow for increased precision when estimating the parameter values. Daily data is also more demanding of the model, as some of the most extreme observations of variance swap rates occur within the month, not at month-end.[6] Finally, estimating the model with daily data allows me to compute model prices and expected returns at a daily frequency, which is useful for analyzing trading strategies and for performing high frequency event studies.

I estimate the model in two steps. First, I estimate the physical parameters $(\hat{\mu}, \hat{\Phi}, \hat{\Sigma}_v)$ from a monthly vector autoregression with overlapping observations,

$$X_{t+h} = \hat{\mu} + \hat{\Phi} X_t + \hat{v}_{t+h}. \tag{21}$$

Second, I estimate the risk-neutral parameters $(\hat{\mu}^{\mathbb{Q}}, \hat{\Phi}^{\mathbb{Q}})$ by minimizing the model's variance swap pricing errors by nonlinear least squares,

$$(\hat{\mu}^{\mathbb{Q}}, \hat{\Phi}^{\mathbb{Q}}) = \arg\min_{(\mu^{\mathbb{Q}}, \Phi^{\mathbb{Q}})} \frac{1}{T \cdot N_\tau} \sum_{t=1}^{T} \sum_{n \in \tau} \left(Y_{t,n} - g_n(X_t, \mu^{\mathbb{Q}}, \Phi^{\mathbb{Q}}, \hat{\Sigma}_v) \right)^2. \tag{22}$$

This two-step approach easily accommodates daily data and is robust to assumptions about the distribution of the variance swap measurement errors. The Appendix considers alternative estimation approaches as a robustness check. Maximum likelihood and Bayesian methods with latent factors deliver similar results.

The state variables in the empirical implementation include log realized variance and the first K_{PC} principal components of log variance swap rates. For numerical stability, I standardize log realized variance and log variance swap rates before computing the principal components. This changes the initial condition for pricing variance swaps to $A_0 = \mu_{\ln RV}$ and $B_0 = [\sigma_{\ln RV}\ \vec{0}]$. I omit the standardization going forward for notational simplicity.[7] In addition to estimating the parameters Θ, the number of principal components K_{PC} must

[6] For example, during the financial crisis in the fall of 2008, the five highest closing values of the VIX were 80.86, 80.06, 79.13, 74.26, and 72.67 on 11/20, 10/27, 10/24, 11/19, and 11/21, none of which are month-end dates. The month-end observations were 59.89 on 10/31 and 55.28 on 11/28.

[7] The state vector is $X_t = [(\ln RV_t - \mu_{lnRV})/\sigma_{\ln RV}\ PC_{1,t} \ldots PC_{K_{PC},t}]$ where the principal components are constructed from the standardized logarithm of variance swap rates.

also be selected. I discuss model selection along with the estimation results below.

3.2 Estimation Data

Table 1 reports summary statistics for the realized variance and variance swap rate data that are used to estimate the model. Figure 1 illustrates the data by plotting realized variance and one-month variance swap rates against the market return. The term structure of volatility is upward sloping on average with realized variance equal to 15.2% and one-month variance swap rates equal to 20.8% in annualized volatility units. This gap reflects the significant unconditional variance risk premium that investors earn by receiving fixed in variance swaps. The average one-year and two-year variance swap rates are even higher at 22.6% and 23.4%. Long dated variance swap rates also have higher one-month and six-month autocorrelations relative to short dated variance swap rates and realized variance. This larger autocorrelations further out on the curve reflect the mean reversion of realized variance. During periods of high stock market volatility, the variance swap curve tends to invert with volatile short-dated rates increasing more than persistent long-dated rates.

Table 1 also reports summary statistics for monthly variance swap returns in percentage units. Variance swap returns are defined by receiving fixed in an n month swap at time t and paying fixed in an $n-1$ month swap at time $t+1$,

$$R_{t+1,n} = VS_{t,n} - RV_{t+1} - VS_{t+1,n-1}. \tag{23}$$

This payoff is an excess return as it costs zero dollars at time t. Table 1 shows that the mean and standard deviation of variance swap returns are increasing in maturity, while the Sharpe ratio and t-statistic are decreasing. These results are consistent with the prior literature on the unconditional term structure of variance swap returns, which highlights how investors demand a larger premium for being exposed to realized variance shocks rather than implied volatility shocks over short horizons (Dew-Becker et al. 2017, Andries et al. 2015). The annualized Sharpe ratio from receiving fixed in one-month variance swaps is 1.82 and the t-statistic for the CAPM alpha is 8.19.[8]

Beyond the significant returns from receiving fixed in short dated variance swaps, the results highlight how variance swap returns are negatively skewed and positively correlated

[8]The high Sharpe ratio and t-statistic in part reflect the use of high frequency data to estimate realized variance RV_t. Using squared daily log returns to define the floating leg payoff lowers the Sharpe ratio to .93 and t-statistic to 4.21. As discussed before, I use high frequency data to estimate realized variance RV_t as this more accurately estimates the quadratic variation that variance swaps are designed to price. Whether investors can capture these returns depends on the setting. By frequently delta-hedging, an option trader may better approximate the theoretical returns with continuous hedging (Bertsimas et al. 2000).

with the market. The CAPM betas are significant and increasing in maturity. The market factor explains about 40% of the variation in variance swap returns for maturities longer than one-month. The positive and significant betas reflect how increases in volatility are negatively correlated with stock market returns, the so-called leverage effect (Black 1976). The bottom plot in Figure 1 illustrates this result by plotting the variance swap returns against CRSP value-weighted market returns. One-month and twelve-month variance swap returns are 51% and 66% correlated with the market at a monthly frequency. As a final observation, the percentage of negative variance swap returns is only 10% at the one-month maturity and 20% at the three-month maturity. The low frequency of negative returns at the short end of the curve supports the interpretation of variance swaps as a form of volatility insurance. In most periods, volatility is low and a premium is collected. However, occasional spikes in volatility can result in large losses and negatively skewed returns.

3.3 Model Selection

Table 2 reports variance swap pricing errors and return forecast errors for alternative specifications that vary the number of principal components in the state vector K_{PC}. The results are averaged across maturities from 1998 to 2016 using 1996 to 1998 as an initial estimation period for the expanding window out-of-sample analysis. A three factor model,

$$X_t = [\ln RV_t \; PC_{level,t} \; PC_{slope,t}], \tag{24}$$

with two principal components $K_{PC} = 2$ computed from the logarithm of variance swap rates performs well relative to the competing models. The principal components can be interpreted as level PC_{level} and slope PC_{slope} factors that explain over 99% of the variation in log variance swap rates.

To see the outperformance of the three-factor logarithmic model, Panels A.I and B.I show that adding the slope factor significantly reduces the in-sample and out-of-sample variance swap pricing errors as measured by either the root-mean-squared-error (RMSE) or the mean-absolute-error (MAE). The decrease in pricing errors from adding the slope factor is approximately .50% to .75% in annualized volatility units, roughly the same size as the average bid-ask spread in the variance swap market according to Markit data. Adding further principal components continues to lower the pricing errors, but with smaller gains. In contrast to the pricing errors, Panels A.II and B.II show that adding additional factors beyond slope actually leads to similar in-sample return predictability and lower out-of-sample return predictability. This suggests that most of the return predictability is being driven by realized variance and the level of variance swap rates. Overall, the results motivate selecting

the three-factor model for its low variance swap pricing errors and significant return forecasts both in-sample and out-of-sample. That said, one could argue for a four-factor model on the grounds of its marginally lower variance swap pricing errors and similar return predictability results. In the interest of parsimony, I select the smaller three-factor model as the baseline specification for the subsequent analysis.

To provide further competition for the three-factor logarithmic model, Table 2 also compares the model's return predictability to reduced form linear forecasts that use realized variance and up to five principal components from the level of variance swap rates as in Van Tassel and Vogt (2016). Panel A.III shows that the unrestricted linear forecasts provide good in-sample fit that improves with the number of factors. At a one-month horizon, Panel A.IV shows that the linear models outperform by as much as 5% to 10% as measured by the in-sample mean squared forecast error. However, Panels A.IV and A.V show that the linear model does not outperform for longer horizons or for a mean-absolute-error criterion. Moreover, the performance of the linear model deteriorates significantly out-of-sample. Panel B.III shows that the out-of-sample explanatory power R^2_{oos} of the linear model is close to zero or negative at a one-month horizon. For longer forecast horizons, the linear model's out-of-sample performance decreases as the number of factors increases. These results suggest that the linear model is unstable and that the larger linear models are overfitting in-sample. In contrast, Panel B.II shows that the three-factor logarithmic model has positive out-of-sample explanatory power R^2_{oos} at all horizons. In addition, Panels B.IV and B.V show that the three-factor logarithmic model outperforms the linear models at out-of-sample return prediction for nearly all combinations of forecast horizons and model sizes according to either a mean-squared-error or mean-absolute-error criterion, sometimes by as much as 10% to 20%. The relative stability of the three-factor logarithmic model and its superior out-of-sample performance provide empirical support for the decision to model the logarithm as opposed to the level of realized variance.

3.4 Variance Term Premia Estimates

Figure 2 plots the time-varying price of volatility risk as measured by the one-month and twelve-month variance term premia estimates $VTP_{t,n}$ in the three factor model with $K_{PC} = 2$ principal components. Variance term premia represent the term-structure of expected holding period returns from receiving fixed in variance swaps. The plot reveals several interesting features of variance term premia. First, variance term premia can be as high as 5% to 10% during periods of financial distress such as the Asian financial crisis, LTCM crisis, financial crisis, and European sovereign debt crisis. This contrasts the unconditional

one-month and twelve-month variance term premia of 2.20% and 2.70% over the 1996 to 2016 sample period. Second, the plot indicates that the twelve-month term premium is more persistent than the one-month term premium. After a negative shock, long dated variance term premia tend to remain elevated while short dated variance term premia mean revert more quickly. Finally, the term structure of variance term premia has changed since the financial crisis. Prior to the crisis, the one-month and twelve-month term premia had a similar magnitude, sometimes below or above each other. After the crisis, it appears that investors repriced the variance term premia, with long-dated term premia now consistently higher than short-dated term premia.[9]

Variance term premia exhibit significant time variation that is driven by each of the state variables.[10] Table 3 reports the estimated model parameters. To measure significance, the table reports Newey-West t-statistics for the physical parameters using 50 lags to account for the overlapping monthly observations and block bootstrapped t-statistics for the prices of risk that take into account both the overlapping observations and the sampling uncertainty from the first step estimation of the physical parameters. The results indicate that each of the state variables contributes significantly to the time variation in the realized variance forecasts and variance term premia estimates. The mean of the physical parameters $\hat{\mu}$ is close to zero reflecting how the state variables are standardized. The first row of $\hat{\Phi}$ shows that higher levels of the state variables forecast higher levels of log realized variance. The second and third rows show that the level and slope factors are relatively uncorrelated with the other variables, which effectively leaves them following their own first order autoregressions. Meanwhile, the prices of risk in the first row of $\hat{\Lambda}_1$ are all significant, indicating that each of the state variables contributes to the time variation in the variance term premia. Interpreting the price of risk estimates $\hat{\Lambda}_0$ and $\hat{\Lambda}_1$ beyond these observations is somewhat challenging because the model is nonlinear. Instead, I present a decomposition below that shows how changes in realized variance and the first two principal components of variance swap rates are related to changes in variance term premia.

Before discussing the variance term premia estimates further, it is also important to inspect the model fit. Table 4 reports the model fitting errors for variance swap rates and returns. Figure 3 plots the model variance swap rates against the observed rates. The mean and standard deviation of the variance swap fitting errors are -.003% and .36% in

[9]From 1996 to 2006 the average one-month and twelve-month term premia were 2.07% and 2.12%, with the one-month term premia above the twelve-month term premia on 50% of days. From 2010 to 2016 the average one-month and twelve-month term premia were 1.65% and 2.94%, with the one-month term premia above the twelve-month term premia on only 5% of days.

[10]In the Appendix, Figure A.4 adds 95% pointwise confidence intervals to the term premia estimates. The movement in variance term premia is significantly larger than the confidence bands.

annualized volatility units averaged across maturities. This magnitude is small relative to the average bid-ask spread of .77% reported by Markit from September 2006 to December 2015 on month-end dates, but somewhat high relative to the .05% tick size for highly liquid VIX futures contracts. To interpret these results in different units, the model also exhibits small pricing errors for variance swap returns.[11] For example, the standard deviations of the one-month and twelve-month return pricing errors are 3.1 and 19.3 basis points in contrast to an unconditional standard deviation of 33 and 161 basis points. Beyond the size of the pricing errors, the results also indicate that the pricing errors are persistent and fat-tailed as measured by the one-month and six-month autocorrelations and excess kurtosis.

3.5 Variance Swap Return Predictability

Table 5 investigates whether the model expected variance swap returns predict realized returns by running the regressions,

$$R_{t+h,n} = \beta_0 + \beta_1 \hat{E}_t[R_{t+h,n}] + \epsilon_{t+h,n}. \tag{25}$$

The dependent variable $R_{t+h,n}$ is the excess return from receiving fixed in an n-month variance swap over an h-month horizon. The independent variable is the estimated expected return from the model $\hat{E}_t[R_{t+h,n}]$. For example, one-month expected returns are equal to

$$\begin{aligned} E_t[R_{t+1,n}] &= E_t[VS_{t,n} - RV_{t+1} - VS_{t+1,n-1}] \\ &= VS_{t,n} - E_t\left[\textstyle\sum_{i=0}^{n-1} e^{A_i + B_i' X_{t+1}}\right] \\ &= VS_{t,n} - \textstyle\sum_{i=0}^{n-1} e^{A_i + B_i'(\mu + \Phi X_t) + \frac{1}{2} B_i' \Sigma B_i}. \end{aligned} \tag{26}$$

I also compute three-month and six-month expected returns. It is not immediate that these expected return estimates will be significant at forecasting realized returns. Recall that the model is estimated by minimizing variance swap pricing errors, not by minimizing return forecast errors. Despite this, the results indicate that the model's expected returns are significant at forecasting realized returns over all horizons h and maturities n with an average explanatory power of 17%, 21%, and 30% for one-, three-, and six-month horizons as measured by the in-sample R^2_{adj}. Relating this back to the model selection analysis, these numbers closely match the results from Panel A.II in Table 2. As Panel A.III indicates, the model also provides significant explanatory power out-of-sample R^2_{oos} equal to 5%, 14%, and

[11] Variance swap return errors are defined as $u_{t+1,n} = R_{t+1,n} - (\hat{VS}_{t,n} - \hat{RV}_{t+1} - \hat{VS}_{t+1,n-1})$ where $R_{t+1,n}$ is the observed return for the n-month swap rate and $(\hat{VS}_{t,n}, \hat{RV}_{t+1}, \hat{VS}_{t+1,n-1})$ are the estimated model prices. Note that $\hat{RV}_{t+1} = RV_{t+1}$ as the model prices realized variance exactly.

24% averaged across maturities.

Figure 4 illustrates the return predictability by plotting the one-month and six-month expected returns alongside realized returns over the subsequent horizons.[12] The explanatory power of the estimates are 23% and 35% as measured by in-sample R^2_{adj}. As the plot makes clear, variance swap return predictability is driven in part by the periods of financial distress. On one hand, the onset of distress tends to result in large forecast errors, decreasing the explanatory power of the model expected returns. On the other hand, distress tends to be followed by periods with persistently high realized returns that match the model's variance term premia estimates. In addition to periods of distress, the term premia also match the relatively low but positive returns during the mid-2000s as well as the low but positive returns in recent years.

Figure 4 also illustrates how the model's variance term premia and realized variance forecasts have evolved over time using area plots. The bottom area in blue represents the realized variance forecast while the remaining area in red represents the variance term premia. The area plots indicate that the long horizon realized variance forecasts are more persistent than short dated realized variance forecasts, reflecting the mean-reversion of realized variance under the physical measure. An implication of this result is that movements in long dated variance swap rates primarily reflect changes in term premia rather than volatility forecasts, whereas movements in short dated variance swap rates reflect both changes in term premia and volatility forecasts. Table 6 quantifies this observation by decomposing the variance of variance swap rates into percentage contributions from variance term premia and realized variance forecasts over different horizons. Panel A shows that the contribution from the realized variance forecasts decreases from 58% at the one-month maturity to 20% at the two-year maturity, while the contribution from variance term premia increases from 42% at the one-month maturity to 80% at the two-year maturity. While the volatility term structure only extends out to two years in calendar time, these results suggest that two years may be a long amount of economic time from the perspective of stock market volatility. For example, because realized variance mean reverts much faster than interest rates, it is not immediate that a two-year volatility term structure is shorter than a thirty-year fixed income term structure.

Reverting back to the price of risk estimates, one challenge the nonlinear model poses is interpreting the point estimates $\hat{\Lambda}_0$ and $\hat{\Lambda}_1$ in Table 3. To understand the importance of the different pricing factors in driving changes in the variance term premia, Panel B of Table 6 provides perspective from a linear model by regressing the monthly change in the variance

[12]Note that the one-month and six-month variance term premia are equal to one-month expected returns for a one-month variance swap and six-month expected returns for a six-month variance swap.

term premia estimate from the nonlinear model onto z-scored changes in realized variance and the first two principal components from the level of variance swap rates. The results indicate that each variable is significant in explaining the changes in variance term premia on average, consistent with the interpretation of the significant price of risk estimates.

Of course, the partial derivatives of the variance term premia with respect to the state variables will change depending on the level of the state vector,

$$\nabla VTP_{t,n} = \frac{12}{n}\left(\sum_{i=1}^{n} B_i \cdot e^{A_i + B_i' X_t} - \sum_{i=1}^{n} B_i^P \cdot e^{A_i^P + (B_i^P)' X_t}\right). \tag{27}$$

Figure 5 investigates this observation by plotting the partial derivatives of the variance term premia for one standard deviation moves in the state variables at different points in time.[13] Similar to the regressions, the top left subplot reports the average partial derivative over the sample period by maturity. The results are similar to the regressions. An increase in realized variance decreases the variance term premia with a magnitude that is larger at the short end of the curve. An increase in the level of variance swap rates increases variance term premia in a roughly parallel manner. An increase in the slope of variance swap rates increases term premia at the short end and decreases term premia at the long end.

The other subplots illustrate how the partial derivatives can change across dates. The top right plot shows the partial derivatives in October 2007 when the state vector $X_t = [-.02 \ -.05 \ -.07]$ was close to its mean $\hat{\mu}$ under the physical measure. The shape of these derivatives is similar to the average derivatives, but with magnitudes that are somewhat lower. The bottom right plot shows the partial derivatives during the financial crisis in November 2008, a period with high realized and implied volatility and an inverted variance swap curve $X_t = [4.07\ 7.83\ .71]$. The magnitude of the partial derivatives during the crisis was much larger than average, with an inverted partial derivative for the level factor. The gray box in this subplot highlights the scale of the other subplots, illustrating the increase in magnitude during the financial crisis. Finally, the bottom right subplot reports the partial derivatives in December 2016, a period of low volatility with an upward sloping variance swap curve. In that plot, the level factor partial derivatives are upward sloping and the realized variance partial derivatives are roughly parallel. Overall the analysis highlights how the linear model provides a good approximation to the average partial derivatives, capturing 94% to 98% of the variation in variance term premia. The exact partial derivatives highlight how the sensitivity of the variance term premia to the state variables can change over time.

[13]The sample standard deviation of the state variables are $\sigma((\ln RV - \mu_{\ln RV})/\sigma_{\ln RV}) = 1$, $\sigma(PC_{level}) = 2.58$, and $\sigma(PC_{slope}) = .55$.

4 Applications and Discussion

4.1 Straddle Return Predictability

The returns from selling delta-hedged straddles are closely related to variance swap returns. A short straddle refers to a position that is short a call option and a put option with the same strike and maturity. A delta-hedge removes the directional exposure to the underlying. The combination of selling a straddle and delta-hedging produces a return that is largely determined by the relationship of implied to realized volatility. In Black-Scholes parlance, delta-hedged short straddle positions have positive theta Θ, negative gamma Γ, and negative vega ν,

$$\begin{aligned} dF &\approx F_S dS + \tfrac{1}{2} F_{SS} dS^2 + F_t dt + F_\sigma d\sigma \\ &= \tfrac{1}{2}\Gamma dS^2 + \Theta dt + \nu d\sigma. \end{aligned} \tag{28}$$

When realized volatility dS^2 is low and implied volatility $d\sigma$ does not increase, short straddle positions are profitable as option writers earn the theta Θdt or carry. When realized volatility dS^2 is high or implied volatility $d\sigma$ increases, short straddle positions can suffer losses. One can draw an analogy to variance swap returns where,

$$\begin{aligned} R_{t+1,n} &= VS_{t,n} - VS_{t+1,n-1} - RV_{t+1} \\ &= \underbrace{(VS_{t,n} - VS_{t,n-1})}_{\Theta dt} + \underbrace{(VS_{t,n-1} - VS_{t+1,n-1})}_{\nu d\sigma} + \underbrace{(-RV_{t+1})}_{\frac{1}{2}\Gamma dS^2}. \end{aligned} \tag{29}$$

Beyond this informal connection there is also a close theoretical relationship between straddle returns and the variance risk premium. Andries et al. (2015) show that delta-hedged straddle returns provide a non-parametric estimate of the variance risk premium in the absence of jumps in the underlying asset. This suggests a natural application for the model: testing whether expected variance swap returns are significant in predicting delta-hedged straddle returns.

To perform this test, I compute delta-hedged straddle returns at a daily frequency from the straddle mid-price S_t, index value P_t, strike price K, and straddle delta Δ_t for all strike-maturity pairs whose delta Δ_t is less than 25% in absolute value. The daily returns are defined as,

$$R_{t+1}^{straddle} = \frac{S_t - S_{t+1} - \Delta_t(P_{t+1} - P_t)}{K}. \tag{30}$$

I then average these returns for all strike-maturity pairs in different maturity buckets to

obtain a term structure of daily returns.[14] For example, the $(1, 3]$ month maturity bucket has an average of 14.5 straddle return observations per day across 2 expirations while the $(9, 15]$ month maturity bucket has an average of 8.8 straddle return observations per day across 1.6 expirations. As a final step, I aggregate the daily returns for each maturity bucket over one, three, and six-month horizons which can be compared to the model's expected variance swap returns.

Table 7 provides a summary comparison of these delta-hedged straddle returns to the synthetic variance swap returns from 1996 to 2016. Panel A begins by reporting the correlation between the straddle returns and variance swap returns at a monthly frequency by maturity. The returns are highly correlated overall with an average pairwise correlation of 80%. This confirms the motivating discussion that straddle returns are closely related to variance swap returns. Moreover, the returns are most highly correlated for similar maturities. Straddle returns for maturity buckets $(1, 3]$ and $(3, 6]$ months are more highly correlated with 3 to 6 month variance swap returns rather than 12 to 24 month variance swap returns. Similarly, straddle returns for maturity buckets $(9, 15]$ and $(15, 24]$ months are more highly correlated with 12 to 24 month variance swap returns rather than 1 to 6 month variance swap returns.

Panels B builds on this analysis by reporting summary statistics for the one-month straddle returns. Selling straddles in the $(1, 3]$ month maturity bucket delivers an average return of 31 basis points per month with a monthly volatility of 1.39 percent. This corresponds to a Sharpe ratio of .22 per month which is comparable to the Sharpe ratio of .25 and .19 for three-month and six-month variance swap returns in Table 1. As with variance swap returns, Panel B indicates that the Sharpe ratios decline with maturity and that straddle returns are negatively skewed and positively autocorrelated at a one-month frequency.

Panel C then reports factor regressions that explain straddle returns with maturity-matched variance swap returns over different horizons. The explanatory power as measured by the R^2_{adj} is around 76% across maturities and horizons, consistent with the high correlation of straddle and variance swap returns in Panel A. This contrasts the CAPM which only explains one-month straddle returns with an average R^2_{adj} of 16% across maturities (unreported). In addition to explanatory power, the intercepts reveal that average straddle and variance swap returns are similar after adjusting for risk. While the straddle intercepts are sometimes negative and significant, the magnitude is usually less than 10 basis per month. The intercepts are also insignificant after averaging across maturities at a one-month and three-month horizon.

[14] I filter the data to only include option prices for standard expiration dates that have a positive bid, offer, implied volatility, and open interest amount. Straddle deltas are computed from OptionMetrics put and call deltas. As a robustness check, Table A.4 in the Appendix reports the return predictability regressions using alternative definitions for straddle returns.

Overall, Table 7 demonstrates that the straddle returns and variance swap returns are closely related. Given the model's ability to forecast variance swap returns, a natural next step is to explore whether the model can also forecast straddle returns. Table 8 reports return predictability regressions to examine this question,

$$R^{straddle}_{t+h,b} = \beta_0 + \beta_1 \hat{E}_t[R_{t+h,n}] + \epsilon_{t+h,b}. \tag{31}$$

The delta-hedged straddle return $R^{straddle}_{t+h,b}$ for maturity bucket b over horizon h is regressed onto the estimated variance swap expected return $\hat{E}_t[R_{t+h,n}]$ for maturity n. The results indicate that the variance swap expected returns are significant in predicting the delta-hedged straddle returns across all maturity buckets and horizons. Despite estimating the expected returns in the model from variance swap and realized variance data, not straddle returns, the model still provides significant predictive power that can be as high as 10% to 15% for the straddle returns.

4.2 Relative Pricing of Variance Swaps and VIX Futures

The model also provides relative prices for variance swaps and VIX futures. As before, the model prices for VIX futures can be interpreted as the fair value implied by variance swaps. Comparing the model prices to observed prices thus provides a quantitative measures of volatility market integration.

Table 9 and Figure 6 summarize the model's pricing errors for the front six VIX futures contracts from 2007 to 2016.[15] The top left subplot in Figure 6 shows that the model closely tracks the front month futures price with a RMSE (MAE) of .90% (.57%). Pricing errors for the other contracts are similar in magnitude, with an average RMSE of .99% as reported in Table 9. The magnitude of the pricing errors is small in comparison to the 8.81% standard deviation of the front month VIX futures contract, but large in comparison to the bid-ask spreads of VIX futures which have a minimum tick size of .05%.

The bottom left plot in Figure 6 provides further analysis by plotting the unconditional term structure of VIX futures prices against the model prices. On average, VIX futures are about .20% cheap relative to the model and within the model bounds.[16] While the model performs well on average, the averages mask substantial variation in the pricing errors over

[15]Model prices are interpolated to match the maturity of the VIX futures contract. The zero-month maturity is the estimate of the VIX in the model. To gain some insight into the magnitude of the interpolation errors, when I compute interpolated prices for the 2, 4, and 6-month maturities using observed prices at the other months, the RMSE for the interpolated prices is .04%, .004%, and .005%. This is small compared to the average RMSE of .99% in Table 9, indicating that interpolation is not driving the pricing errors.

[16]VIX futures are bounded by volatility swap rates and variance swap forward rates, $Fvol_{t,n+1} \leq Fut_{t,n} \leq \sqrt{F_{t,n+1}}$. See the Appendix for a derivation of this result.

time. Table 9 reveals that the lower (upper) bounds are violated on about 30% (16%) of days. The plots on the right in Figure 6 further illustrate this result by plotting the pricing errors and absolute pricing errors over time. According to the model, VIX futures have been as cheap as .50% to 1% in recent years relative to variance swaps. Historically, there have also been prolonged periods with large pricing errors as large as 5% for the front-month contract and as large as 2% to 3% for the other contracts. The magnitude of the futures pricing errors is also large relative to the corresponding variance swap pricing errors (MAE of .73% versus .29% for absolute errors in the bottom right plot).

Finally, the results indicate that the absolute pricing errors for VIX futures are positively correlated with the VIX Index. This observation suggests that a component of the time variation is related to financial distress and a lack of arbitrage capital to trade against the pricing errors. At the same time, financial distress as proxied by the VIX only explains part of the time variation in the pricing errors. The next section examines these results further by exploring how much of the variation in VIX futures prices is explained by the model relative to other reduced form variables.

4.3 Explaining Changes in VIX Futures Prices

Table 10 reports regressions of daily changes in VIX futures prices onto changes in model futures prices and other reduced form explanatory variables. If the model performs well, the explanatory power will be high and the coefficient on the change in the model price will be close to one. If the model fails to capture important variation in futures prices, other variables may enter significantly and increase the explanatory power as measured by the R^2_{adj}. Panel A includes the full sample period from March 2003 to 2016. Panel B is a post-financial crisis subsample from 2010 to 2016. The regressions include all contracts with between three days and one-year to maturity to avoid the roll and longer dated contracts that may be less liquid.

The first specification (1) begins by regressing changes in futures prices onto changes in model futures prices. The model is highly significant and explains around 75% of the variation in futures prices across the two sample periods. At the same time, the coefficient is significantly different from one and about 25% of the variation in prices remains unexplained. This leaves open the possibility that other variables may drive the model out of the regression and increase the explanatory power.

The second specification (2) adds the change in the VIX index between 4pm and 4:15pm on the current and previous day to account for the non-synchronous observation of futures

and option prices.[17] While the high frequency changes in the VIX index enter the regressions significantly with the expected sign, the explanatory power and coefficient on the model are largely unchanged. This result suggests that asynchronicity is not driving the unexplained variation in futures prices.

The third specification (3) adds the change in the VIX index and the CRSP value-weighted excess return. While both of these variables enter significantly with the expected sign, the explanatory power only increases by about 3%. The coefficient magnitudes on the VIX and market return are also relatively small. For example, in Panel A, a 1% increase in the model price is associated with a .57% increase in the observed price in specification (3), down from a .71% increase in specification (1). In contrast, a 1% increase in the VIX and market return are associated with only .03% and -.07% increases in the model price.

The fourth specification (4) adds measures of signed volume to contrast the other variables which are constructed from prices. The signed volume is defined as the daily volume traded multiplied by an indicator for whether the VIX increased (buy) or decreased (sell). This variable is then normalized by a one-month moving average of VIX futures open interest across contracts, standardized, and winsorized at five z-scores to mitigate the impact of outliers. The signed buy and sell variables enter the regressions significantly with the expected sign, but the increase in explanatory power is relatively limited. In contrast to the findings in Dong (2016), this result suggests that demand-based explanations for VIX futures mispricings may have limited scope.[18]

The final specification (5) replaces the VIX and market return variables with time and contract fixed effects. This change increases the explanatory power by around 12%, indicating that a significant part of the variation in futures prices remains unexplained by both the model and reduced form variables. That said, the coefficient on the change in the model price still remains highly significant in specification (5), with only a limited decline in magnitude relative to specification (4).

4.4 VIX Futures Return Predictability

Table 11 reports return predictability regressions for VIX futures. The dependent variable is the excess return from selling futures contracts $Fut^o_{t,n} - Fut^o_{t+h,n}$ over a one-week $h = 5/21$

[17] OptionMetrics reports the best bid and ask at 15:59 EST to be synchronous with the close of equity markets as of March 5th, 2008. VIX futures settlement prices are the average of the best bid and ask at the close of regular trading hours which occurs at 16:15 EST.

[18] In unreported results I compute alternative measures of demand pressure including signed volume from high frequency data for the VXX ETP and short-term VIX ETP demand from equation 2 in Dong (2016). By themselves, daily changes in the high frequency and short-term demand variables explain about 20% and 10% of the variation in futures prices respectively. As in Table 10, these measures remain significant but have limited ability to increase explanatory power when included alongside the change in the model price.

horizon where $Fut^o_{t,n}$ is the daily settlement price for the n-month futures contract on day t. If less than five trading days are left before expiration, the holding period return is computed from the final settlement value for the contract.[19] All of the variables are standardized except for the contemporaneous market return for which the coefficient can be interpreted as a beta or factor loading. As before, Panel A is the full sample and Panel B is a subsample including observations for contracts with between three days and one-year to maturity.

The first specification (1) regresses one-week realized returns from t to $t+h$ onto the model's one-month expected return at time t. Similar to the model price, the model expected return is computed by interpolating over maturities to match the futures contract expiration date. The results indicate that the model expected return delivers a significant forecast with an R^2_{adj} of 3% to 4% over a one-week horizon during both the full sample and post-financial crisis subsample. The magnitude of the coefficient is also large. A one-standard deviation increase in the model expected return predicts a .31% increase in weekly returns. This contrasts the standard deviation of weekly returns which is only 1.50%.

The second specification (2) adds the model pricing error $e_{t,n} = Fut^o_{t,n} - Fut_{t,n}$ as an additional predictor. The model expected return and pricing error are roughly uncorrelated with a 5% (-4%) Pearson (Spearman) correlation over the full sample. As such, the model expected return remains significant and has a similar coefficient across specifications (1) and (2). In addition, the model pricing error is found to deliver significant return forecasts with a coefficient of .23 over the full sample that is comparable in magnitude to the .30 coefficient on the model expected return. This result lends weight to the interpretation that VIX futures are mispriced relative to the model.

The subsequent specifications test the robustness of these findings. The third specification (3) adds the VIX, realized variance over the past month, and the slope of the VIX futures curve as additional predictors. The model expected return and pricing error remain significant in the presence of these variables with the explanatory power increasing to 9%. The fourth specification (4) adds the contemporaneous market return which reflects a beta of about .40. The continued significance indicates that the model expected return and pricing error predict VIX futures returns after adjusting for market risk.

Finally, the fifth (5) and sixth (6) specifications add time and contract fixed effects. While the model expected return remains significant throughout, the coefficient on the model

[19] In contrast to the synthetic variance swap returns, the settlement dates for VIX futures are fixed in calendar time. As a result, the time until settlement for VIX futures varies over time. The final settlement date for the n-month VX serial contract is on the Wednesday that is 30 days prior to the third Friday of the following calendar month. For example, the December 2016 contract (VX Z16) settled on December 21, 2016 using the special opening quotation (SOQ) of the January 2017 SPX options that expired thirty days later. In computing excess returns for the front month contract, if the settlement date comes before the one-week or one-month horizon, I use the holding period return from trade date t until the settlement of the contract.

pricing error remains positive but loses some of its significance. One explanation for this result is the correlation of the model pricing errors across contracts. From 2007 to 2016 the average pairwise correlation of the pricing errors for the front six contracts is 68%. This high degree of comovement can be seen visually in the top right plot in Figure 6. The daily fixed effect absorbs some of this variation, decreasing the significance of the model pricing errors in the final specification.

4.5 VIX Futures Trading Strategy

Figure 7 examines the economic significance the predictability from the perspective of a VIX futures trading strategy based on the return predictability regressions. The trading strategy is constructed as follows. Each day in the sample a hedge ratio is obtained by regressing weekly VIX futures excess returns onto weekly CRSP value-weighted excess returns in rolling one-year pooled OLS regressions. Hedged returns are then defined as,

$$R^{hedge}_{t+h,n} = Fut^{o}_{t,n} - Fut^{o}_{t+h,n} - \beta_t \cdot R^{mkt}_{t+h}, \tag{32}$$

where β_t is the hedge ratio and R^{mkt}_{t+h} is the one-week CRSP value-weighted excess return. Return forecasts are then obtained from rolling one-year regressions of hedged returns onto expanding window estimates of the model expected returns and pricing errors. This delivers a pseudo-out-of-sample forecast $\hat{y}_{t,n}$ for each futures contract that is available in real-time. The portfolio weight for each contract is then set to,

$$\omega_{t,n} = 2 \cdot N\left(\frac{\hat{y}_{n,t}}{\sigma(\hat{y}_t)}\right) - 1, \tag{33}$$

where $\sigma(\hat{y}_t)$ is the standard deviation of the return forecasts across contracts over the previous year. The portfolio weight $\omega_{t,n} \in (-1, 1)$ approaches one for significantly positive return forecasts and negative one for significantly negative return forecasts. The weekly return $h = 5$ for the strategy is then computed each day as,

$$R^{fut}_{t+h} = \frac{1}{N}\sum_{n=1}^{N} \omega_{t,n} \cdot R^{hedge}_{t+h,n}. \tag{34}$$

Figure 7 plots the cumulative sum of these excess returns R^{fut}_{t+h} against the corresponding market returns R^{e}_{t+h} both normalized to have 10% annualized volatility. In addition, the figure plots the strategy's average position $\frac{1}{N}\sum_{n=1}^{N} \omega_{t,n}$ across futures contracts as a one-month moving average to provide a sense for when VIX futures are expensive and cheap, and

to indicate what the turnover of the strategy is like. Finally, the figure plots the hedge ratio β_t over time whose average of about .40 is similar to the point estimate on contemporaneous market returns in Table 11.

The returns from the trading strategy are significantly higher than the corresponding market returns. The VIX futures trading strategy earns an annualized Sharpe ratio of 1.80 versus .50 for the market. The weekly CAPM alpha is .35% with a Newey-West t-statistic of 5.21 using 15 lags. The CAPM beta is close to zero and insignificant, which reflects the hedged nature of the strategy. Overall, the performance suggests that the return predictability results and size of the pricing errors are economically significant. Of course, it is important to caveat this interpretation. Implementing this trading strategy in practice would entail additional transaction, price impact, and funding costs that are not included in this analysis. Nonetheless, the returns from the paper-strategy indicate that the pricing errors and positioning across the different contracts are providing a useful measure of mispricing in the VIX futures market.

5 Conclusion

By modeling the logarithm of realized variance, this paper develops a dynamic term-structure model that provides a closed form relationship between the prices of variance swaps and VIX futures. The paper estimates the model using a detailed dataset that includes realized variance estimates from high frequency data and synthetic variance swap rates constructed from index option prices. Return predictability results support the hypothesis that equity volatility markets are integrated by a common stochastic discount factor, as the model's estimated expected returns significantly forecast the returns from selling volatility through variance swaps, index option straddles, and VIX futures across forecast horizons and maturities. Exploring the integration hypothesis further, the paper also provides a detailed empirical investigation of the relative pricing of variance swaps and VIX futures. While tightly linked to variance swap rates, VIX futures prices exhibit deviations of varying significance from the no-arbitrage prices and bounds implied by the no-arbitrage model. An initial attempt to understand the source of the pricing errors does not identify additional variables that can explain why VIX futures appear mispriced. This finding coupled with the observation that the pricing errors predict VIX futures returns lends weight to the interpretation that VIX futures are mispriced at times relative to variance swaps, pushing back on the integration hypothesis.

References

Aït-Sahalia, Y., M. Karaman, and L. Mancini (2015). The term structure of variance swaps and risk premia. *Working Paper*.

Andersen, T. G., T. Bollerslev, and F. X. Diebold (2007). Roughing it up: Including jump components in the measurement, modeling, and forecasting of return volatility. *The Review of Economics and Statistics 89*(4), 701–720.

Andersen, T. G., T. Bollerslev, F. X. Diebold, and P. Labys (2003). Modeling and forecasting realized volatility. *Econometrica 71*(2), 579–625.

Andries, M., T. Eisenbach, M. Schmalz, and Y. Wang (2015). The term structure of the price of variance risk. *Working Paper*.

Bertsimas, D., L. Kogan, and A. W. Lo (2000). When is time continuous? *Journal of Financial Economics 55*, 173–204.

Black, F. (1976). Studies of stock price volatility changes. *Proceedings of the 1976 Meetings of the American Statistical Association*, 171–181.

Black, F. and M. Scholes (1973). The pricing of options and corporate liabilities. *The Journal of Political Economy*, 637–654.

Bollerslev, T., G. Tauchen, and H. Zhou (2009). Expected stock returns and variance risk premia. *Review of Financial Studies 22*(11), 4463–4492.

Carr, P. and L. Wu (2006). A tale of two indices. *The Journal of Derivatives 13*(3), 13–29.

Carr, P. and L. Wu (2009). Variance risk premiums. *Review of Financial Studies 22*(3), 1311–1341.

Dew-Becker, I., S. Giglio, A. Le, and M. Rodriguez (2017). The price of variance risk. *Journal of Financial Economics 123*(2), 225–250.

Dong, X. S. (2016). Price impact of ETP demand on underliers. *Working Paper*.

Drechsler, I. (2013). Uncertainty, time-varying fear, and asset prices. *The Journal of Finance 68*(5), 1843–1889.

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*(5), 1279–1310.

Eraker, B. and Y. Wu (2017). Explaining the negative returns to volatility claims: An equilibrium approach. *Journal of Financial Economics 125*(1), 72–98.

Filipović, D., E. Gourier, and L. Mancini (2016). Quadratic variance swap models. *Journal of Financial Economics 119*, 44–68.

Giglio, S. and B. Kelly (2017). Excess volatility: Beyond discount rates. *The Quarterly Journal of Economics 133*(1), 71–127.

Gromb, D. and D. Vayanos (2010). Limits of arbitrage. *Annu. Rev. Financ. Econ. 2*(1), 251–275.

Gürkaynak, R. S., B. Sack, and J. H. Wright (2007). The U.S. Treasury yield curve: 1961 to the present. *Journal of Monetary Economics 54*(8), 2291–2304.

Jacquier, E., N. G. Polson, and P. E. Rossi (1994). Bayesian analysis of stochastic volatility models. *Journal of Business & Economic Statistics 12*(4), 69–87.

Johnson, T. L. (2017). Risk premia and the vix term structure. *Journal of Financial and Quantitative Analysis 52*(6), 2461–2490.

Joslin, S., K. J. Singleton, and H. Zhu (2011). A new perspective on gaussian dynamic term

structure models. *Review of Financial Studies 24*(3), 926–970.

Liu, L. Y., A. J. Patton, and K. Sheppard (2015). Does anything beat 5-minute RV? A comparison of realized measures across multiple asset classes. *Journal of Econometrics 187*(1), 293–311.

Longstaff, F. A., P. Santa-Clara, and E. S. Schwartz (2001). The relative valuation of caps and swaptions: Theory and empirical evidence. *The Journal of Finance 56*(6), 2067–2109.

Martin, I. (2017). What is the expected return on the market? *The Quarterly Journal of Economics 132*(1), 367–433.

Merton, R. C. (1973). An intertemporal capital asset pricing model. *Econometrica*, 867–887.

Merton, R. C. (1980). On estimating the expected return on the market: An exploratory investigation. *Journal of Financial Economics 8*(4), 323–361.

Merton, R. C. (1987). A simple model of capital market equilibrium with incomplete information. *The journal of finance 42*(3), 483–510.

Newey, W. and K. West (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica 55*(3), 703–708.

Ross, S. A. (1976). The arbitrage theory of capital asset pricing. *Journal of Economic Theory 13*(3), 341–360.

Shleifer, A. and R. W. Vishny (1997). The limits of arbitrage. *The Journal of Finance 52*(1), 35–55.

Van Tassel, P. and E. Vogt (2016). Global variance term premia and intermediary risk appetite. *Working Paper*.

Zhang, L., P. A. Mykland, and Y. Aït-Sahalia (2005). A tale of two time scales: Determining integrated volatility with noisy high-frequency data. *Journal of the American Statistical Association 100*(472), 1394–1411.

Table 1: Summary Statistics for Variance Swap Rates and Returns

Panel A reports summary statistics for the variance swap and realized variance data from 1996 to 2016. Variance swaps are higher in level and more autocorrelated than realized variance. The term-structure is upward sloping on average, exhibiting higher volatility and more positive skewness at the short-end of the curve. Panel B reports summary statistics for monthly variance swap excess returns in percentage units. From 1996 to 2016, one-month variance swaps have earned .17% with a Sharpe ratio of .52 per month. The mean and standard deviation of variance swap returns are increasing in maturity, while the Sharpe ratio and t-statistics are decreasing. Selling volatility by receiving fixed at the short-end of the curve has earned significant abnormal returns relative to the CAPM. The table reports Newey and West (1987) t-statistics with 50 lags to adjust for the overlapping monthly returns that are observed at a daily frequency. Similar results hold for non-overlapping data at month-end dates.

Panel A: Variance Swap Rates (annualized volatility units)

Maturity in months	RV	1	3	6	9	12	18	24
Mean	15.19	20.82	21.47	22.06	22.37	22.62	23.04	23.38
Standard Deviation	8.16	8.21	7.18	6.42	6.02	5.76	5.48	5.35
Skewness	2.87	2.14	1.79	1.52	1.36	1.23	1.05	0.96
Kurtosis	16.23	10.70	8.43	7.01	6.23	5.67	4.88	4.54
Minimum	5.73	10.05	11.07	12.22	12.92	13.20	13.28	13.32
Median	13.18	19.25	20.31	21.17	21.63	21.88	22.32	22.61
Maximum	77.52	81.85	70.94	62.01	57.77	54.18	50.66	48.96
Autocorrelation 1-month	0.77	0.81	0.86	0.89	0.91	0.92	0.93	0.93
Autocorrelation 6-month	0.30	0.37	0.44	0.50	0.53	0.56	0.60	0.62

Panel B: Variance Swap Returns (one-month returns, percent)

Maturity in months	1	3	6	9	12	18	24
Mean	0.17	0.19	0.21	0.21	0.23	0.24	0.26
Standard Deviation	0.33	0.76	1.11	1.38	1.61	2.09	2.57
Sharpe ratio	0.52	0.25	0.19	0.16	0.14	0.12	0.10
t-statistic	8.12	3.85	2.90	2.40	2.27	1.85	1.67
Skewness	-3.15	-4.41	-3.64	-3.20	-2.75	-2.20	-1.87
Kurtosis	56.60	58.67	43.14	35.59	30.38	23.52	20.58
Minimum	-3.95	-10.25	-13.56	-15.31	-17.09	-20.09	-23.39
Median	0.15	0.20	0.25	0.27	0.29	0.33	0.36
Maximum	3.54	7.13	8.56	10.07	11.44	14.70	19.25
Autocorrelation 1-month	0.23	0.23	0.22	0.20	0.18	0.14	0.12
Autocorrelation 6-month	-0.03	-0.10	-0.12	-0.14	-0.16	-0.16	-0.16
Negative Percent	0.08	0.19	0.27	0.30	0.33	0.36	0.38
CAPM α	0.15	0.12	0.11	0.09	0.09	0.06	0.05
t_α-statistic	7.84	2.98	1.91	1.32	1.16	0.61	0.41
CAPM β	0.03	0.10	0.15	0.19	0.22	0.28	0.33
t_β-statistic	3.98	5.27	6.25	6.77	7.43	7.77	7.91
R^2_{adj}	0.27	0.44	0.46	0.45	0.44	0.43	0.39

Table 2: Model Performance Varying the Number of Principal Components

This table compares alternative specifications that vary the number of principal components in the state vector. A three factor model with two principal components $K_{PC} = 2$ performs well relative to competing models and to reduced form linear forecasts. While the unrestricted linear models provide good in-sample fit, their performance deteriorates out-of-sample leading to overfitting and instability concerns. The three-factor no arbitrage model is relatively stable by way of comparison, outperforming the linear models by as much as 10% to 20% in out-of-sample return prediction over one to six month horizons. The reported results are averaged across maturities from 1998 to 2016. The out-of-sample analysis is performed by estimating the models with an expanding window using 1996 to 1998 as the initial estimation period. Panel I reports variance swap pricing errors measured by the root mean squared error $RMSE$, mean absolute error MAE, and median absolute deviation MAD. Panels II-V report variance swap return forecast errors measured by explanatory power R^2_{adj}, mean squared error MSE, and mean absolute error MAE.

Panel A: In-sample

Panel A.I: Model pricing errors

K_{PC}	1	2	3	4	5
RMSE	1.13	0.37	0.33	0.21	0.16
MAE	0.84	0.27	0.24	0.14	0.10
MAD	0.67	0.21	0.19	0.10	0.08

Panel A.II: Model return forecast $R^2_{Mod,is}$

K_{PC}	1	2	3	4	5
1mn	0.13	0.16	0.18	0.18	0.18
3mn	0.20	0.21	0.21	0.22	0.21
6mn	0.30	0.30	0.31	0.30	0.30

Panel A.III: Linear return forecast $R^2_{Lin,is}$

K_{PC}	1	2	3	4	5
1mn	0.20	0.20	0.23	0.24	0.24
3mn	0.20	0.20	0.21	0.21	0.22
6mn	0.28	0.28	0.29	0.29	0.29

Panel A.IV: $\ln(MSE^{K_{PC}}_{Lin,is}/MSE^{K_{PC}=2}_{Mod,is})$

K_{PC}	1	2	3	4	5
1mn	-0.04	-0.05	-0.08	-0.09	-0.09
3mn	0.02	0.01	-0.00	-0.01	-0.01
6mn	0.04	0.03	0.02	0.02	0.02

Panel A.V: $\ln(MAE^{K_{PC}}_{Lin,is}/MAE^{K_{PC}=2}_{Mod,is})$

K_{PC}	1	2	3	4	5
1mn	0.01	-0.00	0.02	0.01	0.01
3mn	0.04	0.03	0.04	0.04	0.03
6mn	0.06	0.04	0.04	0.04	0.03

Panel B: Out-of-sample

Panel B.I: Model pricing errors

K_{PC}	1	2	3	4	5
RMSE	1.23	0.39	0.24	0.20	0.19
MAE	0.92	0.29	0.16	0.14	0.13
MAD	0.73	0.22	0.12	0.10	0.09

Panel B.II: Model return forecast $R^2_{Mod,oos}$

K_{PC}	1	2	3	4	5
1mn	0.02	0.05	0.05	0.04	0.03
3mn	0.13	0.14	0.12	0.11	0.12
6mn	0.22	0.24	0.23	0.22	0.21

Panel B.III: Linear return forecast $R^2_{Lin,oos}$

K_{PC}	1	2	3	4	5
1mn	0.01	-0.00	0.01	0.01	0.01
3mn	0.15	0.12	0.09	0.09	0.08
6mn	0.22	0.20	0.16	0.14	0.13

Panel B.IV: $\ln(MSE^{K_{PC}}_{Lin,oos}/MSE^{K_{PC}=2}_{Mod,oos})$

K_{PC}	1	2	3	4	5
1mn	0.04	0.06	0.05	0.04	0.04
3mn	-0.01	0.03	0.06	0.06	0.07
6mn	0.02	0.05	0.10	0.13	0.14

Panel B.V: $\ln(MAE^{K_{PC}}_{Lin,oos}/MAE^{K_{PC}=2}_{Mod,oos})$

K_{PC}	1	2	3	4	5
1mn	0.00	0.01	0.07	0.07	0.08
3mn	0.05	0.10	0.18	0.17	0.18
6mn	0.10	0.12	0.17	0.19	0.20

Table 3: Model Estimate

This table reports the estimated model parameters using the realized variance and variance swap rate data from 1996 to 2016. The physical parameters μ and Φ are estimated by a monthly vector autoregression equation by equation from overlapping daily data. The risk-neutral parameters $\mu^{\mathbb{Q}}$ and $\Phi^{\mathbb{Q}}$ are estimated by nonlinear least squares to minimize daily variance swap pricing errors. The table reports the associated prices of risk $\Lambda_0 = \mu - \mu^{\mathbb{Q}}$ and $\Lambda_1 = \Phi - \Phi^{\mathbb{Q}}$. The significant estimates indicate that each state variable contributes to the time-varying price of volatility risk and volatility forecasts. The table reports Newey-West t-statistics for the VAR using 50 lags and block bootstrapped t-statistics for the prices of risk where *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Panel A: Physical Parameters

	μ	$\Phi_{1,1}$	$\Phi_{1,2}$	$\Phi_{1,3}$
$\ln(RV)$	0.00	0.23***	0.21***	0.40***
	[0.01]	[2.95]	[8.43]	[5.76]
PC_{level}	0.01	0.07	0.88***	-0.07
	[0.15]	[0.55]	[23.73]	[-0.69]
PC_{slope}	-0.01	0.02	-0.04***	0.78***
	[-0.33]	[0.58]	[-3.41]	[23.89]

Panel B: Prices of Risk

	Λ_0	$\Lambda_{1,1}$	$\Lambda_{1,2}$	$\Lambda_{1,3}$
$\ln(RV)$	-0.68***	0.21***	-0.07***	-0.12**
	[-13.62]	[3.58]	[-3.23]	[-1.97]
PC_{level}	-0.06	0.11	-0.08*	0.37***
	[-0.79]	[1.01]	[-1.89]	[3.38]
PC_{slope}	-0.04	0.05	-0.03	0.02
	[-0.80]	[0.48]	[-0.76]	[0.23]

Table 4: Variance Swap Pricing Errors

This table summarizes the model fitting errors for variance swap rates and monthly returns from 1996 to 2016 using daily data. The errors are small and unbiased overall. For example, the standard deviation of the pricing errors in Panel A are below the average bid-ask spreads reported by Markit.

Panel A: Variance swap rate pricing errors (annualized volatility units)

Maturity in months	1	3	6	9	12	18	24
Mean	0.01	-0.05	0.02	0.00	0.00	0.00	-0.01
Standard Deviation	0.57	0.36	0.39	0.34	0.29	0.23	0.34
Skewness	1.15	-0.64	-0.98	-1.34	-2.22	-1.29	-0.03
Kurtosis	9.65	10.76	9.51	10.33	20.34	12.73	3.76
Autocorrelation 1-month	0.37	0.33	0.52	0.59	0.50	0.52	0.63
Autocorrelation 6-month	0.18	0.12	0.24	0.22	0.20	0.06	0.36
Bid-ask spread*	1.39	0.82	0.74	0.67	0.60	0.62	0.61

*Markit month-end data 2006:09 to 2015

Panel B: Variance swap return pricing errors (monthly returns in basis points*)

Maturity in months	1	3	6	9	12	18	24
Mean	0.13	0.03	0.28	-0.35	0.46	-0.36	0.46
Standard Deviation	3.17	5.83	11.96	14.68	19.32	20.63	33.37
Skewness	5.32	-4.60	-1.57	-1.36	2.95	0.63	0.18
Kurtosis	75.04	80.28	41.99	50.44	89.19	45.18	16.87
Autocorrelation 1-month	0.23	0.05	-0.18	-0.13	-0.16	-0.10	-0.17
Autocorrelation 6-month	0.04	0.01	-0.05	-0.07	-0.05	-0.02	0.03

*Monthly return from overlapping daily data

Table 5: Variance Swap Return Predictability

This table reports return predictability regressions for variance swap returns. The estimated model expected returns $\hat{E}_t[R_{t+h,n}]$ are significant for all horizons h and maturities n. Newey-West t-statistics indicate significance using $3 \cdot h \cdot 21$ lags to account for the overlapping returns that are computed from daily data from 1996 to 2016.

Variance Swap Return Predictability: $R_{t+h,n} = \beta_0 + \beta_1 \hat{E}_t[R_{t+h,n}] + \epsilon_{t+h,n}$

Maturity	1	3	6	9	12	18	24
One-month returns ($h = 1$, $T = 5,190$)							
β_0	0.02	-0.09	-0.15	-0.18	-0.23	-0.25	-0.12
	[0.78]	[-1.52]	[-1.33]	[-1.20]	[-1.30]	[-1.40]	[-0.68]
β_1	0.85***	1.36***	1.52***	1.60***	1.83***	1.85***	1.36***
	[4.55]	[9.07]	[5.71]	[4.62]	[4.39]	[5.19]	[5.22]
R^2_{adj}	0.23	0.16	0.17	0.16	0.17	0.13	0.12
Three-month returns ($h = 3$, $T = 5,190$)							
β_0		-0.13*	-0.21	-0.26	-0.31	-0.37	-0.30
		[-1.71]	[-1.20]	[-1.12]	[-1.07]	[-0.98]	[-0.64]
β_1		1.14***	1.21***	1.26***	1.32***	1.37***	1.28***
		[12.72]	[10.69]	[11.20]	[11.55]	[10.22]	[8.22]
R^2_{adj}		0.31	0.21	0.20	0.19	0.18	0.18
Six-month returns ($h = 6$, $T = 5,190$)							
β_0			-0.29	-0.40	-0.47	-0.53	-0.51
			[-1.21]	[-0.97]	[-0.90]	[-0.75]	[-0.58]
β_1			1.15***	1.20***	1.23***	1.26***	1.24***
			[12.59]	[11.87]	[11.59]	[10.58]	[9.26]
R^2_{adj}			0.35	0.31	0.29	0.28	0.28

Table 6: Variance Swap Rate Decompositions

Panel A provides a variance decomposition for variance swap rates into percentage contributions from variance term premia and realized variance forecasts. Panel B regresses monthly changes in variance term premia in annualized percentage points onto z-scored changes in realized variance and the first two principal components of the level of variance swap rates.

Panel A: Variance Decomposition for Variance Swaps: $VS_{t,n} = RVF_{t,n} + VTP_{t,n}$

Maturity	1	3	6	9	12	18	24
$\sigma_{VS,RVF}/\sigma^2_{VS}$	0.58	0.54	0.46	0.39	0.33	0.25	0.20
$\sigma_{VS,VTP}/\sigma^2_{VS}$	0.42	0.46	0.54	0.61	0.67	0.75	0.80

Panel B: Variance Term Premia Decomposition: $\Delta VTP_{t,n} = \beta' \Delta f_t + \epsilon_{t,n}$

Maturity	1	3	6	9	12	18	24
ΔRV	-0.77***	-0.40***	-0.17***	-0.15***	-0.13***	-0.07***	-0.04***
	[-13.38]	[-14.93]	[-18.99]	[-16.58]	[-21.25]	[-8.59]	[-3.41]
$\Delta PC1$	2.45***	2.43***	2.23***	2.19***	2.18***	2.09***	2.10***
	[19.55]	[26.78]	[23.68]	[24.33]	[25.67]	[75.61]	[90.76]
$\Delta PC2$	1.53***	0.05	-0.43***	-0.58***	-0.65***	-0.67***	-0.72***
	[31.93]	[1.11]	[-17.71]	[-38.04]	[-44.32]	[-17.64]	[-18.10]
R^2_{adj}	0.96	0.94	0.94	0.96	0.97	0.98	0.97

Table 7: Explaining Straddle Returns with Variance Swap Returns

Panel A reports the correlation of straddle and variance swap returns. Panel B reports straddle return summary statistics. Panel C reports explanatory regressions for straddle returns. The returns are highly correlated overall, with variance swap returns explaining straddle returns with an average R^2_{adj} of 76% over a one-month horizon. Estimated alphas measuring relative performance indicate that straddles and variance swaps deliver similar average returns after adjusting for risk. The sample period is 1996 to 2016.

Panel A: Straddle × VS Correlation (one-month returns)

Maturity	1	3	6	9	12	18	24	Avg.
(1,3]	0.72	0.84	0.82	0.80	0.77	0.74	0.71	0.77
(3,6]	0.68	0.84	0.87	0.86	0.84	0.82	0.80	0.82
(6,9]	0.63	0.81	0.86	0.87	0.87	0.86	0.85	0.82
(9,15]	0.57	0.76	0.83	0.86	0.87	0.88	0.88	0.81
(15,24]	0.51	0.69	0.78	0.82	0.84	0.87	0.88	0.77

Panel B: Straddle Return Summary Statistics (one-month returns, percent)

Maturity	$(1,3]$	$(3,6]$	$(6,9]$	$(9,15]$	$(15,24]$	Avg.
Mean	0.31	0.19	0.12	0.08	0.03	0.15
Standard Deviation	1.39	1.35	1.34	1.36	1.44	1.33
Skewness	-3.57	-3.13	-2.51	-1.87	-1.37	-2.56
Kurtosis	30.25	26.06	19.61	15.07	12.18	20.35
Autocorrelation 1-month	0.26	0.27	0.23	0.21	0.15	0.25

Panel C: Straddle Return Explanatory Regressions $R^{straddle}_{t+h,b} = \alpha + \beta R_{t+h,n} + \epsilon_{t+h,b}$

Straddle Maturity Bucket b	(1,3]	(3,6]	(6,9]	(9,15]	(15,24]	Avg.
Expected Return Maturity n	3	6	9	12	18	Avg.
One-month returns ($h = 1$, $T = 5,190$)						
α	0.02	-0.03	-0.06	-0.09**	-0.12***	-0.04
	[0.31]	[-0.64]	[-1.55]	[-2.12]	[-2.72]	[-0.96]
β	1.53***	1.05***	0.84***	0.73***	0.60***	0.84***
	[18.28]	[22.94]	[25.73]	[23.00]	[23.64]	[26.30]
R^2_{adj}	0.70	0.75	0.76	0.75	0.76	0.76
Three-month returns ($h = 3$, $T = 5,190$)						
α	-0.28	-0.21*	-0.26*	-0.29**	-0.37**	-0.20
	[-1.58]	[-1.69]	[-1.92]	[-2.01]	[-2.41]	[-1.56]
β	2.25***	1.30***	0.97***	0.82***	0.65***	1.01***
	[12.48]	[24.35]	[28.70]	[26.31]	[25.90]	[26.52]
R^2_{adj}	0.69	0.81	0.81	0.81	0.79	0.82
Six-month returns ($h = 6$, $T = 5,190$)						
α		-0.59*	-0.59*	-0.62*	-0.77**	-0.57*
		[-1.78]	[-1.77]	[-1.72]	[-2.07]	[-1.77]
β		1.55***	1.06***	0.87***	0.67***	0.96***
		[13.10]	[15.95]	[14.20]	[13.61]	[15.14]
R^2_{adj}		0.69	0.74	0.74	0.74	0.76

Table 8: Straddle Return Predictability

This table reports return predictability regressions for straddle returns for different maturity buckets b using the estimated variance swap expected return $\hat{E}_t[R_{t+h,n}]$ of maturity n to forecast h-month returns. Newey-West t-statistics indicate significance using $21 \cdot h \cdot 3$ lags.

Straddle Return Predictability from 1996 to 2016: $R^{straddle}_{t+h,b} = \beta_0 + \beta_1 \hat{E}_t[R_{t+h,n}] + \epsilon_{t+h,b}$

Straddle Maturity Bucket b	(1,3]	(3,6]	(6,9]	(9,15]	(15,24]	Avg.
Expected Return Maturity n	3	6	9	12	18	Avg.
One-month returns ($h = 1$, $T = 5,190$)						
β_0	-0.00	-0.14	-0.18	-0.23*	-0.27**	-0.19
	[-0.02]	[-1.10]	[-1.35]	[-1.74]	[-2.27]	[-1.44]
β_1	1.49***	1.42***	1.22***	1.25***	1.15***	1.39***
	[5.79]	[5.15]	[4.35]	[4.47]	[4.97]	[4.81]
R^2_{adj}	0.06	0.10	0.10	0.11	0.11	0.10
Three-month returns ($h = 3$, $T = 5,190$)						
β_0	-0.02	-0.24	-0.35	-0.46	-0.54*	-0.32
	[-0.09]	[-0.82]	[-1.17]	[-1.46]	[-1.68]	[-1.07]
β_1	1.62***	1.22***	1.00***	0.96***	0.80***	1.09***
	[7.55]	[6.77]	[7.17]	[7.24]	[6.41]	[7.40]
R^2_{adj}	0.08	0.10	0.11	0.12	0.11	0.11
Six-month returns ($h = 6$, $T = 5,190$)						
β_0		-0.59	-0.73	-0.87	-1.00	-0.77
		[-0.99]	[-1.17]	[-1.34]	[-1.46]	[-1.22]
β_1		1.41***	1.06***	0.97***	0.77***	1.02***
		[5.85]	[6.52]	[6.50]	[6.19]	[6.55]
R^2_{adj}		0.15	0.16	0.18	0.17	0.17

Table 9: VIX Futures Pricing Errors

This table reports the model pricing errors for VIX futures for the front six contracts from 2007 to 2016 after estimating the model with realized variance and variance swap rate data. Figure 6 plots the pricing errors to illustrate these results.

VIX Futures Pricing Errors 2007-2016: $e_{t,n} = Fut^o_{t,n} - Fut_{t,n}$

Contract (n)	1	2	3	4	5	6	Avg.
Mean	-0.25	-0.18	-0.16	-0.20	-0.17	-0.12	-0.18
RMSE	0.90	1.21	1.09	0.95	0.90	0.87	0.99
Mean t-statistic	-4.26	-1.54	-1.36	-1.75	-1.51	-1.10	-1.92
Skewness	-2.31	-2.13	-1.19	-0.18	-0.07	-0.02	-0.98
Kurtosis	17.50	17.71	10.72	4.27	3.54	3.14	9.48
Autocorrelation 1-week	0.38	0.75	0.82	0.85	0.85	0.83	0.75
Autocorrelation 1-month	0.08	0.25	0.39	0.58	0.61	0.61	0.42
Percent below lower bound	0.35	0.28	0.30	0.33	0.28	0.26	0.30
Percent above upper bound	0.26	0.22	0.18	0.13	0.09	0.07	0.16
RMSE lower bound	0.89	1.34	1.26	1.13	1.10	1.11	1.14
RMSE upper bound	1.01	1.39	1.37	1.40	1.43	1.46	1.34
Corr($\lvert e_{t,n}\rvert$, $\lvert e^{VS}_{t,n}\rvert$)	0.59	0.44	0.40	0.38	0.22	0.14	0.36
Corr($\lvert e_{t,n}\rvert$, VIX)	0.49	0.59	0.49	0.31	0.22	0.20	0.38
Corr($\Delta\lvert e_{t,n}\rvert$,ΔVIX) 5d chg.	0.49	0.44	0.41	0.29	0.15	0.07	0.31

Table 10: Explaining Daily Changes in VIX Futures Prices

This table reports regressions of daily changes in VIX futures prices onto the change in the model futures price and other explanatory variables. Panel A includes the full sample period from March 2003 to 2016. Panel B is a subsample from 2010 to 2016 that excludes the financial crisis. The change in the model futures price remains highly significant across specifications and sample periods. For example, in Panel A, the model explains 74.5% of the variation in futures prices by itself. Including additional explanatory variables that enter significantly with the expected sign only increases the R^2_{adj} to 76.9%. The t-statistics in brackets are double clustered by date and contract. The final specification includes date and contract fixed effects.

Panel A: 2004.03-2016	(1)	(2)	(3)	(4)	(5)
Δ Model Futures Price	0.71***	0.71***	0.57***	0.53***	0.54***
	[39.85]	[38.28]	[25.04]	[22.92]	[21.81]
Δ VIX 4pm to 415pm		0.14**	0.10**	0.09*	
		[2.60]	[2.12]	[1.97]	
Δ VIX 4pm to 415pm Lag		-0.12**	-0.09*	-0.09*	
		[-2.38]	[-1.79]	[-1.85]	
Δ VIX			0.03***	0.03***	
			[2.95]	[3.06]	
Market Return			-0.07***	-0.08***	
			[-4.91]	[-5.43]	
Volume Buy				0.04***	0.03***
				[6.02]	[4.41]
Volume Sell				-0.04***	-0.01**
				[-5.53]	[-2.49]
Observations	23681	23681	23681	23681	23681
Adjusted R^2	0.745	0.748	0.764	0.769	0.896
Fixed Effects	No	No	No	No	Yes

Panel B: 2010-2016	(1)	(2)	(3)	(4)	(5)
Δ Model Futures Price	0.77***	0.78***	0.61***	0.58***	0.55***
	[36.53]	[33.54]	[22.51]	[19.26]	[20.91]
Δ VIX 4pm to 415pm		0.38***	0.31***	0.30***	
		[5.19]	[4.30]	[4.13]	
Δ VIX 4pm to 415pm Lag		-0.27***	-0.24***	-0.24***	
		[-4.35]	[-3.98]	[-4.00]	
Δ VIX			0.04***	0.04***	
			[3.22]	[3.35]	
Market Return			-0.09***	-0.09***	
			[-5.10]	[-5.42]	
Volume Buy				0.03***	0.03***
				[3.77]	[4.63]
Volume Sell				-0.02***	-0.01
				[-2.91]	[-1.34]
Observations	14650	14650	14650	14650	14650
Adjusted R^2	0.776	0.794	0.813	0.814	0.932
Fixed Effects	No	No	No	No	Yes

Table 11: VIX Futures Return Predictability

This table reports return predictability regressions for weekly VIX futures returns. Panel A includes the full sample period from March 2003 to 2016. Panel B is a subsample from 2010 to 2016 that excludes the financial crisis. The model expected return and pricing error significantly predict returns in the first three specifications which use time-t data to forecast the t to $t+5$ return. Even after adding the contemporaneous CRSP value-weighted market return or date and contract fixed effects in the next three specifications, the model expected return and pricing error remain significant. All of the explanatory variables are standardized except for the market return for which the coefficient can be interpreted as a beta or factor loading. The t-statistics are clustered by date and contract and use 5 lags to adjust for autocorrelation.

Panel A: 2004.03-2016	(1)	(2)	(3)	(4)	(5)	(6)
Model Expected Return	0.31***	0.30***	0.22***	0.21***	0.34***	0.30***
	[5.24]	[5.52]	[4.64]	[4.47]	[4.04]	[3.71]
Model Pricing Error		0.23***	0.19***	0.15***	0.14**	0.06*
		[4.21]	[4.63]	[3.96]	[2.50]	[1.78]
VIX Futures Slope			0.10*	0.12*		0.15**
			[1.88]	[1.93]		[2.50]
VIX			0.53***	0.03		
			[3.45]	[0.20]		
Realized Variance			-0.64***	-0.13		
			[-3.61]	[-1.02]		
Market Return				0.39***		
				[14.66]		
Observations	23828	23828	23828	23828	23828	23828
Adjusted R^2	0.041	0.063	0.093	0.524	0.779	0.784
Fixed Effects	No	No	No	No	Yes	Yes

Panel B: 2010-2016	(1)	(2)	(3)	(4)	(5)	(6)
Model Expected Return	0.30***	0.26***	0.20***	0.24***	0.51***	0.48***
	[4.36]	[3.64]	[2.70]	[4.05]	[4.85]	[4.34]
Model Pricing Error		0.13**	0.14***	0.08**	0.10**	0.07*
		[2.61]	[2.65]	[2.35]	[2.43]	[1.88]
VIX Futures Slope			0.01	0.04		0.05
			[0.34]	[0.84]		[1.36]
VIX			0.46**	-0.09		
			[2.60]	[-0.68]		
Realized Variance			-0.43**	-0.03		
			[-2.05]	[-0.22]		
Market Return				0.50***		
				[24.51]		
Observations	14735	14735	14735	14735	14735	14735
Adjusted R^2	0.030	0.036	0.046	0.601	0.830	0.830
Fixed Effects	No	No	No	No	Yes	Yes

Figure 1: Realized Variance and Variance Swap Rates

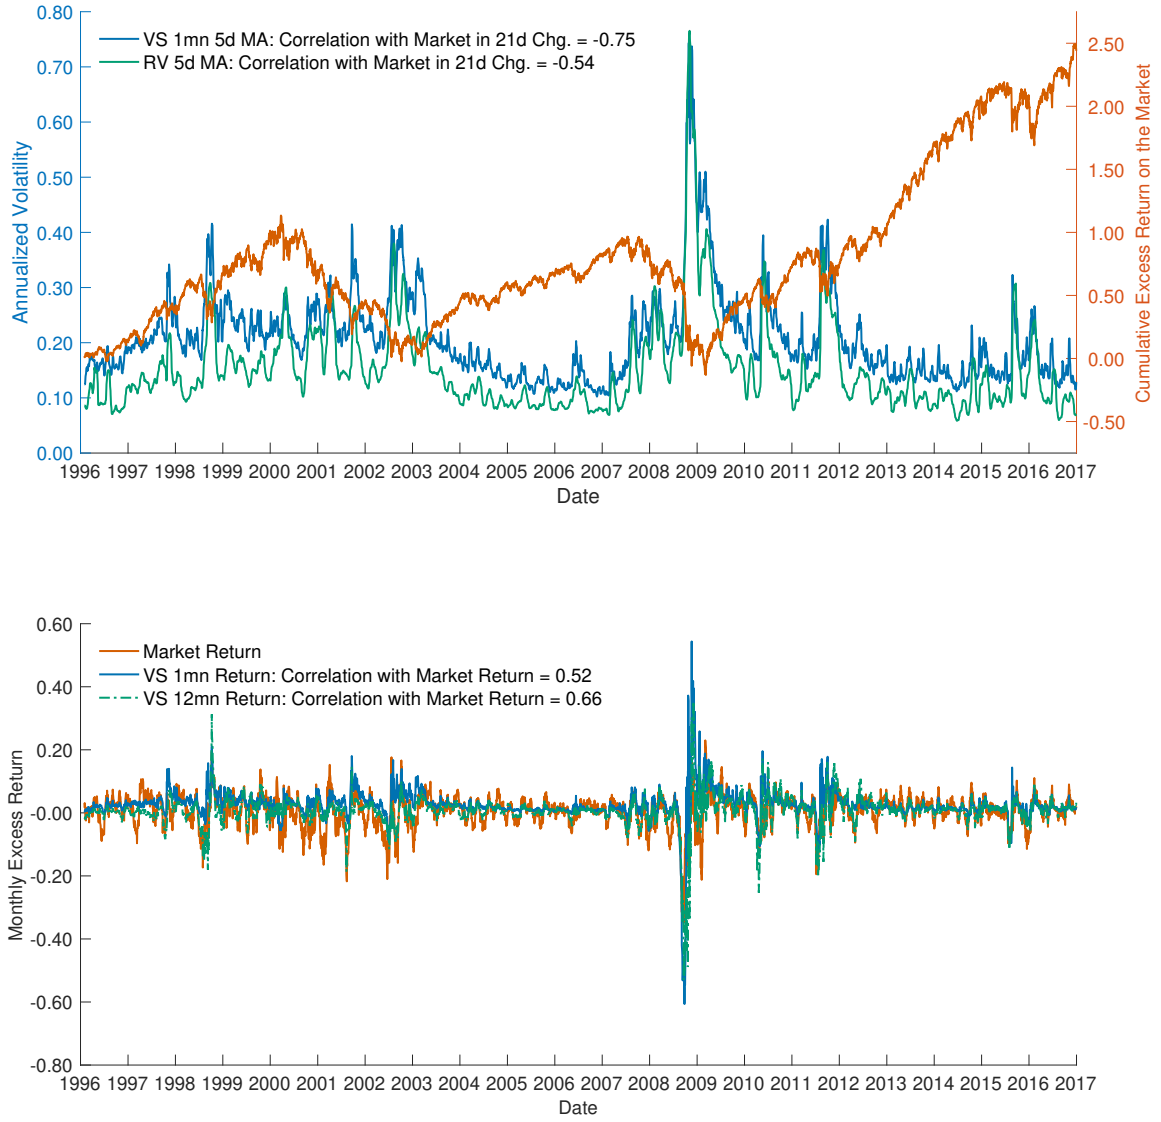

The top figure plots the time series of realized variance and one-month variance swap rates against the cumulative excess return on the market. Increases in volatility are negatively correlated with stock market returns, a result known as the leverage effect due to Black (1976). Selling volatility insurance by receiving fixed in variance swaps produces returns that are significantly positively correlated with stock market returns. Equivalently, paying fixed provides a hedge for stock market declines. The bottom figure illustrates this result by plotting overlapping monthly returns for the stock market and for receiving fixed in one-month and twelve-month variance swaps at a daily frequency. Market returns are defined as the CRSP value-weighted return in excess of the one-month Treasury bill rate obtained from Ken French's website. The returns are normalized to have 5% monthly volatility from 1996 to 2016 for comparison.

Figure 2: The Time-Varying Price of Volatility Risk

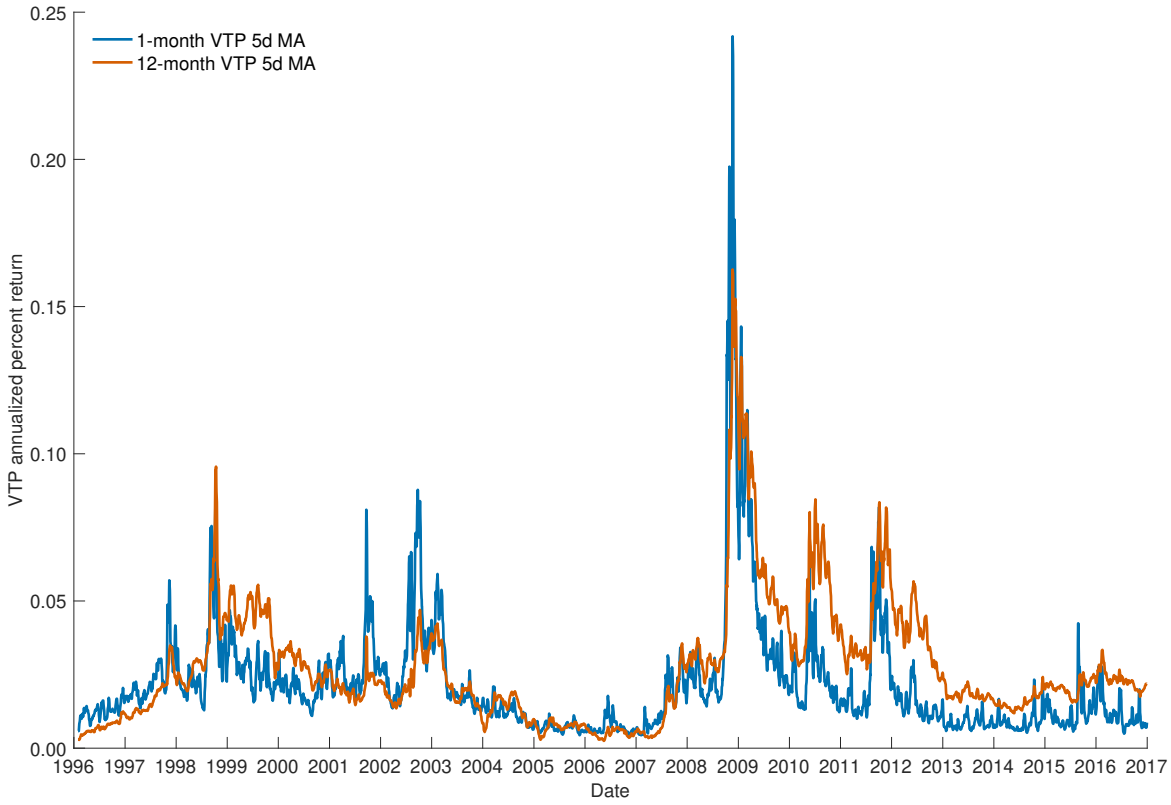

This figure plots the time-varying price of volatility risk as measured by the one-month and twelve-month variance term premia estimates $VTP_{t,n}$. Variance term premia represent the expected holding period return from receiving fixed in variance swaps. Variance term premia can be interpreted as the cost of insuring against realized variance shocks over different horizons. High levels of variance term premia predict high returns from selling volatility by receiving fixed in variance swaps, selling VIX futures, or selling straddles on SPX options.

Figure 3: Variance Swap Pricing

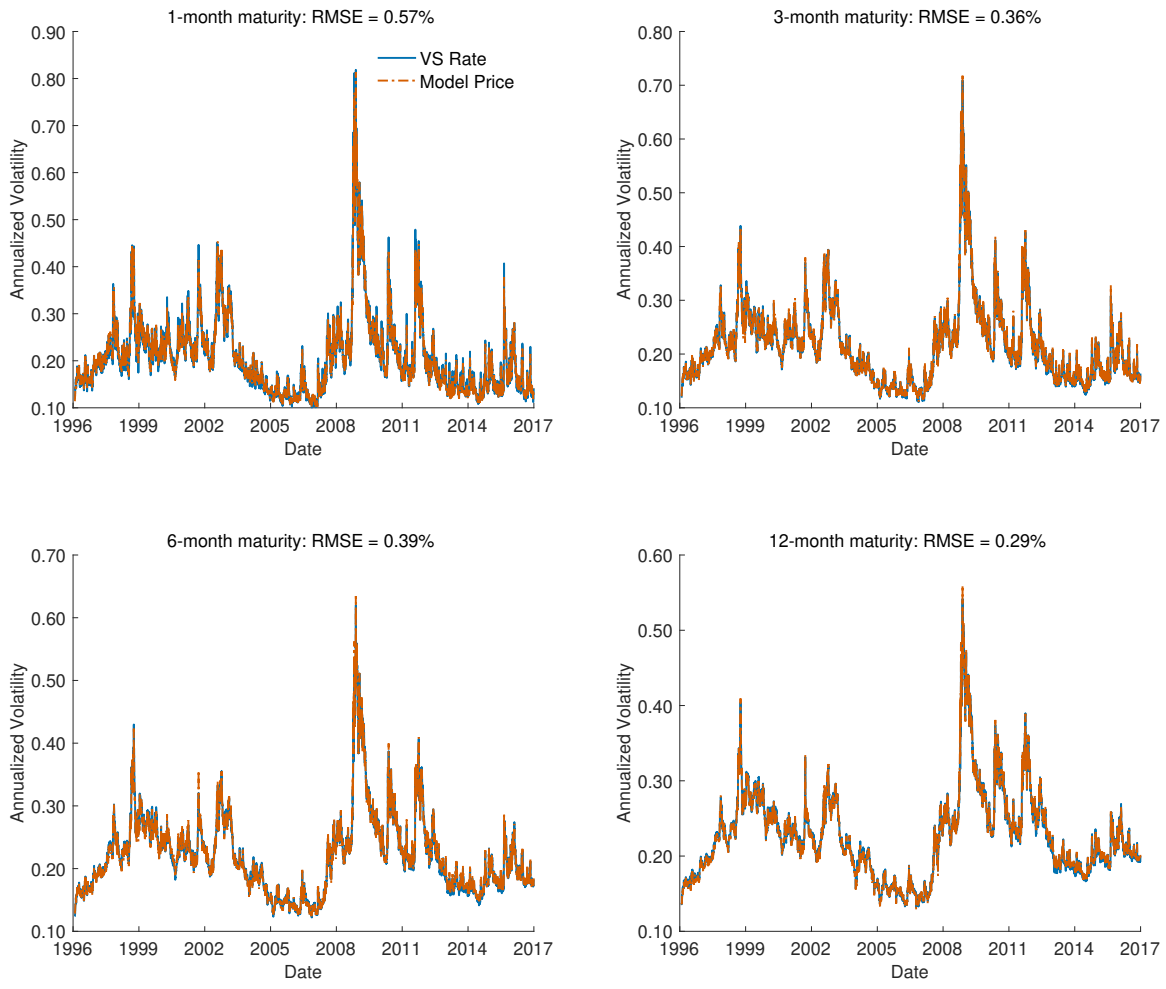

This figure plots observed variance swap rates against the estimated model prices for 1, 3, 6, and 12-month maturities at a daily frequency from 1996 to 2016. A three-factor logarithmic model with two principal components fits the cross section of variance swap rates with small pricing errors as measured by the root-mean-squared-errors (RMSEs) in the titles of the subplots. The model's risk neutral parameters $\mu^{\mathbb{Q}}$ and $\Phi^{\mathbb{Q}}$ are estimated by nonlinear least squares to minimize variance swap pricing errors for 1, 3, 6, 9, 12, 18, and 24-month maturities.

Figure 4: Variance Swap Return Predictability

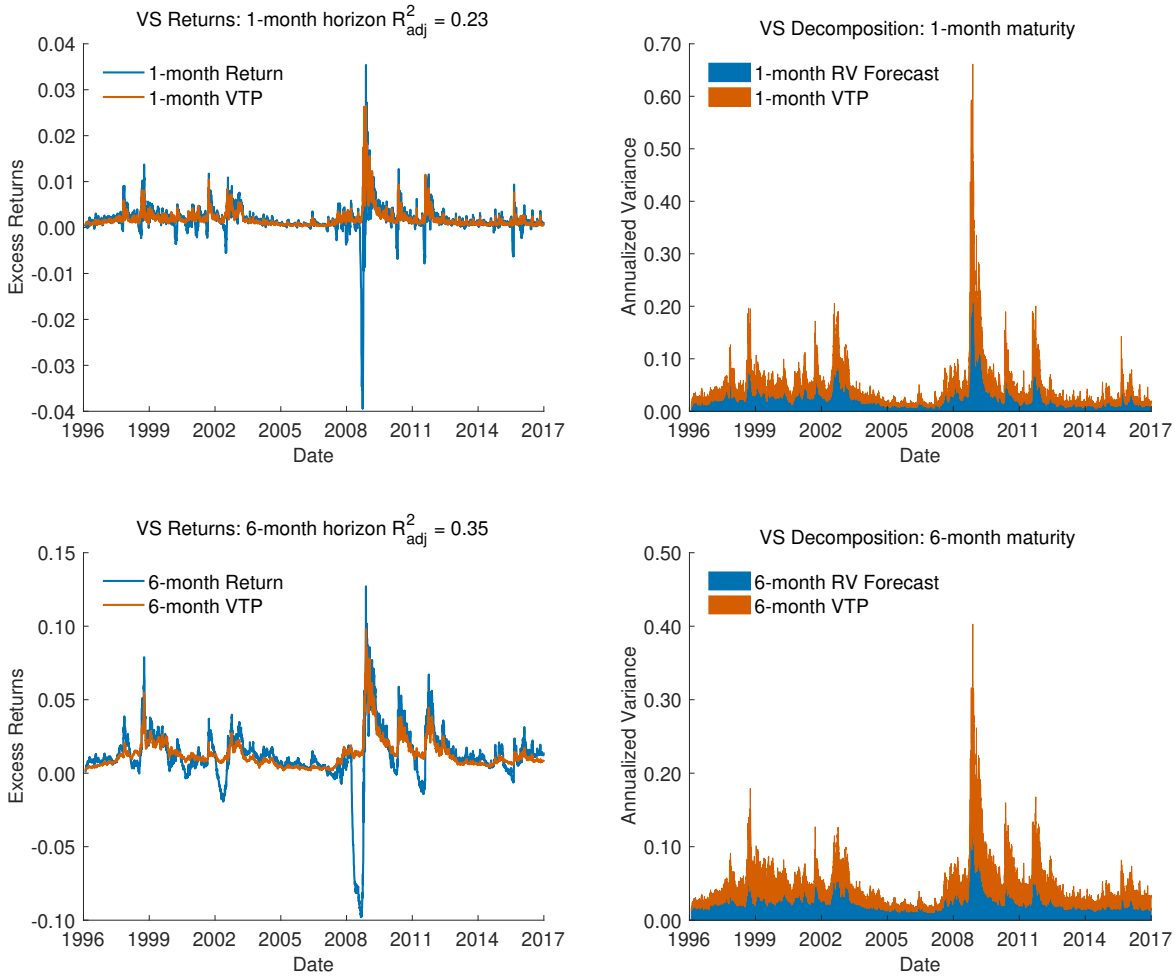

The plots on the left report the excess return from receiving fixed in variance swaps against the variance term premium over 1-month and 6-month horizons. High levels of the variance term premium forecast high returns from receiving fixed in variance swaps, consistent with the in-sample return predictability results in Table 5. The plots on the right decompose variance swap rates into realized variance forecasts and variance term premia. The realized variance forecasts are less volatile and the variance term premia are more persistent over longer horizons, consistent with the variance decompositions in Table 6.

Figure 5: Variance Term Premia Parital Derivatives

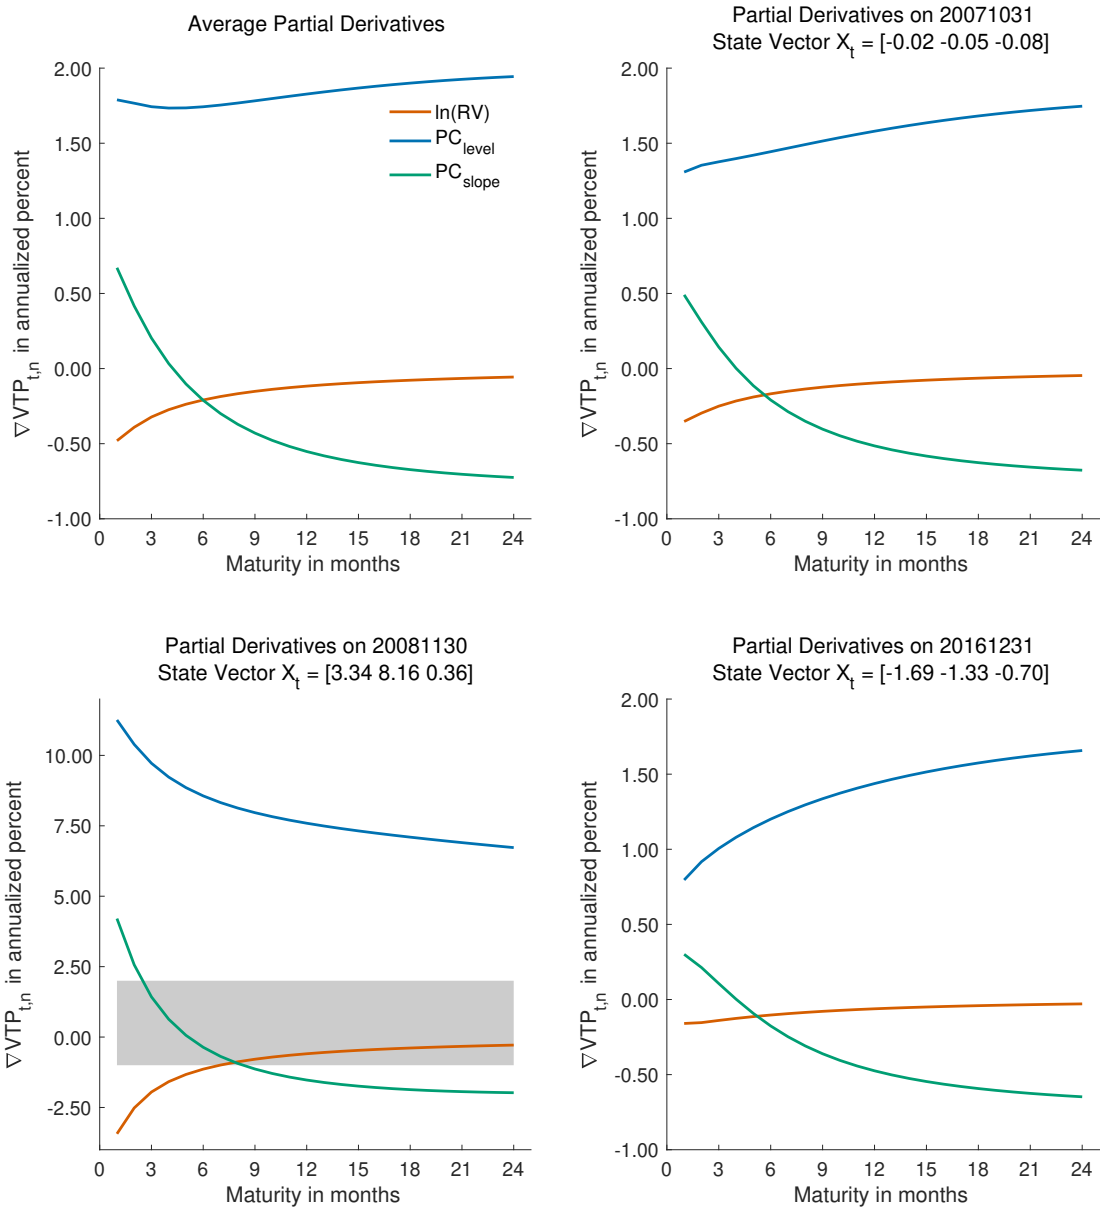

This figure plots the partial derivative of the variance term premia $\nabla VTP_{t,n}$ on average and at different points in time for a one standard deviation change in the state variables. The average partial derivatives in the top left subplot are similar to the coefficients from the regression analysis in Panel B of Table 6. The other plots highlight how the partial derivatives change for different values of the state vector. The top right subplot from October 2007 represents a period when the state vector is close to its unconditional mean $\hat{\mu}$ of zero. The bottom left subplot from November 2008 represents a high volatility state with an inverted variance swap curve. The bottom right subplot from December 2016 represents a low volatility state with an upward sloping variance swap curve. The gray box for the November 2008 plot highlights the scale of the other subplots, illustrating the increase in magnitude of the partial derivatives during the financial crisis.

Figure 6: Relative Pricing of VIX Futures

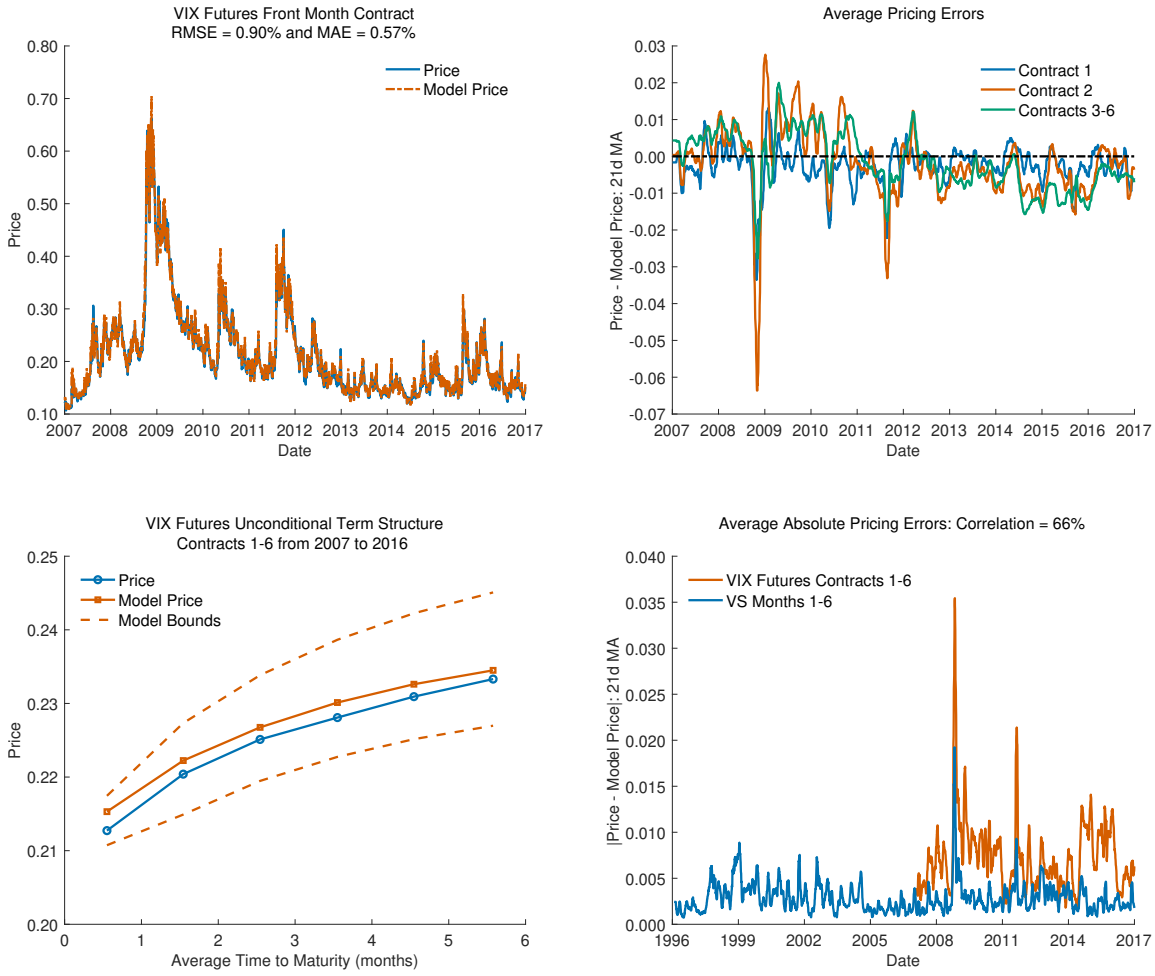

This figure plots the model fit for VIX futures after estimating the model with variance swap rate and realized variance data. The top left plot shows the model fit for the front month contract. The top right plot shows the pricing errors as a one-month moving average. The bottom left plot reports the unconditional term structure. The bottom right plot shows the magnitude of the pricing errors for VIX futures relative to the in-sample pricing errors for variance swaps (the MAE of .73% versus .29% is about 2.5 times larger). The model is relatively unbiased with VIX futures falling within the model bounds on average. In recent years, however, the cheapness of VIX futures has been as large as .50% to 1% according to the model.

Figure 7: VIX Futures Trading Strategy

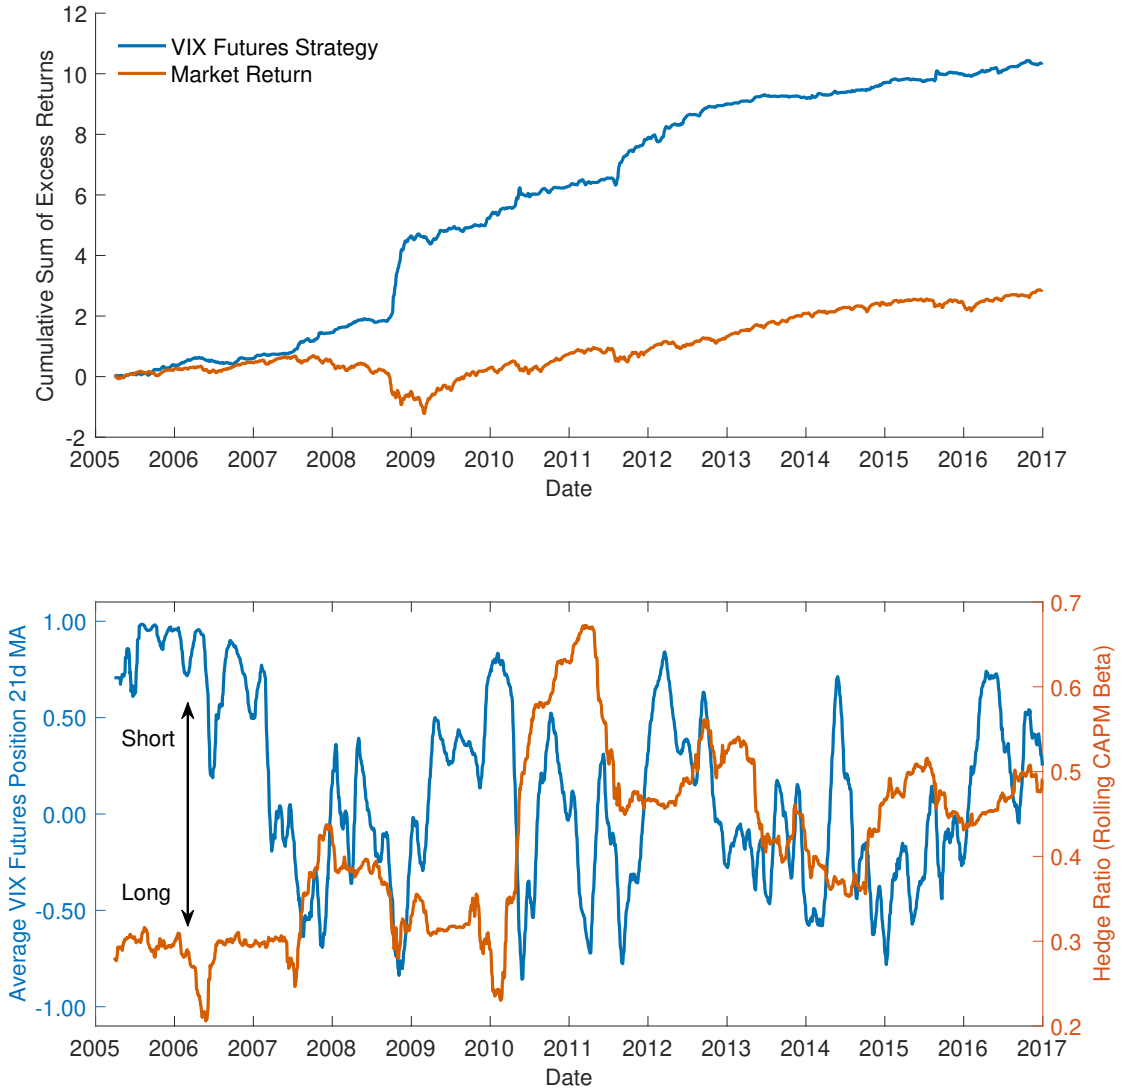

This figure plots the performance of a pseudo-out-of-sample VIX futures trading strategy described in the paper that exploits the return predictability of the model expected returns and pricing errors. The corresponding CRSP value-weighted return is included for comparison. Both return series are normalized to have 10% annualized volatility. The VIX futures trading strategy earns an annualized Sharpe ratio of 1.80 versus .50 for the market with a weekly CAPM alpha of .35% [5.21]. The bottom figure plots the average position across futures contracts to highlight when VIX futures are expensive (short) and cheap (long) and the rolling hedge ratio for the strategy.

A Appendix

A.1 Synthetic Variance Swap Rates

I compute synthetic variance swap rates from the price of a replicating portfolio that takes a static position in a continuum of out-of-the money European options (Carr and Wu 2009).[20] I perform this computation every day for standard expirations between ten calendar days and three years to maturity with at least five out-of-the money call and put options whose Black and Scholes (1973) deltas are greater than or equal to 1%.[21] For each date-maturity pair satisfying this filter, I fit a flexible implied volatility function by local linear regression to out-of-the money option prices with positive implied volatility as reported by OptionMetrics. I determine which options are out-of-the money using the forward rate implied by put-call parity.[22] To extrapolate beyond the observed strikes and for deep out-of-the money options with a delta less than 1%, I append log-normal tails using flat implied volatility functions. I then compute synthetic variance swap rates as a weighted average of out-of-the money option prices, using the fitted implied volatility functions to compute option prices. To obtain the term-structure on a constant grid, I interpolate the synthetic variance swap rates at the observed maturities onto a monthly grid from one-month to two-years. The interpolation is linear in total variance following Carr and Wu (2009) and the CBOE volatility indexes.

Figure A.1 provides an example of this procedure on July 31, 2015 for SPX options written on the S&P 500 Index. As the top plot illustrates, the implied volatility fitting errors are small while the range of observed moneyness is large. The close fit indicates that the implied volatility functions provide an accurate estimate of the risk-neutral distribution. The bottom plot reports the resulting synthetic variance swap rates alongside over-the-counter (OTC) rates from Markit Totem and the CBOE volatility indexes.[23] The rates from the

[20] The replicating portfolio is exact when there are no jumps and interest rates are constant. Martin (2017) shows that the assumption of a continuous underlying can be relaxed by computing realized variance with simple returns as opposed to log returns for the floating leg payoff, a result that is particularly important for pricing variance swaps on single stocks whose prices can go to zero in the event of default.

[21] I use traditional SPX options with an AM settlement on the third Friday of the month. In recent years, the CBOE has introduced SPX Weekly, End-of-Month, and PM options. Liquidity and contract specifications can differ for these products, introducing noise in implied volatility curve fitting. To mitigate this issue, I require expiration dates be on the third Friday or Saturday of the month. In addition, I require the first four characters of the option symbol field to match "SPX " from 2011 on. This differs from the "SPXW" symbol for weekly options that is used from February 2010 on in the US OptionMetrics data.

[22] I define the forward rate as the median forward rate implied from as many as ten strike prices that are closest to the strike price that minimizes the absolute difference between call and put prices. The forward rate implied by different strike prices is $F(\tau, K) = K + Z(\tau)^{-1}(C(\tau, K) - P(\tau, K))$ where $Z(\tau)$ is the risk free discount factor. Note that if call and put prices are equal, the implied forward rate is the strike price.

[23] In addition to the VIX which tracks the one-month implied volatility of SPX index options, the CBOE also tracks three-month and six-month implied volatility with the VIX3M (formerly VXV) and VXMT indexes.

different sources closely align with the knot points indicating the observed maturities. To show this example is representative of the full sample, Table A.1 reports summary statistics for the option prices used in the curve construction over time. Similar to the example, the range of moneyness is large and the implied volatility fitting errors are small. The average RMSE for the implied volatility fitting errors are less than .20% for short-dated maturities and .10% for long-dated maturities over the full sample. On an average day I am able to compute synthetic variance swap rates for as many as nine to ten maturities with a minimum maturity shorter than one-month and a maximum maturity longer than two-years. This range of maturities and the liquidity of the corresponding options, as evidenced by the large open interest amounts and small bid-ask spreads, supports my empirical approach which interpolates the observed synthetic rates onto a monthly grid from one-month to two-years for estimating the dynamic term-structure model.

Table A.2 illustrates the external validity of my approach by comparing my synthetic variance swap rates to synthetic rates from the CBOE and Bloomberg as well as to OTC rates from Markit Totem and a hedge fund. Across the different datasets, maturities, and sample periods, my synthetic rates closely track the variance swap rates from the alternative datasets. For example, the correlation between my one-month synthetic variance swap rate and the VIX is 96% in one-day changes, 98% in one-week changes, and 99% in one-month changes from 1996 to 2016. Similar results hold for the other maturities and datasets. Even for the OTC rates, the correlation is still as high as 96% to 99% in monthly changes across maturities. Beyond these large time-series correlations, Table A.2 also confirms that my synthetic rates are relatively unbiased. The lack of bias is perhaps surprising given the observation that OTC variance swaps embed additional jump, liquidity, and counterparty risks that are not present in the synthetic rates. It is possible that some of these differences offset. For example, negative jump risk and illiquid OTC markets might push OTC rates up while counterparty risk may push OTC rates down relative to synthetic rates. Empirically, my synthetic rates are within .20% implied volatility units of the OTC rates on average for maturities greater than one-month both for the hedge fund data from 2000:01-2013:11 and for the Markit data from 2006:09 to 2015. I do observe a larger difference at the one-month maturity, where my synthetic rate is .26% higher than the VIX from 1996 to 2016 and .62% higher than the OTC rates from Markit from 2006:09 to 2015. That said, this difference is still relatively small compared to option bid-ask spreads. For context, the average bid-ask spread for OTC variance swap rates declines from 1.39% at the one-month maturity to .60% for the one-year to two-year maturities according to Markit data. Average bid-ask spreads for index options are similar and equal to about .80% to 1% in recent years according to OptionMetrics data as reported in Table A.1. Finally, Table A.2 indicates that my synthetic

rates are positively skewed in comparison to the OTC rates from Markit, but less so in comparison to the other rates.

Figures A.2 and A.3 illustrate these results. Figure A.2 plots my synthetic variance swap rates against the CBOE volatility indexes at month-end and the daily change in my synthetic rates against the daily changes in the indexes from 2008 to 2016. As in Table A.2, the daily changes are over 95% correlated. In addition to these time series dynamics, Figure A.3 plots the average term-structure of the variance swap rates from 2009 to 2015 using month-end data for my synthetic rates, the CBOE volatility indexes, Bloomberg's synthetic rates, and Markit's OTC rates. Figure A.3 also plots the average bid-ask spread from Markit and 95% confidence intervals for the synthetic rates, which are computed separately for each maturity using Newey-West standard errors with 36 monthly lags. Overall, the differences between the average term-structures are relatively small when compared to either the bid-ask spread from Markit or the 95% confidence intervals. Of course, I should emphasize that the different rates are not expected to match each other exactly. OTC rates embed additional risks that are not present in the synthetic rates. In addition, when computing my synthetic variance swap rates, I avoid the strike truncation and discretization error that are present in the CBOE volatility indexes. While these differences could in principal drive large variations between my synthetic variance swap rates and the alternative datasets, I find empirically that the different rates tend to track each other quite closely.

A.2 Two-Scale Realized Variance Estimation

I estimate realized variance for the S&P 500 index each day in the sample following the two-scale approach described in Zhang et al. (2005). In the first step, I use a sparse five-minute sampling frequency to compute realized variance estimates from five different subsamples whose intraday observations are spaced one-minute apart. For example, if the market closes at 4:00pm, the realized variance estimate for the first subsample is the sum of squared log returns from the previous close at 4:00pm to 9:30am, 9:30am to 9:35am, ..., 3:55pm to 4:00pm. The realized variance estimate for the second subsample is the sum of squared log returns from the previous close at 4:00pm to 9:31am, 9:31am to 9:36am, ..., 3:56pm to 4:00pm. The realized variance for the third subsample is the sum of squared log returns from the previous close at 4:00pm to 9:32am, 9:32am to 9:37am, ..., 3:57pm to 4:00pm, etc. Each subsample uses the same starting and ending prices to estimate the daily realized variance. The two-scale realized variance estimate, or second stage estimate, is the average of the first stage realized variance estimates across subsamples to reduce sampling variability.

I apply standard data cleaning techniques for high frequency data when implementing

the two-scale estimator empirically. I use one-minute intraday prices during regular market hours from 9:30am to 4:00pm from TRTH. I filter these observations by dropping prices that are below the daily low or above the daily high as reported in TRTH's end-of-day data. In addition, I follow Liu et al. (2015) by excluding short days with fewer than 60% of the expected observations during regular market hours (days with less than 235 observations) to remove early closes from the sample. On each of the remaining days, I interpolate the observed prices onto a one-minute grid from 9:30am to 4:00pm using the previous tick method (previous neighbor interpolation).[24] I then compute the two-scale realized variance estimate as described above.

A.3 Variance Term Premia Estimation

A.3.1 VTP confidence intervals

Figure A.4 plots the one-month and twelve-month variance term premia alongside 95% pointwise confidence intervals. The top plots report NLLS confidence intervals from daily data that are block bootstrapped treating the state variables as observable. The bottom plots report Bayesian MCMC confidence intervals from non-overlapping monthly data allowing for latent state variables. I discuss Bayesian estimation of the model in the next section. The confidence intervals are small relative to the time variation in the variance term premia. The null hypothesis of a constant variance risk premium at a one-month or twelve-month horizon is easily rejected. Beyond illustrating how the variance term premia varies significantly over time, the plot also quantifies how much wider the confidence intervals are for the twelve-month versus one-month estimates and for the Bayesian versus NLLS estimates. For the MCMC (NLLS) estimates, the average lengths of the one-month and twelve-month confidence intervals are .76% (.60%) and 1.56% (1.25%). Longer forecast horizons and the MCMC estimates, which use month-end data and allow for latent state variables, produce wider confidence intervals.

A.3.2 MLE and MCMC Estimation

I estimate the model parameters in the paper by two steps. First, I estimate the physical parameters $(\hat{\mu}, \hat{\Phi}, \hat{\Sigma}_v)$ from a monthly vector autoregression with overlapping daily data. Second, I estimate the risk-neutral parameters $(\hat{\mu}^{\mathbb{Q}}, \hat{\Phi}^{\mathbb{Q}})$ by nonlinear least squares using

[24] On most days I observe a price every minute so the interpolation is not required. The mean (median) number of observations per day is 389.5 (391) out of $6.5 \cdot 60 + 1 = 391$ possible observations. In total there are 354 days with fewer than 391 observations, most of which occur earlier in the 1996 to 2016 sample period. On these days, the mean (median) number of observations is 368 (390).

daily variance swap rates. This approach easily accommodates daily data and only requires that the observation errors be mean zero conditioned on the state vector $E[e_{t,n}|X_t] = 0$. With the additional assumption that the measurement errors are conditionally normal, I can also estimate the model by maximum likelihood (MLE) and by Bayesian methods using a Markov Chain Monte Carlo (MCMC) algorithm. For these estimation strategies I use non-overlapping month-end data. As before, the model can be summarized as,

$$\begin{array}{rclcrcl} X_{t+1} & = & \mu + \Phi X_t + v_{t+1}, & & v_{t+1}|\mathcal{F}_t & \sim & N(0, \Sigma_v) \\ Y_{t,n} & = & g_n(X_t, \mu^{\mathbb{Q}}, \Phi^{\mathbb{Q}}, \Sigma_v) + e_{t,n}, & & e_{t,n}|X_t & \sim & N(0, \sigma_e^2). \end{array} \tag{35}$$

The likelihood function from the forecast error decomposition is,

$$\begin{array}{rcl} f(Y_t, X_t|X_{t-1}, \Theta) & = & f(Y_t|X_t, \Theta) f(X_t|X_{t-1}, \Theta) \\ & = & f(Y_t|X_t, \mu^{\mathbb{Q}}, \Phi^{\mathbb{Q}}, L_v) f(X_t|X_{t-1}, \mu, \Phi, L_v). \end{array} \tag{36}$$

The resulting log likelihood function (conditioned on $t = 0$ information) is,

$$\begin{array}{rcl} LL & = & \sum_{t=1}^{T} \ln f(Y_t|X_t, \Theta) + \sum_{t=2}^{T} \ln f(X_t|X_{t-1}, \Theta) \\ & = & -\frac{T \cdot N_\tau}{2} \ln(2\pi\sigma_e^2) - \frac{1}{2} \sum_{t=1}^{T} \sum_{n \in \tau} ((Y_{t,n} - g_n(X_t, \mu^{\mathbb{Q}}, \Phi^{\mathbb{Q}}, \Sigma_v))/\sigma_e)^2 \\ & & -\frac{T \cdot K}{2} \ln(2\pi) - \frac{T}{2} \ln|\Sigma_v| - \frac{1}{2} \sum_{t=1}^{T} (X_t - \mu - \Phi X_{t-1})' \Sigma_v^{-1} (X_t - \mu - \Phi X_{t-1}). \end{array} \tag{37}$$

The separation of the physical parameters that govern the conditional mean of the state vector and the risk-neutral parameters that govern variance swap pricing is emphasized by Joslin et al. (2011). Because of this separation, one can show that the maximum likelihood estimates for μ and Φ are the ordinary least squares estimates from a vector autoregression of the state variables. This greatly simplifies maximum likelihood estimation as the likelihood function only needs to be maximized over the remaining parameters $(\mu^{\mathbb{Q}}, \Phi^{\mathbb{Q}}, L_v)$. In addition, the separation implies that variance swap pricing errors do not inform the model's predictability for the state variables in the VAR. Forecasts are entirely driven by the estimated VAR under the physical measure. Of course, variance swap rates still impact the model's predictability as they are included in the state vector through the level and slope factors $X_t = [\ln RV_t \; PC_{level,t} \; PC_{slope,t}]$.

In addition to maximum likelihood, one can also estimate the model by MCMC methods following the literature on the Bayesian estimation of stochastic volatility models (Jacquier et al. 1994). In this case, the posterior distribution is,

$$\begin{array}{rcl} p(X, \Theta|VS) & \propto & p(Y, X, \Theta) \\ & = & p(Y|X, \Theta) p(X|\Theta) p(\Theta), \end{array} \tag{38}$$

which suggests an MCMC algorithm that cycles through:[25]

- $p(\Theta|Y,X)$ draw $\mu, \mu^{\mathbb{Q}}, \Phi, \Phi^{\mathbb{Q}}, L_v$ sequentially
- $p(X_t|Y_t, X_{t-1}, X_{t+1}, \Theta)$ draw X_t for $t = 1, \ldots, T$.

In implementing the estimation I use diffuse priors for $(\mu, \Phi, \mu^{\mathbb{Q}}, \Phi^{\mathbb{Q}})$ centered around the null of no return predictability. The advantage of the MCMC algorithm is that it accommodates latent factors and the nonlinear model for pricing variance swaps, addressing a potential criticism of the baseline estimation strategy. Variance swaps in the model are a nonlinear function of the state vector. While realized variance is priced exactly by the model, the level and slope factors are not. Despite this, I have assumed that the level and slope factors are observable, i.e. the model should price these factors exactly so that the factors can be found by inverting the model. Whether this criticism matters for the results is an empirical question. In brief, I find that it doesn't matter. Using observable or latent state variables delivers similar estimates of variance swap rates and variance term premia. Given that model estimation is significantly faster for the two-step approach described in the paper as compared to the MCMC estimation strategy, I assume the state variables are observable for the baseline analysis.[26] In particular, this helps to facilitate the detailed analysis of the model's out-of-sample return predictability and to allow for the out-of-sample VIX futures trading strategy.

Table A.3 summarizes these results by reporting the model parameters estimated by the different methods using month-end non-overlapping data. The results are similar across the estimation methods. The differences from the baseline parameter estimates are predominantly insignificant and economically small. Panel D illustrates this result for the MCMC estimates. Similar results hold for the NLLS and MLE estimates (not shown). The t-statistics for the MLE are from the asymptotic sandwich covariance matrix estimated numerically using double-sided finite differences for the Hessian and gradient. Figure A.5 plots the model variance swap rates and variance term premia from the different estimation methods alongside each other to interpret these differences from an economic perspective. The differences in the variance swap rates versus the baseline model are small and less than Markit's reported bid-offer spreads. The differences in the variance term premia estimates are also

[25] Drawing from $P(X_1|Y_1, \Theta)$ and $P(X_T|Y_T, X_{T-1}, \Theta)$ for the initial and final conditions.

[26] An alternative estimation approach to accommodate latent factors is to estimate the model using the Kalman filter with log variance swap forward rates in the observation equation. I originally pursued this approach but found that pricing log variance swap forwards resulted in larger pricing errors for variance swap rates, motivating the MCMC approach. The performance deterioration potentially stemmed from increased measurement noise when computing synthetic variance swap forward rates, which requires differencing the interpolated synthetic variance swap rates across the maturity dimension.

small in comparison to the confidence intervals discussed in the previous section. As a final observation, Figure A.6 plots the level $PC_{level,t}$ and slope $PC_{slope,t}$ factors from the baseline estimation against the posterior distribution of the latent state variables from the MCMC estimation. The highest posterior density (HPD) regions for the first latent factor are very tightly centered around $PC_{level,t}$. The HPD regions are somewhat wider for the slope factor, but still close $PC_{slope,t}$ used in the baseline estimation. Despite this uncertainty, the variance swap rates, term premia, and term premia confidence intervals are quite similar across Figures A.4 and A.5.

A.4 Model Extensions

A.4.1 Volatility swap rates and bounds on VIX futures prices

Bounds on VIX futures prices can be derived from Jensen's inequality (Carr and Wu 2006). Using the notation in this paper, the upper bound for the n-month futures contract is the square root of the $n+1$ month variance swap forward,

$$\begin{aligned} Fut_{t,n} &= E_t^{\mathbb{Q}}\left[\sqrt{VS_{t+n,1}}\right] \\ &\leq \sqrt{E_t^{\mathbb{Q}}\left[VS_{t+n,1}\right]} \\ &= \sqrt{F_{t,n+1}} \\ &\equiv UB_{t,n}. \end{aligned} \tag{39}$$

Proceeding along similar lines, the lower bound for the n-month futures contract is the price of an $n+1$ month volatility swap forward,

$$\begin{aligned} Fut_{t,n} &= E_t^{\mathbb{Q}}\left[\sqrt{E_{t+n}^{\mathbb{Q}}\left[RV_{t+n+1}\right]}\right] \\ &\geq E_t^{\mathbb{Q}}\left[E_{t+n}^{\mathbb{Q}}\left[\sqrt{RV_{t+n+1}}\right]\right] \\ &= E_t^{\mathbb{Q}}\left[\sqrt{RV_{t+n+1}}\right] \\ &= E_t^{\mathbb{Q}}\left[e^{\frac{1}{2}A_0+\frac{1}{2}B_0'X_{t+n+1}}\right] \\ &= E_t^{\mathbb{Q}}\left[Vol_{t+n+1}\right] \\ &= Fvol_{t,n+1} \\ &\equiv LB_{t,n}. \end{aligned} \tag{40}$$

Similar to variance swap forwards, volatility swap forwards are exponential affine in the state vector $Fvol_{t,n} = e^{A_n^{Vol}+(B_n^{Vol})'X_t}$. The coefficients A_n^{Vol} and B_n^{Vol} follow the same recursions as before with the adjusted initial conditions $A_0^{Vol} = \frac{1}{2}A_0$ and $B_0^{Vol} = \frac{1}{2}B_0$. Note at $n = 0$

the upper bound becomes an equality,

$$Fut_{t,0} = VIX_t = \sqrt{VS_{t,1}} = \sqrt{F_{t,1}}. \tag{41}$$

In contrast, the lower bound at $n = 0$ remains an inequality,

$$Fut_{t,0} = VIX_t = \sqrt{E_t^{\mathbb{Q}}[RV_{t+1}]} \geq E_t^{\mathbb{Q}}\left[\sqrt{RV_{t+1}}\right]. \tag{42}$$

Summarizing these results, VIX futures prices are bounded by,

$$LB_{t,n} \equiv Fvol_{t,n+1} \leq Fut_{t,n} \leq \sqrt{F_{t,n+1}} \equiv UB_{t,n}. \tag{43}$$

A.4.2 VIX Options

The VIX is conditionally log-Normal in the model. As a result, VIX options can be priced by the Black (1976) formula for pricing options on futures contracts. Assuming that interest rates are constant and equal to r, the price of an n-month call option on the n-month VIX futures contract is,

$$\begin{aligned} C_{t,n} &= E_t^{\mathbb{Q}}\left[e^{-r\cdot\frac{n}{12}}\left(Fut_{t+n,0} - K\right)^+\right] \\ &= E_t^{\mathbb{Q}}\left[e^{-r\cdot\frac{n}{12}}\left(VIX_{t+n} - K\right)^+\right] \\ &= E_t^{\mathbb{Q}}\left[e^{-r\cdot\frac{n}{12}}\left(e^{A_0^F + (B_0^F)'X_{t+n}} - K\right)^+\right]. \end{aligned} \tag{44}$$

The conditional distribution of the logarithm of the VIX is,

$$\ln VIX_{t+n} = A_0^F + (B_0^F)'X_{t+n} | \mathcal{F}_t \overset{\mathbb{Q}}{\sim} N(\mu_{t,n}, \sigma_n^2). \tag{45}$$

with mean and variance that are equal to,

$$\begin{aligned} \mu_{t,n} &= A_0^F + \left(B_0^F\right)'\left(\textstyle\sum_{i=1}^n \left(\Phi^{\mathbb{Q}}\right)^{i-1}\right)\mu^{\mathbb{Q}} + \left(B_0^F\right)'\left(\Phi^{\mathbb{Q}}\right)^n X_t \\ \sigma_n^2 &= \textstyle\sum_{i=1}^n \left(B_0^F\right)'\left(\Phi^{\mathbb{Q}}\right)^{n-i}\Sigma_v\left(\left(\Phi^{\mathbb{Q}}\right)^{n-i}\right)' B_0^F. \end{aligned} \tag{46}$$

It follows that VIX futures prices are equal to,

$$\begin{aligned} Fut_{t,n} &= E_t^{\mathbb{Q}}\left[VIX_{t+n}\right] \\ &= E\left[e^{\mu_{t,n}+\sigma_n Z}\right] \\ &= e^{\mu_{t,n}+\frac{1}{2}\sigma_n^2}. \end{aligned} \tag{47}$$

and that VIX options prices are equal to,

$$C_{t,n} = C_{Black}\left(Fut_{t,n}, K, \sigma_n\sqrt{\frac{12}{n}}, r, \frac{n}{12}\right). \tag{48}$$

The option price is expressed using the Black formula,

$$\begin{aligned} C_{Black}(F_t, K, \sigma, r, \tau) &= e^{-rT}\left[F_t \cdot N(d_1) - K \cdot N(d_2)\right] \\ d_1 &= \frac{\ln(\frac{F_t}{K}) + \frac{1}{2}\sigma^2\tau}{\sigma\sqrt{\tau}} \\ d_2 &= d_1 - \sigma\sqrt{\tau}. \end{aligned} \tag{49}$$

This result has several implications. First, note that the model admits an unconditional term structure of implied volatility for VIX options that is equal to $\sigma_n\sqrt{12/n}$ in annualized units. The corresponding term structure for the realized volatility of the VIX can be computed by replacing $\Phi^{\mathbb{Q}}$ with Φ in equation 46 above. Figure A.7 plots these term structures alongside their empirical counterparts. Overall, the model captures the level and downward sloping term structure for the volatility of volatility.

However, as Figure A.7 indicates, the volatility of the VIX has increased in recent years. The model assumes that the logarithm of realized variance follows a vector autoregression with homoskedastic shocks. As a result, the Black formula applies for pricing VIX options with the implication that the model features no time variation in the volatility of the VIX and no volatility smile for VIX options. These features of the model are rejected by the data. The CBOE VVIX Index provides direct evidence that the implied volatility of VIX options is time varying. In addition, VIX options feature an implied volatility smile that reflects an asymmetric conditional distribution under the risk-neutral measure. High strike call options that hedge against increases in the VIX tend to have a higher implied volatility than low strike call options. Figure A.7 illustrates these results. I leave extensions of the model to accommodate these features of the data to future work.

A.4.3 Continuous Time Model

The baseline approach assumes that the logarithm of realized variance follows a discrete time first order VAR with innovations that are conditionally Normal. An analogous model can be derived in continuous time. As a motivating example, suppose that spot log variance $\ln v_t = y_t$ follows a univariate Gaussian process,

$$dy_t = \kappa(\bar{y} - y_t)dt + \sigma dW_t^{\mathbb{Q}}, \tag{50}$$

with the conditional distribution,

$$y_T|y_t \overset{\mathbb{Q}}{\sim} N\left(\bar{y} + e^{-\kappa\tau}(y_t - \bar{y}),\ \frac{\sigma^2}{2\kappa}(1 - e^{-2\kappa\tau})\right), \tag{51}$$

where $\tau = T - t$. In this example, variance swap rates can be derived by applying Fubini's theorem and the moment generating function of a Normally distributed random variable to evaluate the expectation,

$$\begin{aligned} VS(t,T) &= E_t^{\mathbb{Q}}\left[\int_t^T v_s ds\right] \\ &= \int_t^T E_t^{\mathbb{Q}}\left[v_s\right] ds \\ &= \int_t^T E_t^{\mathbb{Q}}\left[e^{y_s}\right] ds \\ &= \int_t^T e^{\bar{y} + e^{-\kappa(s-t)}(y_t - \bar{y}) + \frac{\sigma^2}{4\kappa}(1 - e^{-2\kappa(s-t)})} ds. \end{aligned} \tag{52}$$

The integral above is analogous to the summation of exponential affine variance swap forwards from the discrete time model. Moreover, this approach readily extends to a multivariate setting.

Suppose there is a $K \times 1$ state vector X_t following the multivariate Gaussian process,

$$dX_t = \kappa(\mu - X_t) + \sigma dW_t, \tag{53}$$

with the instantaneous covariance matrix,

$$E_t^{\mathbb{Q}}\left[dX_t dX_t'\right] = \sigma\sigma' dt = \Sigma \cdot dt. \tag{54}$$

As before, the conditional distribution under the risk-neutral measure is,

$$X_T|X_t \overset{\mathbb{Q}}{\sim} N\left(\left(I - e^{-\kappa\tau}\right)\mu + e^{-\kappa\tau}X_t, \Sigma(\tau)\right), \tag{55}$$

where $\tau = T - t$ and $\Sigma(\tau) \equiv \int_0^\tau e^{\kappa(u-\tau)}\Sigma e^{\kappa'(u-\tau)} du$. If the spot variance is an exponential affine function of the state vector $\ln v_t = A_0 + B_0' X_t$, variance swap forward rates are equal to,

$$\begin{aligned} F(t,T) &= E_t^{\mathbb{Q}}\left[v_T\right] \\ &= E_t^{\mathbb{Q}}\left[e^{A_0 + B_0' X_T}\right] \\ &= e^{A_0 + B_0'\left(\left(I - e^{-\kappa\tau}\right)\mu + e^{-\kappa\tau}X_t\right) + \frac{1}{2}B_0'\Sigma(\tau)B_0} \\ &= e^{A_0 + B_0'\left(I - e^{-\kappa\tau}\right)\mu + \frac{1}{2}B_0'\Sigma(\tau)B_0 + B_0' e^{-\kappa\tau}X_t} \\ &\equiv e^{A(\tau) + B(\tau)' X_t} \end{aligned} \tag{56}$$

Variance swap rates are then equal to,

$$VS(t,T) = \int_t^T F(t,s)ds. \tag{57}$$

As a final observation, note that closed-form VIX futures prices are not available in the continuous time model. To see this, note that,

$$VIX_t = 100 \cdot \sqrt{12 \cdot VS\left(t, t+\frac{1}{12}\right)} = 100 \cdot \sqrt{12 \cdot \int_t^{t+\frac{1}{12}} e^{A(s-t)+B(s-t)'X_t} ds}. \tag{58}$$

The one-month variance swap rate is obtained by integrating over spot variance in continuous time, which prevents the exponential function from being able to absorb the convexity adjustment as it does in discrete time.

Table A.1: Index Option Summary Statistics

This table reports summary statistics for the out-of-the money SPX option prices that are used to construct the synthetic variance swap rates for the S&P 500 Index. Each day the minimum, average, and maximum maturity τ in months and the number of maturities N_τ are averaged across maturity buckets. Similarly, the minimum, average, and maximum absolute moneyness $|x| = |K/F - 1|$ and absolute Black-Scholes delta $|\Delta|$ are averaged across maturity-date pairs by maturity bucket along with the number of out-of-the money options N_{opt}, the root-mean-squared implied volatility fitting error $RMSE$ and option price bid-ask spread $SPD = (Ask - Bid)/Vega$ both equal-weighted and value-weighted by outstanding vega, and open interest OI and volume VLM reported in thousands of contracts and Black-Scholes vega in millions of dollars. The second half of the sample (2009-2016) features a larger number of option price observations and a larger quantity of option trading as measured by volume and open interest.

Sample Period	2000-2016			2009-2016			2000-2008		
Maturity (τ)	[0,3)	[3,9)	[9,36]	[0,3)	[3,9)	[9,36]	[0,3)	[3,9)	[9,36]
τ_{min}	0.82	3.99	10.63	0.82	3.55	10.40	0.82	4.38	10.84
τ_{avg}	1.62	5.64	17.61	1.64	5.39	18.56	1.61	5.85	16.76
τ_{max}	2.43	7.43	26.55	2.47	7.51	29.80	2.40	7.37	23.66
N_τ	2.62	2.42	3.70	2.66	2.85	4.38	2.58	2.03	3.09
N_{opt}	69.87	43.58	41.05	103.43	55.06	53.53	39.02	29.24	25.28
$\|x\|_{min}$	0.00	0.00	0.01	0.00	0.00	0.00	0.00	0.00	0.01
$\|x\|_{avg}$	0.10	0.16	0.23	0.11	0.17	0.26	0.09	0.15	0.19
$\|x\|_{max}$	0.25	0.40	0.57	0.27	0.42	0.64	0.23	0.36	0.49
$\|\Delta\|_{min}$	0.01	0.01	0.02	0.01	0.01	0.01	0.01	0.01	0.03
$\|\Delta\|_{avg}$	0.16	0.18	0.21	0.14	0.17	0.20	0.18	0.19	0.22
$\|\Delta\|_{max}$	0.49	0.50	0.52	0.50	0.50	0.53	0.49	0.49	0.49
$RMSE$	0.17	0.08	0.08	0.17	0.07	0.06	0.18	0.09	0.10
$RMSE_{vega}$	0.14	0.07	0.08	0.13	0.05	0.05	0.16	0.09	0.11
SPD	1.47	1.05	1.03	1.62	1.13	1.20	1.33	0.95	0.81
SPD_{vega}	1.04	0.81	0.81	1.04	0.82	0.89	1.05	0.80	0.72
OI	919.40	531.84	239.92	1250.65	624.23	289.80	614.88	416.44	176.91
OI_{vega}	92.09	118.61	91.44	134.07	144.77	116.48	53.50	85.93	59.79
VLM	87.00	18.19	4.09	115.99	21.73	4.64	60.35	13.78	3.40
VLM_{vega}	10.04	4.54	1.71	14.50	5.61	2.09	5.94	3.19	1.22

Table A.2: Synthetic Variance Swap Rates Versus Alternative Datasets

This table compares the estimated synthetic variance swap rates to over-the-counter rates from a hedge fund and Markit Totem as well as synthetic rates from Bloomberg and the CBOE. Summary statistics for the differences between the rates are reported below in annualized volatility units. For example, the standard deviation of the difference between my one-month synthetic rate and the VIX index is .47%. The mean and standard deviation of the differences are small across the various datasets and maturities. The synthetic rates are also highly correlated with the alternative datasets across maturities, sample periods, and horizons.

Maturity	1	3	6	9	12	18	24
Panel A: Hedge fund over-the-counter rates (daily data 2000:01-2013:11)							
Mean	0.67	0.06	-0.08	-0.19	-0.12	-0.09	0.08
Standard Deviation	1.18	0.73	0.57	0.54	0.53	0.62	0.68
Skewness	0.31	0.74	0.44	0.51	0.12	0.43	0.26
Minimum	-14.61	-5.95	-4.77	-3.81	-3.83	-2.85	-3.15
Median	0.62	0.01	-0.09	-0.20	-0.13	-0.08	0.07
Maximum	10.99	6.77	5.63	4.15	3.87	5.06	5.68
Correlation of monthly changes	0.96	0.97	0.97	0.97	0.97	0.97	0.96
Correlation of weekly changes	0.87	0.90	0.91	0.91	0.90	0.89	0.87
Correlation of daily changes	0.71	0.74	0.73	0.73	0.72	0.69	0.65
Panel B: Bloomberg synthetic rates (daily data 2008:11-2016:12)							
Mean	0.73	-0.06	-0.22	-0.26	-0.30	-0.47	-0.57
Standard Deviation	0.84	0.37	0.34	0.35	0.33	0.40	0.33
Skewness	0.73	-1.05	-1.17	-1.67	-0.43	-1.41	-1.01
Minimum	-4.23	-2.90	-3.02	-2.79	-2.11	-3.11	-2.32
Median	0.54	-0.05	-0.17	-0.20	-0.27	-0.39	-0.52
Maximum	5.14	2.28	1.50	1.49	1.69	0.69	0.56
Correlation of monthly changes	0.99	0.99	0.99	0.99	0.99	0.99	0.99
Correlation of weekly changes	0.98	0.98	0.98	0.98	0.97	0.97	0.97
Correlation of daily changes	0.94	0.94	0.94	0.93	0.92	0.92	0.92
Panel C: Markit over-the-counter rates (monthly data 2006:9-2015:12)							
Mean	0.62	0.02	-0.07	-0.12	-0.05	0.02	0.20
Standard Deviation	1.06	0.64	0.45	0.42	0.50	0.48	0.49
Skewness	6.35	4.34	2.11	0.78	0.70	0.25	0.28
Minimum	-0.91	-0.96	-0.81	-0.96	-1.10	-1.11	-0.91
Median	0.50	-0.10	-0.15	-0.18	-0.12	-0.02	0.14
Maximum	10.04	4.93	2.63	1.68	1.95	1.29	1.66
Correlation of monthly changes	0.96	0.98	0.99	0.99	0.99	0.99	0.99
Panel D: CBOE synthetic rates (daily data*)							
Mean	0.26	-0.10	-0.22				
Standard Deviation	0.47	0.44	0.31				
Skewness	0.29	3.80	0.86				
Minimum	-4.20	-2.31	-3.41				
Median	0.24	-0.11	-0.23				
Maximum	4.85	6.17	2.48				
Correlation of monthly changes	0.99	0.99	1.00				
Correlation of weekly changes	0.98	0.98	0.99				
Correlation of daily changes	0.96	0.96	0.96				

*VIX 1996:01-2016:12; VXV and VXMT 2008:01-2016:12

Table A.3: Model Estimation By Different Methods

This table reports the model parameters estimated by different methods using month-end non-overlapping data from 1996 to 2016 ($T = 252$). Panel A reports nonlinear least squares (NLLS) estimates with bootstrapped t-statistics. Panel B reports maximum likelihood estimates (MLE) with t-statistics from the asymptotic sandwich covariance matrix. Panel C reports MCMC results from a Bayesian estimation of the model with latent state variables. Panel D tests for a difference between the baseline estimates in Table 3 and the MCMC estimates using the baseline standard errors. Overall, the results are very similar across the different methods. As Panel D indicates, allowing the state variables to be latent has very little impact on the results. The differences in the coefficients are predominantly insignificant and economically small. Moreover, similar results hold for the differences between the baseline estimates and the NLLS and MLE estimates.

Panel A.I: Physical Parameters (NLLS bootstrap)

	μ	$\Phi_{1,1}$	$\Phi_{1,2}$	$\Phi_{1,3}$
$\ln(RV)$	0.00	0.32***	0.18***	0.41***
	[0.00]	[3.55]	[5.09]	[3.98]
PC_{level}	0.01	0.28*	0.80***	-0.13
	[0.13]	[1.88]	[14.00]	[-0.75]
PC_{slope}	-0.01	0.06	-0.06***	0.70***
	[-0.26]	[1.40]	[-3.32]	[12.79]

Panel A.II: Prices of Risk (NLLS bootstrap)

	Λ_0	$\Lambda_{1,1}$	$\Lambda_{1,2}$	$\Lambda_{1,3}$
$\ln(RV)$	-0.70***	0.30***	-0.10***	-0.07
	[-11.24]	[3.20]	[-2.66]	[-0.63]
PC_{level}	-0.03	0.29	-0.16**	0.29
	[-0.18]	[1.62]	[-2.17]	[1.28]
PC_{slope}	-0.04	0.06	-0.04	-0.05
	[-0.17]	[0.25]	[-0.36]	[-0.17]

Panel B.I: Physical Parameters (MLE)

	μ	$\Phi_{1,1}$	$\Phi_{1,2}$	$\Phi_{1,3}$
$\ln(RV)$	-0.00	0.34***	0.17***	0.36***
	[-0.02]	[3.81]	[5.51]	[4.08]
PC_{level}	0.01	0.27*	0.83***	-0.16**
	[0.46]	[1.81]	[16.49]	[-2.16]
PC_{slope}	-0.00	0.07	-0.05***	0.74***
	[-0.36]	[1.56]	[-3.26]	[16.64]

Panel B.II: Prices of Risk (MLE)

	Λ_0	$\Lambda_{1,1}$	$\Lambda_{1,2}$	$\Lambda_{1,3}$
$\ln(RV)$	-0.70***	0.32***	-0.10***	-0.13
	[-31.63]	[3.59]	[-3.21]	[-1.47]
PC_{level}	-0.06*	0.31**	-0.14***	0.25***
	[-1.78]	[2.01]	[-2.72]	[3.23]
PC_{slope}	-0.04***	0.10**	-0.04***	-0.04
	[-2.83]	[2.22]	[-2.68]	[-0.91]

Panel C.I: Physical Parameters (MCMC)

	μ	$\Phi_{1,1}$	$\Phi_{1,2}$	$\Phi_{1,3}$
$\ln(RV)$	0.03	0.33***	0.18***	0.41***
	[0.73]	[3.85]	[5.67]	[4.53]
PC_{level}	0.02	0.24**	0.84***	-0.19
	[0.28]	[1.99]	[18.46]	[-1.60]
PC_{slope}	-0.00	0.09*	-0.06***	0.70***
	[-0.16]	[1.88]	[-3.31]	[13.19]

Panel C.II: Prices of Risk (MCMC)

	Λ_0	$\Lambda_{1,1}$	$\Lambda_{1,2}$	$\Lambda_{1,3}$
$\ln(RV)$	-0.69***	0.30***	-0.10***	-0.15*
	[-20.18]	[3.61]	[-3.14]	[-1.65]
PC_{level}	-0.05	0.29**	-0.14***	0.17
	[-0.82]	[2.38]	[-3.13]	[1.29]
PC_{slope}	-0.03	0.10**	-0.05***	-0.05
	[-1.10]	[2.30]	[-2.88]	[-1.07]

Panel D.I: Difference from baseline (MCMC)

	μ	$\Phi_{1,1}$	$\Phi_{1,2}$	$\Phi_{1,3}$
$\ln(RV)$	-0.02	-0.09	0.03	-0.02
	[-0.66]	[-1.18]	[1.22]	[-0.22]
PC_{level}	-0.01	-0.17	0.04	0.12
	[-0.11]	[-1.35]	[1.18]	[1.29]
PC_{slope}	-0.00	-0.07**	0.02*	0.08**
	[-0.12]	[-2.34]	[1.85]	[2.42]

Panel D.II: Difference from baseline (MCMC)

	Λ_0	$\Lambda_{1,1}$	$\Lambda_{1,2}$	$\Lambda_{1,3}$
$\ln(RV)$	0.01	-0.09	0.03	0.03
	[0.12]	[-1.46]	[1.19]	[0.51]
PC_{level}	-0.02	-0.17	0.06	0.20*
	[-0.20]	[-1.53]	[1.54]	[1.83]
PC_{slope}	-0.02	-0.06	0.02	0.08
	[-0.28]	[-0.56]	[0.75]	[0.76]

Table A.4: Robustness of Straddle Return Predictability using Various Definitions to Compute Straddle Returns

The estimated variance swap expected returns $\hat{E}_t[R_{t+h,n}]$ provide significant forecasts for straddle returns using various definitions to compute daily straddle returns. The results are out-of-sample in the sense that the model is estimated with realized variance and variance swap data, not straddle returns. Newey-West t-statistics indicate significance using $21 \cdot h \cdot 3$ lags to account for overlapping observations from daily data. The sample period is 1996 to 2016 ($n = 5,190$) and the return horizon is one month ($h = 1$).

Straddle Return Predictability Regressions ($h = 1$): $R_{t+h,b} = \beta_0 + \beta_1 \hat{E}_t[R_{t+h,n}] + \epsilon_{t+h,b}$

Straddle Maturity Bucket b	(1,3]	(3,6]	(6,9]	(9,15]	(15,24]	Average
Expected Return Maturity n	3	6	9	12	18	Average
Daily Returns: $R_{t+1} = (S_t - S_{t+1} - \Delta_t(P_{t+1} - P_t))/K$						
β_0	-0.00	-0.14	-0.18	-0.23*	-0.27**	-0.19
	[-0.02]	[-1.10]	[-1.35]	[-1.74]	[-2.27]	[-1.44]
β_1	1.49***	1.42***	1.22***	1.25***	1.15***	1.39***
	[5.79]	[5.15]	[4.35]	[4.47]	[4.97]	[4.81]
R^2_{adj}	0.06	0.10	0.10	0.11	0.11	0.10
Daily Returns: $R_{t+1} = (S_t - S_{t+1} - \Delta_t(P_{t+1} - P_t) + rf_t \cdot (S_t + K))/K - rf_t$						
β_0	0.01	-0.13	-0.16	-0.20	-0.24*	-0.17
	[0.09]	[-0.96]	[-1.18]	[-1.52]	[-1.96]	[-1.26]
β_1	1.49***	1.42***	1.21***	1.24***	1.13***	1.38***
	[5.79]	[5.14]	[4.32]	[4.43]	[4.87]	[4.78]
R^2_{adj}	0.06	0.10	0.10	0.11	0.11	0.10
Daily Returns: $R_{t+1} = (S_t - S_{t+1} - \Delta_t(P_{t+1} - P_t) + rf_t \cdot (S_t + K - \Delta_t P_t))/K - rf_t$						
β_0	0.01	-0.12	-0.15	-0.20	-0.23*	-0.16
	[0.12]	[-0.94]	[-1.15]	[-1.48]	[-1.91]	[-1.22]
β_1	1.49***	1.42***	1.21***	1.24***	1.13***	1.38***
	[5.79]	[5.14]	[4.32]	[4.42]	[4.84]	[4.76]
R^2_{adj}	0.06	0.10	0.10	0.11	0.10	0.10
Daily Returns: $R_{t+1} = (S_t - S_{t+1})/S_t$						
β_0	3.21**	-0.15	-0.62	-0.99	-1.19**	-0.10
	[2.16]	[-0.14]	[-0.81]	[-1.54]	[-2.39]	[-0.11]
β_1	11.41***	9.45***	6.16***	5.45***	4.37***	7.80***
	[3.52]	[4.29]	[4.68]	[4.75]	[5.33]	[4.37]
R^2_{adj}	0.02	0.05	0.05	0.07	0.08	0.05
Daily Returns: $R_{t+1} = (S_t - S_{t+1} - \Delta_t(P_{t+1} - P_t))/S_t$						
β_0	3.07**	-0.09	-0.76	-1.08	-1.26**	-0.21
	[2.08]	[-0.09]	[-0.94]	[-1.63]	[-2.44]	[-0.23]
β_1	11.16***	9.16***	6.27***	5.59***	4.54***	7.93***
	[3.46]	[4.11]	[4.22]	[4.62]	[5.18]	[4.26]
R^2_{adj}	0.02	0.05	0.06	0.07	0.08	0.05

Figure A.1: Synthetic Variance Swap Rate Construction

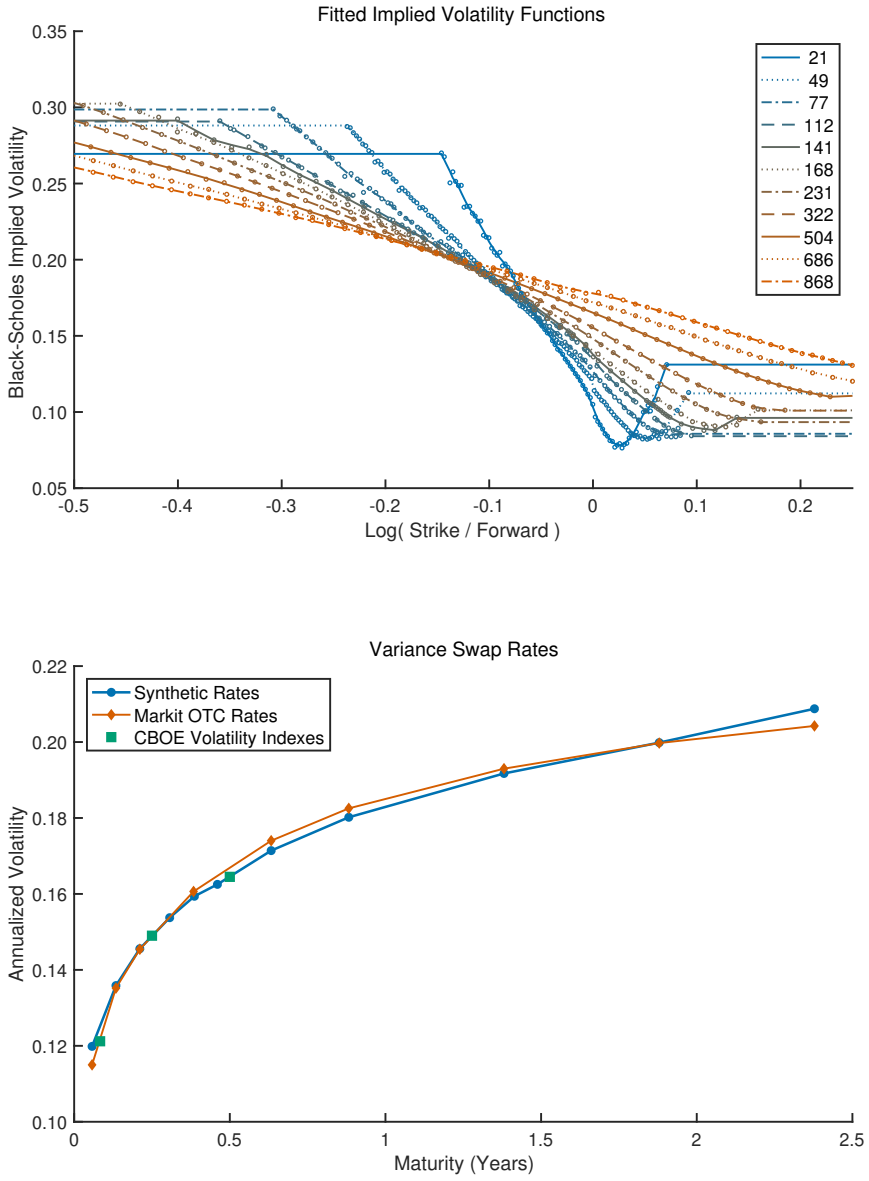

This figure illustrates my computation of synthetic variance swap rates for the S&P 500 Index on July 31, 2015. The top plot reports the fitted implied volatility functions against log-moneyness for different expirations whose time-to-maturity is reported in the legend. For each maturity, I compute the synthetic variance swap rate as a weighted average of out-of-the money option prices following Carr and Wu (2009). As the plot indicates, I extrapolate the price deep-out-of-the-money options with Black-Scholes deltas less than 1% by appending log-Normal tails with flat implied volatility functions. The bottom plot shows the resulting synthetic variance swap rates alongside over-the-counter rates from Markit and the CBOE volatility indexes. When estimating the model, I interpolate between these synthetic rates at the observed maturities onto a monthly grid from one-month to two years.

Figure A.2: Synthetic Variance Swap Rates Versus the CBOE Indexes

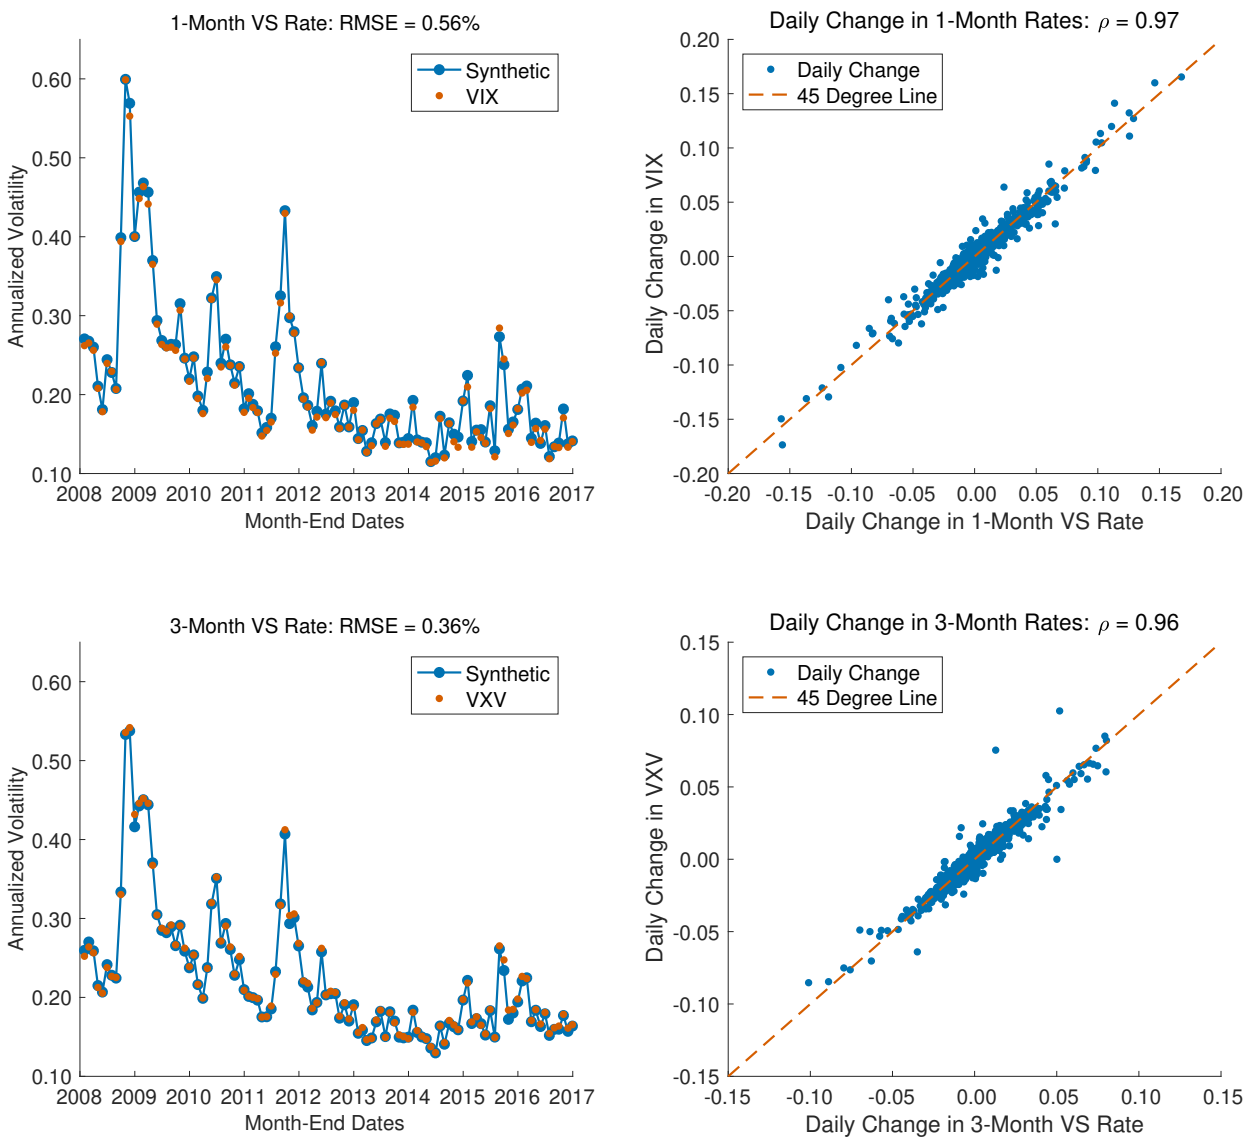

This figure plots my one-month and three-month synthetic variance swap rates against the CBOE volatility indexes from 2008 to 2016. The left plots report the time series dynamics at month-end dates. The right plots report the daily changes in my synthetic rates against the daily changes in the CBOE indexes which are over 95% correlated. The plots highlights the synthetic variance swap rates move and CBOE indexes move with a nearly perfect correlation throughout the sample. As Table A.2 confirms, similar results hold across the other maturities and datasets.

Figure A.3: Unconditional Variance Swap Term Structure from 2009-2015

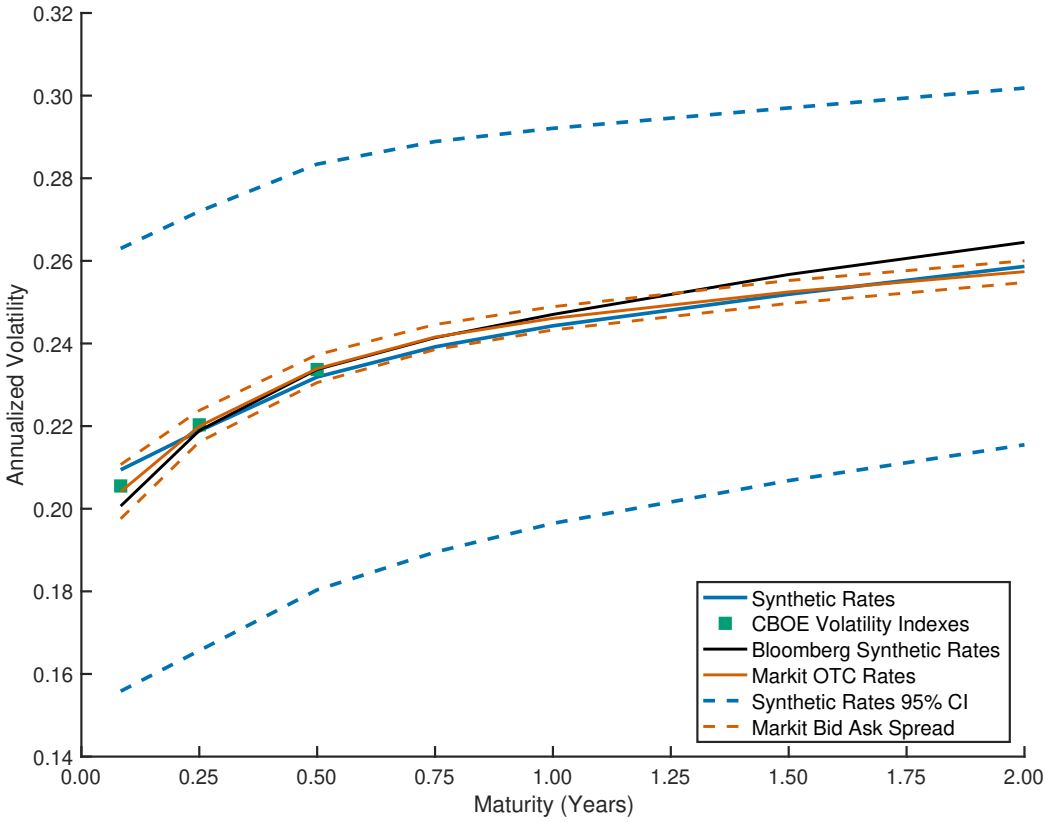

This figure plots the average term-structure of variance swap rates from 2009 to 2015 to compare my synthetic variance swap rates to the synthetic rates from the CBOE and Bloomberg as well as the over-the-counter rates from Markit. The plot also reports the average bid-ask spread from Markit as well as 95% confidence intervals for my synthetic rates, which are computed separately for each maturity using Newey-West standard errors with 36 lags. As the plot indicates, the variance swap rates from the different datasets are within the typical bid-ask spread from the Markit data, and they are well within the 95% confidence intervals for the synthetic rates. That said, the plot also highlights where some differences occur. For example, my synthetic rates are about .50% higher than the other rates at the one-month maturity, consistent with the results in Table A.2. The Bloomberg rates also appear to be about .50% higher than the other rates at the long end of the curve. Overall, however, these differences are relatively small in comparison to Markit's average bid-ask spreads.

Figure A.4: Variance Term Premia Confidence Intervals

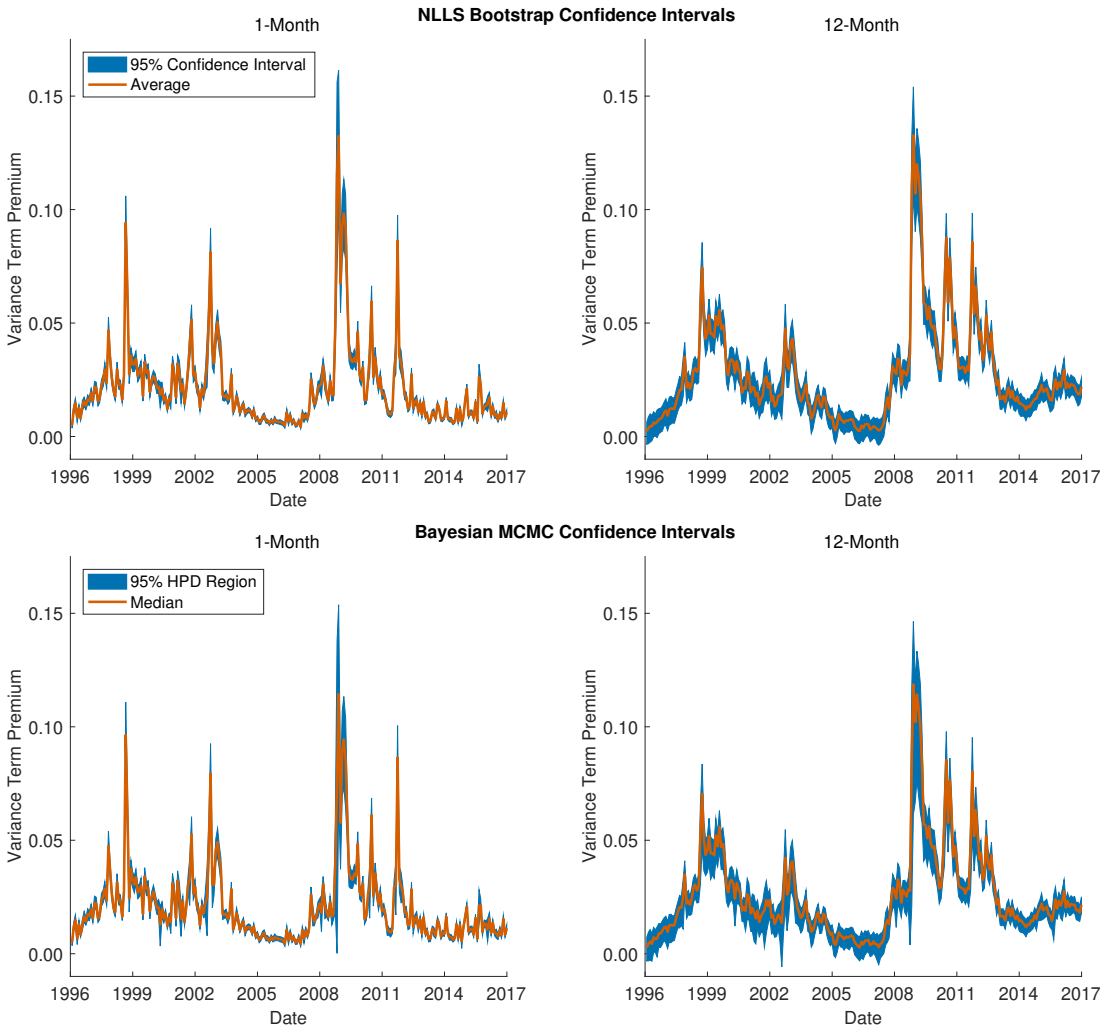

This figure plots the variance term premium alongside 95% confidence intervals for the one-month and twelve-month maturities at month-end dates from 1996 to 2016. The top plots report the baseline estimates that are obtained by nonlinear least squares using daily data with block bootstrapped 5th and 95th quantiles treating the state variables as observable. The bottom plots report the Bayesian MCMC estimates from non-overlapping monthly data that allow for latent state variables. While the confidence intervals can be wide at times, the null hypothesis of a constant variance risk premium is easily rejected. There is substantial time variation in the variance term premium relative to both the bootstrapped and the MCMC confidence intervals. In addition, the plots illustrate that the incremental uncertainty that stems from allowing for latent state variables is economically small.

Figure A.5: Model Variance Swap Rates and Term Premia: Robustness Across Estimation Methods

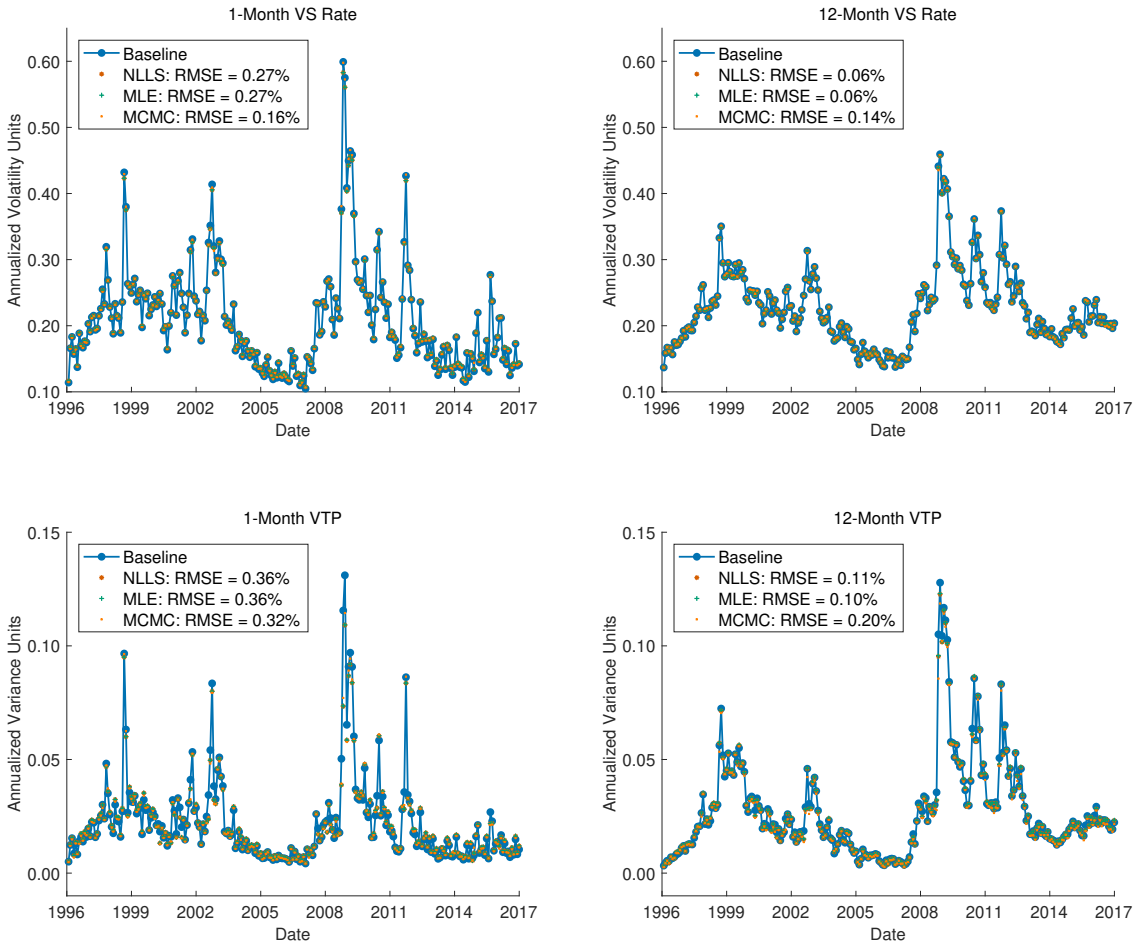

This figure plots the estimated one-month and twelve-month variance swap rates and term premia at month-end dates from 1996 to 2016. The baseline estimates using the parameters from Table 3 are plotted in blue. The alternative estimates from Table A.3 are plotted in red, green, and orange for the nonlinear least squares (NLLS), maximum likelihood (MLE), and Bayesian (MCMC) estimates respectively. The plot confirms that the small differences in the parameter estimates from the different methods do not translate into economically meaningful differences in variance swap rates or term premia. The mean absolute errors (MAE) in the legends are all small. Moreover, there is also a small difference between using observable versus latent state variables. The MCMC estimates are the median from the posterior distribution allowing the state variables to be latent. In contrast, the other estimates plug in the estimated parameters assuming the state variables are observable.

Figure A.6: Observable versus Latent State Variables

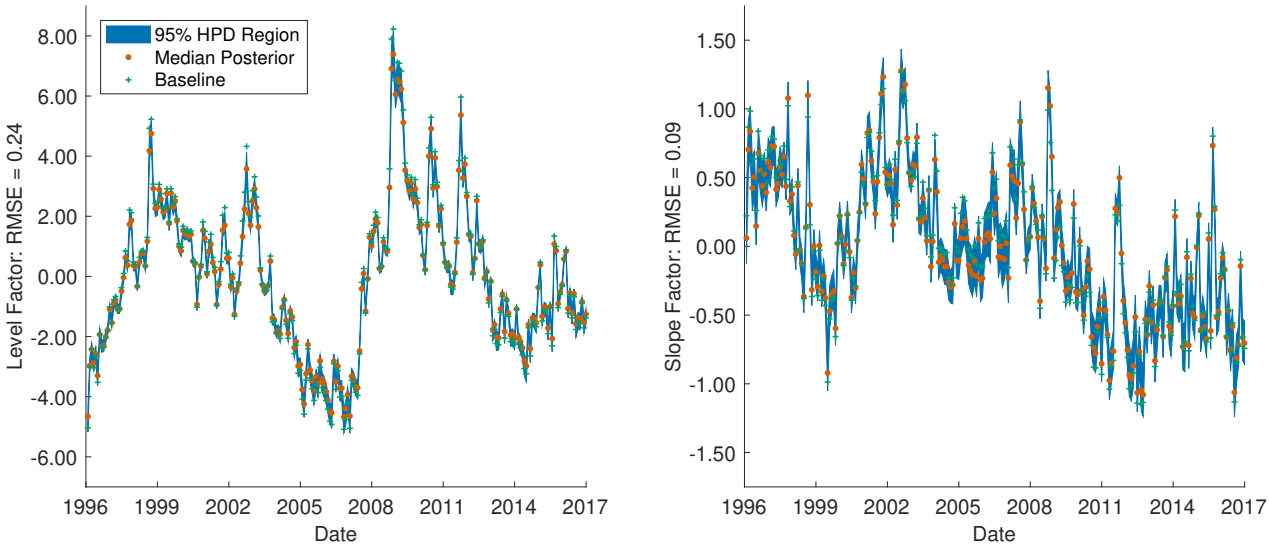

This figure plots the standardized logarithm of the first two principal components of variance swap rates denoted as level and slope that are used in the baseline estimation The plot also reports the smoothed estimates of the corresponding latent variables from the Bayesian MCMC estimation including the posterior median and the 95% highest posterior density region. The smoothed median is very close to the observed principal components. In addition, the highest posterior density regions fall very tightly around the observed principal components, particularly for the level factor.

Figure A.7: Volatility of Volatility

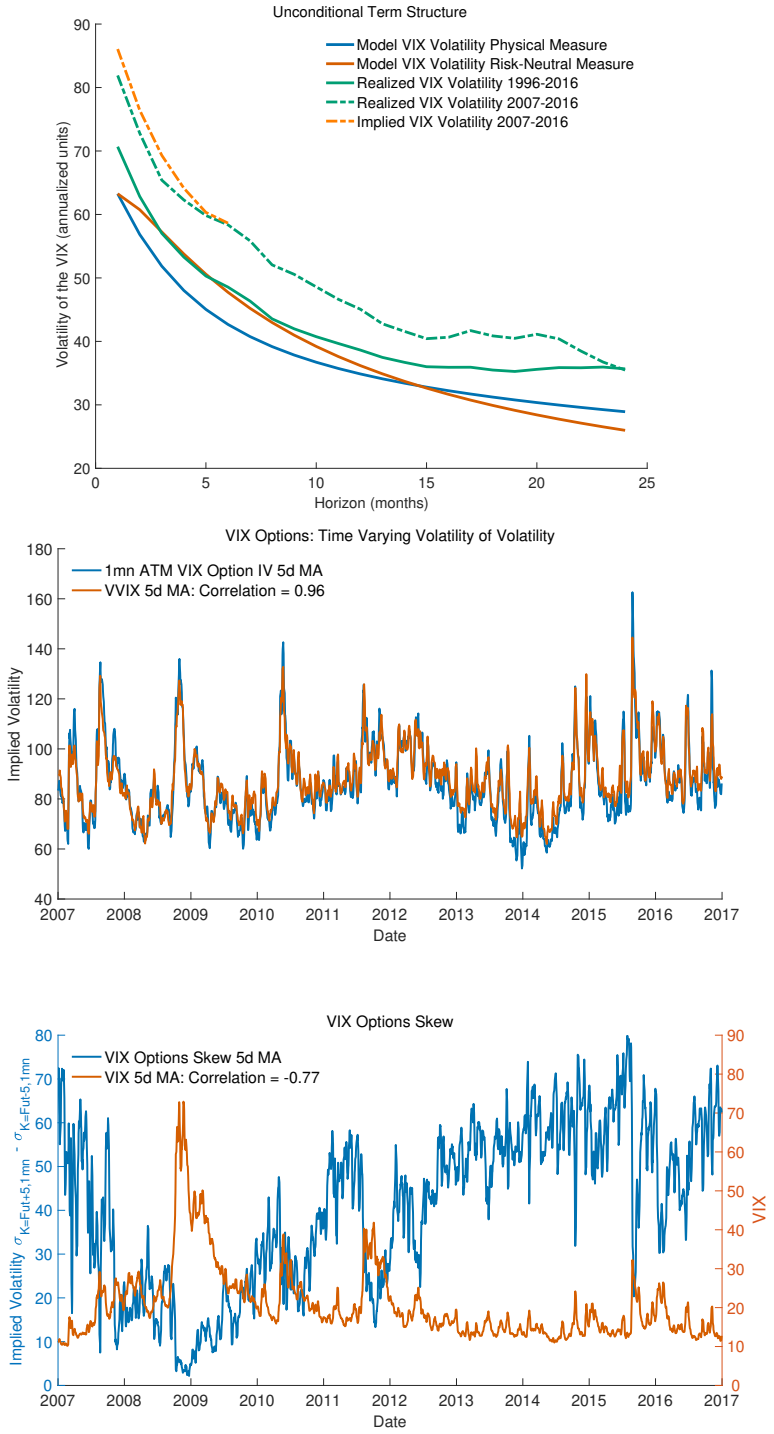

The top plot reports the unconditional term structure of volatility for the VIX in the model against the data. The middle plot reports the VVIX Index alongside the interpolated one-month at-the-money implied volatility for VIX options. The bottom plot reports VIX options skew, measured as the difference in implied volatility for one-month options with high and low strike prices: $K_{High} = Fut + 5$ and $K_{Low} = Fut - 5$.