STAFF REPORTS

NO. 867
OCTOBER 2018

REVISED
APRIL 2022

Equity Volatility Term Premia

Peter Van Tassel

FEDERAL RESERVE BANK of NEW YORK

Equity Volatility Term Premia
Peter Van Tassel
Federal Reserve Bank of New York Staff Reports, no. 867
October 2018; revised April 2022
JEL classification: G12, G13, C58

Abstract

This paper estimates the term structure of volatility risk premia for the stock market. By modeling the logarithm of realized variance, the paper derives a closed-form relationship between the prices of variance swaps and VIX futures. Term premia estimates for realized variance and implied volatility predict variance swap and VIX futures returns. When systematic risk increases, realized variance term premia increase but implied volatility term premia decline or exhibit a muted response. VIX futures are cheaper than the option-implied model prices by .50 percent on average and they tend to cheapen even further relative to the model when the VIX increases.

Key words: term structure, volatility, term premia, variance swap, VIX futures, 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, Ing-Haw Cheng, Richard Crump, Thomas Eisenbach, Robert Engle, Bjorn Eraker, Cam Harvey, David Lando, David Lucca, Lasse Pedersen, Ivan Shaliastovich, George Tauchen, Erik Vogt, and seminar participants at Duke University, the Federal Reserve Bank of New York, the NYU Stern QFE Seminar, the OptionMetrics Research Conference, the Wisconsin-Madison Junior Finance Conference, and Copenhagen Business School for comments.

1 Introduction

This paper develops a new approach for estimating the term-structure of volatility risk premia. By modeling realized variance as log-affine in discrete time, the paper derives a closed-form relationship between the prices of variance swaps and VIX futures. Model estimates of volatility term premia measure the time-varying cost of insuring against realized variance and implied volatility shocks over different horizons. Traders obtain exposures to these shocks for numerous reasons such as risk management, earning volatility risk premia, obtaining leverage, and arbitrage trading, among others.

The paper estimates the log term-structure model over a 25-year sample period from 1996 to 2020 using option-implied variance swaps and realized variance data. The log transformation guarantees that the model prices and volatility forecasts are non-negative.[1] In contrast, affine and quadratic models that are prevalent in the literature provide closed-form prices for variance swaps but not for VIX futures, and they do not guarantee non-negative prices without restrictive assumptions. Model estimation is fast and tractable by applying insights from the regression-based methodology from Adrian et al. (2015). A three-factor model is selected by standard information criteria. This benchmark model matches realized variance dynamics exactly and performs well at pricing variance swaps and forecasting variance swap returns.

The paper considers two empirical applications of the model. The first application studies the risk-return tradeoff in volatility markets. When analyzing the dynamics of volatility risk premia, many studies do not distinguish between the risk premium for realized variance and the risk premium for implied volatility. Variance swap studies often focus on variance swap expected returns which measure the realized variance risk premium (Egloff et al. 2010; Filipović et al. 2016; Amengual and Xiu 2018; Aït-Sahalia et al. 2020). Other studies focus on the implied volatility risk premium (Cheng 2018) or the realized variance risk premium (Lochstoer and Muir 2022) in isolation. Dew-Becker et al. (2017) is an exception which emphasizes the different types of volatility risk premia, finding that realized variance earns a much larger unconditional risk premium than implied variance as measured by the average returns of variance swap forwards.

The model in this paper provides a straightforward way to estimate the conditional risk premium for realized variance and implied volatility in a way that is consistent with the absence of arbitrage. Figures 1 and 2 plot the realized variance term premia (RVTP) and implied volatility

[1]Negative variance swap rates or VIX futures prices are arbitrage opportunities similar to zero lower bound violations in fixed income settings.

term premia (IVTP) estimates. To validate their interpretation as risk premia measures, the paper shows that the RVTP and IVTP estimates significantly predict variance swap and VIX futures returns.[2] By studying the time-series dynamics of the estimates, the paper finds that there is a positive risk-return tradeoff for realized variance but not for implied volatility. When systematic risk increases, the RVTP estimates increase across maturities and then decline slowly over time. In contrast, the IVTP estimates exhibit a negative or muted response for maturities up to six months, increasing only at a lag. These findings are novel because they are inconsistent with the prediction from many equilibrium models of a positive risk-return tradeoff for realized and implied volatility. The analysis highlights the importance of testing whether different models are consistent with the separate dynamics of realized variance and implied volatility risk premia.

The second application in the paper studies relative pricing across volatility markets, contributing to the limits-to-arbitrage literature (Gromb and Vayanos 2010). This analysis directly benefits from the model's closed-form pricing relationships. After being estimated with option-implied variance swap rates and realized variance data, the model is applied to price VIX futures. Figure 3 illustrates the model-based pricing error for the front-month VIX futures contract. While the model tracks the futures price quite well throughout most of the sample, the analysis reveals differences in average prices and periods with large dislocations in the prices of relative claims. VIX futures are found to be cheaper than the option-implied model prices by .50% on average and exhibit pricing errors as large as several percentage points at times. When the VIX increases, futures tend to cheapen relative to the model, exacerbating the anomalous response of implied volatility risk premia to increases in risk.

The paper compares the term-structure model's ability to price VIX futures to a non-parametric approach that uses index options and VIX options following Hülsbusch and Kraftschik (2018) and Park (2020). The pricing errors from the model are highly correlated with the pricing errors from the non-parametric approach, providing external validity to the analysis. In addition, the pricing errors from the model are found to be smaller in magnitude for longer-dated contracts and more significant at forecasting VIX futures returns than the pricing errors from the non-parametric approach. The results illustrate the ability of the model to identify valuation differences across the VIX futures and index options markets. Compared to the non-parametric approach, the term-structure model is particularly valuable for pricing longer-dated contracts because it does not rely on the use of less

[2] The return predictability results are significant across contract maturities and forecast horizons and hold in-sample and out-of-sample. The results are robust to accounting for conditional heteroskedasticity, Stambaugh (1999) bias, and Hodrick (1992) standard errors following the approach in Johnson (2019).

liquid or unavailable VIX options prices.

To interpret the results in more detail, one can observe that the behavior of the IVTP estimates is unexpected. A natural hypothesis is that the risk premium for selling volatility should be increasing in systematic risk due to the leverage effect for both variance swaps and VIX futures (Black 1976).[3] In theories for why investors hedge volatility, risk premia for realized variance and implied volatility are increasing in the economy's systematic risk factors such as volatility-of-volatility and jump risk (Bollerslev et al. 2009; Drechsler and Yaron 2011; Eraker and Wu 2017). In these models, volatility term premia exhibit a positive time-series correlation with the VIX. For example, in the Eraker and Wu (2017) model, volatility term premia for realized variance and implied variance increase after a shock to the VIX and then decline slowly over time, as the paper demonstrates. More broadly, a similar result holds in any theoretical model in which the returns from selling volatility are negatively correlated with the stochastic discount factor, so that the negative correlation condition (NCC) from Martin (2017) is satisfied.

The behavior of the IVTP estimates also relates to recent studies on the dynamics of volatility risk premia such as Cheng (2018) and Lochstoer and Muir (2022). Cheng (2018) finds that the risk premium for selling VIX futures is decreasing in various measures of systematic risk and provides evidence that demand effects are contributing to the unexpected behavior. If volatility hedging demand falls or if hedgers take profit when risk increases, downward demand pressure on futures prices may dampen the implied volatility risk premium. Exploring a different channel, Lochstoer and Muir (2022) consider a model in which the representative agent has slow-moving beliefs about stock market volatility that leads to initial underreaction to volatility shocks followed by delayed overreaction. Embedding these subjective beliefs into an otherwise standard long-run risk model results in volatility term premia that initially decrease and go negative after a shock to expected variance and then rise and become positive at a lag.

Compared to the reduced-form statistical approach in Cheng (2018), this paper finds similar results for implied volatility risk premia in a fully-fledged no-arbitrage model. Since the term-structure model is estimated without VIX futures data and because the sample extends back to 1996 before the start of VIX futures trading in 2004, it seems unlikely that demand for VIX futures is driving the IVTP estimates. At the least, the impact of VIX futures demand on the IVTP

[3] The leverage effect refers to the negative correlation between increases in volatility and changes in stock prices. Kalnina and Xiu (2017) show how to non-parametrically estimate the leverage effect using prices alone and using the VIX as a volatility instrument. From a factor model perspective, selling volatility should earn a positive risk premium as the returns from selling variance swaps and VIX futures load positively on stock market returns (Merton 1973; Ross 1976).

estimates would need to arise through an indirect channel that overcomes the much larger size of the index options market relative to the VIX futures market throughout most of the sample period.[4] Demand for VIX futures is more likely an important factor for the relative pricing analysis, where declining futures price relative to the model when the VIX increases is consistent with the demand channel that Cheng (2018) identifies.

An alternative possibility is that the mechanism in Lochstoer and Muir (2022) explains the IVTP estimates. However, there is a difficulty with this description as well. In Lochstoer and Muir (2022), realized and implied volatility risk premia are driven by the same, single-factor, which is the mistake agents make in their variance expectation. When there is a shock to expected variance in the Lochstoer and Muir (2022) model, the volatility risk premium declines, resulting in a negative risk-return tradeoff for both types of volatility shocks. In contrast, the RVTP estimates from the term-structure model exhibit a positive risk-return tradeoff that differs from the negative or muted tradeoff for the IVTP estimates. The empirical results in this paper suggest that the dynamics of realized variance and implied volatility risk premia may be less intertwined than what is predicted by many models.

From a historical time-series perspective, modeling the logarithm of realized variance is motivated by numerous studies including French et al. (1987) who show that the log-transformation reduces skewness, Nelson (1991) who develops an exponential ARCH model, and Andersen et al. (2003) and Andersen et al. (2007) who forecast volatility using nonlinear transformations of realized variance including the log transformation, which empirically tends to be approximately unconditionally Normally distributed across asset classes. While Box-Cox transformations can further reduce the skewness of realized variance in finite samples (Gonçalves and Meddahi 2011), this paper focuses on the log transform as a special case because it allows for closed-form pricing of variance swaps and VIX futures.

The literature on volatility term premia and variance swap pricing have received significant attention in recent years. In fixed income, a large literature on term-structure modeling decomposes yield curves into expected paths for interest rates and term premia components while seeking to determine the number of factors that drive the yield curve (Dai and Singleton 2003; Piazzesi 2010). Building on these results, a more recent literature focuses on variance swap term-structures, esti-

[4]S&P 500 index options had an average open interest of $4.65 billion in Black-Scholes-Merton vega in 2020. This was over 15-times larger than the average VIX futures open interest of 302 thousand contracts or $302 million of "vega" for one-point changes in VIX futures prices given the contract multiplier of $1000. Since 2010 the index options market has been around 9-times larger than the VIX futures market as measured by average open interest in units of vega. Earlier in the sample the index options market was even larger on a relative basis.

mating affine and quadratic models that are adapted from fixed income settings (Egloff et al. 2010; Filipović et al. 2016; Amengual and Xiu 2018; Aït-Sahalia et al. 2020). The variance swap models are applied to study portfolio choice problems and to estimate expected returns for variance swaps and the stock market.

Among existing variance swap models, Egloff et al. (2010), Dew-Becker et al. (2017), and Aït-Sahalia et al. (2020) estimate two-factor affine models with spot- and long-run volatility factors to match variance swap rates and a time-series factor that is either realized variance or stock market returns. Giglio and Kelly (2017) estimate a two-factor affine model to price variance swaps, but do not target realized variance dynamics. Filipović et al. (2016) develop a class of quadratic variance swap models and find evidence in favor of a bivariate model without volatility jumps that is estimated to match variance swap prices and variance swap and S&P 500 return dynamics. Closer to this paper, Fusari and Gonzalez-Perez (2013) model spot variance in a two-factor log-affine specification that outperforms affine models. Amengual and Xiu (2018) estimate a two-factor non-affine model with double-sided volatility jumps, finding evidence of negative volatility jumps in a log-specification for spot variance. Since spot variance drives the continuous component of realized variance, these approaches are similar but distinct from modeling realized variance as log-affine, as realized variance reflects both the continuous and jump components of the price process. In addition, since these models are set in continuous time with latent volatility factors, rolling estimation for out-of-sample analysis is more challenging and pricing of VIX futures is less tractable because variance swaps are only available in (quasi) closed-form.[5]

2 Pricing Variance Swaps and VIX Futures

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

[5] Variance swap rates in Amengual and Xiu (2018) are computed numerically by integrating over the solutions of ordinary differential equations. Pricing VIX futures by applying Zhu and Lian (2012) requires an additional numerical integration step.

swaps of different maturities, 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 dataset of synthetic variance swap rates on a monthly grid from one-month to two-years from S&P 500 index option quotes using OptionMetrics data. The sample period is from January 4, 1996 to December 31, 2020. The estimation approach exploits the well-known no-arbitrage relationship for pricing variance swaps from option portfolios that serves as the basis for the VIX index (Carr and Wu 2009).[6] The Appendix includes a detailed description of the estimation procedure and a comparison of the estimated synthetic variance swap rates to the CBOE volatility indices and Bloomberg synthetic rates. This paper's estimated rates closely align with the alternative datasets. The advantage of this study's estimated rates is their availability over the full sample period for a wide range of maturities.

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. In practice, variance swap contracts must provide a definition for computing realized variance. For example, contracts 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, the definition should be chosen so that the floating leg payoff produces an accurate estimate of the quadratic variation of the underlying asset. The no-arbitrage replication argument that is used for pricing variance swaps relies on computing the risk-neutral expectation of an asset's quadratic variation through an application of Itô's lemma. Thus, it is desirable for the floating leg payoff to provide an accurate estimate of an asset's quadratic variation or realized variance during the life of the swap.

Based on these observations, the realized variance payoff for this paper is defined using the two-scale realized variance estimator from Zhang et al. (2005). The two-scale estimator is computed

[6]The Appendix shows that the estimated synthetic variance swap rates closely track the volatility indices from the Chicago Board Options Exchange (CBOE) and the synthetic variance swap rates from Bloomberg. The advantage of the estimated rates is their availability from 1996 to 2020 for a wide range of maturities, whereas Bloomberg data is only available from November 2008 and the CBOE indices are only available for 1-, 3-, and 6-month maturities with the VIX3M and VIX6M starting in December 2007 and January 2008. This study focuses on traditional variance swaps rather than simple variance swaps from Martin (2017) for two reasons. First, traditional variance swaps can be decomposed into variance swap forwards. For example, a one-year variance swap is the sum of a six-month variance swap and a six-month forward six-month variance swap. These properties are convenient for deriving prices in dynamic term-structure settings and they do not hold for simple variance swaps. Second, VIX futures are based on the CBOE definition for the VIX. The squared-VIX index is an estimate of a one-month traditional variance swap as described in Carr and Wu (2006) and CBOE (2019). Simple variance swaps have the advantage that their replicating portfolio is robust to jumps in the underlying asset (Martin 2017).

using high frequency data for the S&P 500 Index from Thomson Reuters Tick History (TRTH). The two-scale estimator 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 the bid-ask spread and price discreteness, which can severely bias estimation (Aït-Sahalia and Jacod 2014). The two-scale estimator balances these trade-offs by averaging realized variance estimates from a sparse sampling frequency across subsamples on a finer grid.

In practice, this paper applies the two-scale estimator by computing first stage estimates of realized variance that are equal to the sum of squared five-minute intraday log returns plus the squared overnight log return across different one-minute subsamples for each trading day.[7] The choice of a five-minute intraday sampling frequency is common in the empirical literature and is motivated by Liu et al. (2015). The second stage averages the first stage estimates across subsamples for a daily estimate of realized variance with reduced sampling variability. The daily realized variance estimates are then summed for each trading day within the month to obtain a monthly estimate of realized variance. The payoff to the floating leg of an n-month variance swap traded at time t is defined as the sum of the monthly realized variance estimates from month t to month $t+n$. The Appendix describes the steps for cleaning the high frequency data which follow Liu et al. (2015).

2.3 VIX Futures

The CBOE introduced trading in VIX futures in 2004 and VIX options in 2006. The payoff to a VIX futures contract is the difference between the futures price and a special opening quotation of the VIX index at maturity. From its definition, the VIX index upon which VIX futures are based is equal to the square root of a one-month synthetic variance swap rate. The swap rate is "synthetic" because it is computed from the price of a portfolio of S&P 500 index options following the no-arbitrage formula for pricing variance swaps. The square-root adjustment expresses the VIX index in the same units as the Black-Scholes-Merton implied volatility parameter which is familiar to option traders. By providing exposure to the VIX index which is a measure of option prices, VIX futures allow investors to hedge and speculate on shocks to implied volatility. This contrasts

[7] The realized variance estimate for the first subsample is the sum of the squared log returns from the previous close to 9:30am, 9:30am to 9:35am, ..., 3:55pm to 4:00pm. The realized variance estimate for the second subsample is sum of squared log returns from the previous close to 9:31am, 9:31am to 9:36am, ..., 3:56pm to 4:00pm, etc. Each subsample uses the same starting and ending prices to estimate the daily realized variance.

variance swaps which provide exposure to realized volatility.

2.4 Modeling Variance Swaps and VIX Futures

This paper models variance swaps as the risk-neutral ($\mathbb{Q}$-measure) expected value of realized variance 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. 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 under the physical measure $\mathbb{P}$ that follow a stationary vector autoregression,

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

with shocks v_{t+1} that are conditionally Normal $v_{t+1}|\mathcal{F}_t \overset{\mathbb{P}}{\sim} N(0, \Sigma_v)$.[8] This specification can be motivated by the intertemporal capital asset pricing model of Merton (1973) or the arbitrage pricing theory of Ross (1976).

I set the first element of the state vector to the logarithm of realized variance $\ln RV_t$,

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

This allows the model to match the dynamics of realized variance exactly and ensures that the model spans variance swap payoffs. The subsequent variables Y_t can be any financial or macroeconomic variables that help to explain the cross-sectional and time-series variation of realized variance and variance swaps.

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,

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

[8] The assumption of a first-order VAR is without loss of generality as the analysis encompasses the case in which (2) is the companion form of a higher-order VAR(p). The model selection analysis provides empirical evidence in favor of using a first-order VAR.

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

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

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 can be 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 realized variance $F_{t,0} = RV_t$. Variance swap forwards decompose the variance swap curve into one-month forward rates,

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

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 at time t, it is equivalent to the risk-neutral pricing equation,

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

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. The risk-neutral pricing equation for the one-month variance swap rate is then 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 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 then equal to the sum of variance swap forward rates as noted above,

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

The adjustment $\sqrt{12/n \cdot VS_{t,n}}$ expresses variance swap rates in annualized volatility units. Variance swap rates may be decomposed into realized variance term premia $RVTP_{t,n}$ and realized variance forecasts $RVF_{t,n}$ as follows,

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

Realized variance term premia are equal to the expected holding period return from receiving fixed in variance swaps over an n-month horizon,

$$
\begin{aligned}
RVTP_{t,n} &\equiv E_t^{\mathbb{Q}}\left[\textstyle\sum_{i=1}^{n} RV_{t+i}\right] - E_t^{\mathbb{P}}\left[\textstyle\sum_{i=1}^{n} RV_{t+i}\right] \\
&= \textstyle\sum_{i=1}^{n} e^{A_i + B_i' X_t} - \sum_{i=1}^{n} e^{A_i^P + (B_i^P)' X_t}.
\end{aligned} \tag{16}
$$

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.

The model also admits closed-form prices for VIX futures. The exponential affine price for variance swap forwards that follows from modeling the logarithm of realized variance naturally absorbs the square-root adjustment that is needed to price VIX futures.[9] 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 $B_0^F = \frac{1}{2}B_1$. Analogous to realized variance term premia, implied volatility term premia are defined as,

$$\begin{aligned} IVTP_{t,n} &\equiv E_t^{\mathbb{Q}}\left[VIX_{t+n}\right] - E_t^{\mathbb{P}}\left[VIX_{t+n}\right] \\ &= e^{A_n^F+(B_n^F)'X_t} - e^{A_n^{P,F}+(B_n^{P,F})'X_t}. \end{aligned} \tag{19}$$

Implied volatility term premia $IVTP_{t,n}$ are equal to the expected holding period return from selling the n-month VIX futures contract and thus represent the market price of risk for bearing exposure to implied volatility shocks. This differs from realized variance term premia $RVTP_{t,n}$ which represent the market price of risk for bearing exposure to realized variance shocks.

3 Model Estimation

3.1 Data and Estimation Approach

I estimate the model with realized variance and synthetic variance swap rate data from 1996 to 2020. Realized variance is estimated from high frequency data for the S&P 500 index. Synthetic variance swap rates are estimated from SPX index option quotes for {1, 2, 3, 6, 9, 12, 18, 24}-month maturities. The baseline estimation results are from non-overlapping monthly data. Realized

[9] In affine models, pricing VIX futures requires computing the risk-neutral expectation $E_t^{\mathbb{Q}}[\sqrt{A_1 + B_1'X_{t+n}}]$ by simulation or numerical integration. Quadratic models pose a similar problem. In the log term-structure model, the VIX and VIX futures price can be expressed in annualized volatility units by multiplying $\sqrt{VS_{t,1}}$ and $Fut_{t,n}$ by $100 \cdot \sqrt{12}$.

variance is the monthly sum of daily realized variance estimates from high frequency data. Synthetic variance swap rates are observed at month-end. In addition to the baseline results, the Appendix discusses estimating the model with an overlapping VAR and daily data. The paper uses the daily estimates for the empirical applications including the return predictability analysis and the relative value analysis of variance swap and VIX futures pricing.[10]

Table 1 reports summary statistics for the estimation data. The term-structure of volatility is upward sloping on average. The higher level of variance swap rates compared to realized variance reflects the unconditional variance risk premium that investors earn by bearing exposure to realized variance shocks. Further out on the curve, variance swaps are more persistent and less volatile. During periods of elevated stock market volatility, the variance swap curve tends to invert with realized variance and short-maturity variance swaps increasing more than the persistent long-maturity variance swaps. When stock market volatility declines, the curve tends to revert to its unconditional upward sloping state.

The model can be summarized by the following system of equations,

$$\begin{aligned} X_{t+1} &= \mu + \Phi X_t + v_{t+1}, & v_{t+1}|\mathcal{F}_t &\sim N(0, \Sigma_v) \\ Y_{t,n} &= g_n(X_t, \mu^{\mathbb{Q}}, \Phi^{\mathbb{Q}}, \Sigma_v) + e_{t,n}, & e_{t,n}|\mathcal{F}_t &\sim (\rho \cdot e_{t-1,n}, \sigma^2_{e,n}). \end{aligned} \tag{20}$$

The state vector X_t follows a monthly vector autoregression. Variance swap rates $Y_{t,n}$ are assumed to be observed with measurement errors $e_{t,n}$ that are mutually independent across maturities but serially correlated conditioned on the state vector with mean $\rho \cdot e_{t-1,n}$ and variance $\sigma^2_{e,n}$. The model price expressed in annualized volatility units is,

$$g_n(X_t, \mu^{\mathbb{Q}}, \Phi^{\mathbb{Q}}, \Sigma_v) = \sqrt{\frac{12}{n}\sum_{i=1}^{n} e^{A_i + B_i' X_t}}. \tag{21}$$

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

The estimation approach proceeds in three steps. First, I estimate the physical parameters $(\hat{\mu}, \hat{\Phi})$ from a monthly vector autoregression by ordinary least squares,

$$X_{t+1} = \hat{\mu} + \hat{\Phi} X_t + \hat{v}_{t+1}, \tag{22}$$

[10] The model can accommodate daily data by allowing the state vector to follow a monthly vector autoregression with overlapping observations and a horizon of $h = 21$ trading days. The parameter estimates from overlapping daily data are similar to the baseline estimates from non-overlapping monthly data.

and set $\hat{\Sigma}_v$ to the sample covariance matrix of the innovations $\hat{v}_{t+1}$. This step follows the regression-based approach in Adrian et al. (2013). Second, I estimate the risk-neutral parameters $(\hat{\mu}^{\mathbb{Q}}, \hat{\Phi}^{\mathbb{Q}})$ and $\hat{\rho}$ by minimizing the model's variance swap pricing errors by nonlinear least squares,

$$(\hat{\mu}^{\mathbb{Q}}, \hat{\Phi}^{\mathbb{Q}}) = \arg \min_{(\mu^{\mathbb{Q}}, \Phi^{\mathbb{Q}}, \rho)} \frac{1}{T \cdot N_\tau} \sum_{t=1}^{T} \sum_{n \in \tau} \left(Y_{t,n} - g_n(X_t, \mu^{\mathbb{Q}}, \Phi^{\mathbb{Q}}, \hat{\Sigma}_v) - \rho \cdot \hat{e}_{t-1,n} \right)^2 . \tag{23}$$

Third, I compute the maximum likelihood estimates $\hat{\Theta}_{MLE}$ using the OLS and nonlinear least squares estimates as an initial condition, assuming the pricing errors are conditionally Normal. The Appendix reports the log-likelihood function.[11]

3.2 Model Selection

The baseline model selected for this paper is a VAR(1) with three factors,

$$X_t = [\ln RV_t \; PC1_t \; PC2_t]. \tag{24}$$

The realized variance factor, $\ln RV_t$, is the log of monthly realized variance. The principal component factors, $PC1_t$ and $PC2_t$, are the first two PCs of log variance swap rates.[12] Figure 4 plots the state variables and the PC loadings. The top plot shows that the state variables are increasing during periods of financial distress and mean reverting. The bottom plot shows that the principal components can be interpreted as level and slope factors. High values for the slope factor indicate an inverted variance swap curve.

The model selection and specification analysis is included in the Appendix. The lag length for the VAR(p) and number of principal components is selected using standard AIC and BIC information criteria measures that trade off the improved fit from a larger model against the increasing number of model parameters. The best model according to the information criteria measures has one lag in the VAR and two PC factors. The paper uses this three-factor VAR(1) specification as the benchmark model for the empirical analysis. However, since it is challenging to precisely identify the number of factors in term-structure settings (Crump and Gospodinov 2019), the Appendix provides detailed robustness analysis showing that the estimation results and the dynamics of the volatility term

[11] Nonlinear least squares, maximum likelihood, and Bayesian MCMC estimation with latent factors deliver similar parameter estimates and results for model prices and volatility term premia. The estimation results from the alternative approaches are unreported and available upon request.

[12] I standardize log realized variance and log variance swap rates for numerical stability when estimating the model. 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 in this study for notational simplicity.

premia estimates are similar across a range of model specifications, including a five-factor model with daily and weekly realized variance factors.[13]

3.3 Estimation Results

Table 2 reports the model parameter estimates. The physical parameters in Panel A show that each of the state variables contributes significantly to the realized variance forecasts. The first row of $\hat{\Phi}$ reveals that higher levels of log realized variance $\ln RV$, the level factor $PC1$, and the slope factor $PC2$ forecast higher levels of log realized variance next month. A one unit increase in $\ln RV$, $PC1$, and $PC2$ forecasts a .22, .52, and .25 standard deviation higher value of $\ln RV$ next month.[14] The mean of the physical parameters $\hat{\mu}$ is close to zero because log realized variance and log variance swap rates have been z-scored when constructing the state vector. The second and third row of $\hat{\Phi}$ show that the level and slope factors are persistent, but that their forecasts have only a limited dependence on the other variables.

Panel B reports the price of risk estimates. The results indicate that each of the state variables contributes significantly to the time variation in volatility term premia. Since the model is nonlinear, it is difficult to interpret the price of risk estimates $\hat{\Lambda}_0$ and $\hat{\Lambda}_1$ directly. Instead, I present decompositions below to show how volatility term premia load on the different factors in the estimated model. As a preview, realized variance term premia are decreasing in $\ln RV$, increasing in $PC1$, and increasing (decreasing) in $PC2$ at the short-end (long-end) of the curve. Implied volatility term premia are increasing in $PC1$ and decreasing in $PC2$. Panel C reports the parameter estimates for the standard deviation and autocorrelation of the variance swap pricing errors and the Cholesky decomposition of the VAR residual covariance matrix. The results show that the model fits the data with small pricing errors that are significantly autocorrelated.

[13] Table A.2 reports the model specification analysis. The estimated VAR is stationary with a maximum eigenvalue of .85 (.03) and a 95% bootstrapped confidence interval from (.81, .93) that is significantly below 1. Augmented Dickey-Fuller tests reject the null hypothesis of a unit-root process for each of the state variables in favor of the alternative that the data is stationary. The VAR residuals in the benchmark three-factor model exhibit little autocorrelation. Ljung-Box tests of the null hypothesis that the residual autocorrelations are jointly equal to zero are not rejected at the 5% level for any of the state variables. In addition, since $\ln RV_t$ reflects realized variance over an entire month (thus spanning variance swap payoffs), whereas variance swap rates are observed at month-end, the Appendix estimates a five-factor model with additional RV factors including the log of daily and weekly realized variance computed during the last day and week of the month. Model prices and term premia estimates are similar in the baseline three-factor and extended five-factor specification with weekly and daily RV factors.

[14] The standard deviations of the state variables X_t are 1, 2.75 and .65 for $\ln RV$, $PC1$, and $PC2$.

3.4 Variance Swap Pricing

Table 3 reports summary statistics for the model variance swap pricing errors using month-end data from the baseline estimation results and daily data from the overlapping VAR. The mean of the errors is close to zero across maturities showing that the model is unbiased. The standard deviation of the errors is around .50% in annualized volatility units on average across maturities. In comparison, the median bid-ask spread for synthetic variance swap rates is around 1% as estimated in Table A.1 in the Appendix. The low standard deviation of the variance swap pricing errors relative to typical bid-ask spreads shows that the model provides a good fit to the data. Figure 5 illustrates the model fit by plotting variance swap rates against the model prices over time. The model provides a close fit throughout the sample period. In comparison to existing variance swap models, the baseline three-factor logarithmic model performs well, obtaining similar in-sample and good out-of-sample performance.[15]

4 Discussion

4.1 Equity Volatility Term Premia

This section investigates the relationship between the risk-return tradeoff and the volatility term premia estimates. The results include: time-series plots of the term premia estimates, return predictability regressions confirming that the term premia estimates do predict returns, regression analysis documenting how the term premia estimates are related to the different state variables through the estimated prices of risk, and a discussion of how the term premia estimates behave in response to changes in systematic risk with a comparison to the predictions from theoretical equilibrium models.

4.1.1 Time-Series Dynamics

Figure 1 plots the realized variance term premia (RVTP) estimates for one-month and twelve-month horizons in annualized variance units. The RVTP estimates tend to increase during periods of financial distress and heightened systematic risk. During the financial crisis, the twelve-month

[15] A direct comparison to existing studies is challenging because other papers use proprietary datasets and different sample periods, variance swap maturities, and estimation methods. However, reviewing existing studies for affine and quadratic models , typical in-sample RMSEs are around .40%-.60% from 1996 to 2007 and out-of-sample RMSEs are around .50%-1.50% from 2007 to 2010. For these sample periods and for {2, 3, 6, 12, 24}-month maturities, the three-factor logarithmic model obtains an in-sample RMSE of .40% and an expanding window out-of-sample RMSE of .45% on average across maturities.

RVTP reaches levels as high as 5% to 15% versus an unconditional average of 2.6%. Since the financial crisis, the term structure of RVTP has steepened with long-dated term premia increasing relative to short-dated term premia. Figure 2 plots the implied volatility term premia (IVTP) estimates for different maturities in monthly volatility units. The IVTP estimates are large in magnitude relative to typical VIX futures bid-ask spreads of .05% and switch signs between being positive and negative. During the financial crisis, IVTP reach levels as high as 1% to 3% per month, but the IVTP estimates were negative in the years leading up to the financial crisis and for different periods in recent years.

4.1.2 Variance Swap and VIX Futures Return Predictability

In the model, RVTP forecast variance swap returns and IVTP forecast VIX futures returns. The paper tests this prediction by running return predictability regressions of the form,

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

The dependent variable, $R_{t+h,n}$, for the variance swap regressions is the excess return from receiving fixed in an n-month variance swap over an h-month horizon,

$$R_{t+h,n} = VS_{t,n} - VS_{t+h,n-h} - \sum_{i=1}^{h} RV_{t+i}. \tag{26}$$

The independent variable is the expected return from the model $\hat{E}_t[R_{t+h,n}]$. For example, the one-month expected return for an n-month variance swap is equal to,

$$\begin{aligned} \hat{E}_t[R_{t+1,n}] &= \hat{E}_t[VS_{t,n} - RV_{t+1} - VS_{t+1,n-1}] \\ &= \textstyle\sum_{i=1}^{n} e^{\hat{A}_i + \hat{B}_i' X_t} - \sum_{i=0}^{n-1} e^{\hat{A}_i + \hat{B}_i'(\hat{\mu} + \hat{\Phi} X_t) + \frac{1}{2}\hat{B}_i' \hat{\Sigma}_v \hat{B}_i}. \end{aligned} \tag{27}$$

Note that the expected return for an n-month variance swap over an n-month horizon is the realized variance term premium $\hat{E}_t[R_{t+n,n}] = RVTP_{t,n}$. Thus, when $h = n$, these regressions test whether RVTP predict returns.

Table 4 reports the return predictability regressions for variance swaps using overlapping returns from daily data and Newey-West standard errors. The RVTP estimates significantly forecast returns over 1-, 3-, and 6-month horizons and exhibit an explanatory power of 18%, 27% and 34% as measured by the in-sample R^2_{adj}. The out-of-sample R^2_{oos} relative to a historical mean model is

12%, 23%, and 29% for the corresponding maturities, showing that the model continues to perform well in expanding window out-of-sample analysis from 2000 on.[16] Beyond the RVTP estimates, the table shows that the model expected return is significant at forecasting variance swap returns across forecast horizons and maturities with positive out-of-sample R^2_{oos}.

While this initial look at the model forecasting performance seems promising, several concerns might be raised regarding the robustness of the results. First, the coefficient estimate on the model expected return may be biased up (Stambaugh 1999). Second, Newey-West standard errors (SEs) may be biased down in small samples due to the heteroskedasticity of variance swap returns (Hodrick 1992). Third, the predictability may be driven by a few periods when the conditional volatility of returns is high, which OLS regressions may place a large emphasis upon (Johnson 2019). Each of these concerns could lead to the predictability of the model being overstated in Table 4.

Table 5 addresses these concerns by following the approach in Johnson (2019) and finds that the return predictability results are robust.[17] The model continues to significantly forecast returns in Hodrick-style overlapping return regressions and the null hypothesis that the model is unbiased is harder to reject after accounting for Stambaugh bias and the larger Hodrick SEs. In Table 5, daily variance swap returns are regressed onto the model expected return, $X_t = \hat{E}_t[R_{t+h,n}]$, summed h days into the past for each horizon h and maturity n. The point estimate and standard error are scaled by $Var(\sum_{s=0}^{h-1} X_{t-s})/Var(X_t)$. This approach follows Hodrick (1992) and exploits the observation that overlapping return regressions are asymptotically equivalent to non-overlapping regressions of one-period returns onto the sum of the predictor variable h periods into the past under the null hypothesis of no predictability and the assumption of stationarity. The point estimate on the model expected return is reported as Stambaugh β_{adj} after the bias adjustment and Unadjusted β without the bias adjustment. The other rows report the standard errors and p-values from the Hodrick regressions using asymptotic Newey-West and bootstrapped standard errors. The Appendix reports the corresponding weighted least squares (WLS) estimates that account for the conditional

[16] For the expanding window analysis, the model is re-estimated each day in the sample starting in 2000 using an overlapping VAR and daily data since the start of the sample in 1996. All aspects of the estimation including standardizing (z-scoring) the state variables, computing the principal components of variance swap rates, estimating the term-structure model, and estimating the return predictability regressions are repeated for each day so that the estimates match what would be available to an econometrician in real-time. The R^2_{oos} for horizon h and maturity n is defined as $R^2_{oos,h,n} = 1 - \sum_{t=1}^{T}(e^{ew,model}_{t,h,n})^2 / \sum_{t=1}^{T}(e^{ew,avg}_{t,h,n})^2$ where $e^{ew,model}_{t,h,n}$ is the expanding window forecast error from the return predictability regression using the model expected return and $e^{ew,avg}_{t,h,n}$ is the expanding window average realized return. Results are similar if $\beta_0 = 0$ and $\beta_1 = 1$ is imposed for computing $e^{ew,model}_{t,h,n}$. In this case, the average R^2_{oos} across maturities is 6%, 13% and 21% for 1-month, 3-month, and 6-month horizons versus 6%, 14%, and 22% in the table.

[17] I thank an anonymous referee for suggesting these robustness checks.

heteroskedasticity of returns to improve efficiency following Johnson (2019).

In relation to the concerns above, the point estimate on the model expected return is similar in Tables 4 and 5 despite the different estimation approaches. Adjusting for Stambaugh bias decreases the point estimates and moves the coefficients closer to the null hypothesis of an unbiased model. The asymptotic and bootstrap SEs from the Hodrick regressions range from being similar in magnitude to being as much as 2-3 times larger than the Newey-West SEs. Despite this, the model continues to significantly predict variance swap returns. For example, for a six-month forecast horizon, the average t-statistic across maturities declines from 10 using Newey-West SEs to 3.8 using the Hodrick bootstrapped SEs. In relation to the third concern about conditional heteroskedasticity, Johnson (2019) finds that the significance of the variance risk premium for predicting stock market returns is sensitive to accounting for the time-varying volatility of returns. Table A.3 in the Appendix shows that the model remains significant at forecasting variance swap returns in WLS regressions that downweight the importance of periods with high volatility such as the financial crisis and Covid-19 crisis. The WLS results show that the model's predictability for variance swap returns is not sensitive to the high volatility periods. The term premia estimates predict returns throughout the sample period in low and high volatility periods.

Tables 6 and 7 provide similar return predictability analysis for VIX futures returns. Table 6 reports regressions of VIX futures holding period returns onto IVTP estimates that are interpolated to match the futures contract maturity. The IVTP estimates significantly forecast VIX futures returns using overlapping returns from daily data and Newey-West SEs. Table 7 checks the robustness of these results by reporting Hodrick-style regressions. As before, the model continues to significantly predict returns in OLS and WLS Hodrick regressions, showing that the IVTP estimates deliver robust return forecasts. Combined with the variance swap return predictability regressions, the results show that the RVTP and IVTP estimates do forecast returns, supporting their interpretation as risk premia measures.

4.1.3 Relationship of Volatility Term Premia to the State Variables

What is the relationship between the volatility term premia estimates and the state variables? The nonlinear nature of the model makes it difficult to interpret the price of risk estimates $\hat{\Lambda}_0$ and $\hat{\Lambda}_1$ directly. To better understand how the RVTP and IVTP estimates are related to the state variables, Table 8 regresses the monthly change in volatility term premia onto the standardized (z-scored) change in the level of realized variance (RV) and the first two principal components from the level of variance swap rates ($PC1_{lvl}$ and $PC2_{lvl}$). The regressions provide a linear approximation for how realized variance, the level of variance swaps, and the slope of the variance swap curve contribute to equity volatility term premia on average. In Panel A, the linear approximation for RVTP provides high explanatory power with each of the factors contributing significantly. An increase in RV decreases RVTP with a magnitude that is largest for shorter maturities. An increase in $PC1_{lvl}$ increases RVTP in a roughly parallel manner across the curve. An increase in $PC2_{lvl}$ increases (decreases) RVTP at the short-end (long-end) of the curve. In Panel B, the linear approximation for IVTP provides moderate (high) explanatory power for shorter (longer) maturities. IVTP are increasing in $PC1_{lvl}$ with coefficients that are increasing in maturity and decreasing in $PC2_{lvl}$ with coefficients that are hump-shaped. IVTP are also increasing in RV, but with a smaller magnitude compared to that of the PC factors. In addition to these linear approximations, the Appendix computes the exact sensitivities of term premia to the model state variables at each point in time. The time-series average of the exact sensitivities is qualitatively similar to the regression results. During periods of heightened volatility, the response of term premia to the state vector is magnified due to the high values of the state variables and the nonlinear nature of the model.

4.1.4 Volatility Term Premia and the Risk-Return Tradeoff

When systematic risk increases, the model can accommodate a positive or negative response in volatility term premia depending on the estimated prices of risk and the response of the state variables. Figures 1 and 2 provide an early glimpse of this result. During periods of heightened risk like the financial crisis and Covid-19 crisis, it appears that RVTP tend to increase whereas IVTP exhibit downward spikes and subsequent increases.

Table 9 investigates this observation in more depth by regressing weekly changes in RVTP and IVTP onto weekly changes in systematic risk factors including the VIX in Panel A and weekly CRSP value-weighted stock market returns in Panel B. The results show that RVTP are significantly positively (negatively) correlated with changes in the VIX (stock market returns), consistent with the notion that RVTP increase during periods of heightened systematic risk. In contrast, the changes in IVTP for one-month and three-month maturities are significantly negatively (positively) correlated with the VIX (stock market returns), the opposite of what one might expect. At a six-month maturity the changes in IVTP exhibit an insignificant contemporaneous relationship with either of the systematic risk factors and at a twelve-month maturity the relationship changes signs to become positive.

The downward spikes in the one-month and three-month IVTP estimates when the VIX increases and muted response of six-month IVTP are surprising. Many theoretical models with a risk-averse representative agent predict that investors should earn a positive risk premium for bearing exposure to realized and implied volatility shocks that is increasing in systematic risk. In any model where the returns from receiving fixed in variance swaps and from selling VIX futures are negatively correlated with the stochastic discount factor, so that the negative correlation condition from Martin (2017) holds, the lower bound for volatility term premia is increasing in the SVIX^2, not decreasing or uncorrelated as with the IVTP estimates.[18]

To further investigate the relationship between equity volatility term premia and systematic risk, Figure 6 reports impulse response functions (IRFs) of RVTP and IVTP to a VIX shock in a bivariate VAR.[19] The IRFs for RVTP increase in response to a VIX shock and then decline over time. There is a more pronounced effect for the one-month RVTP in the first weeks after the shock and a more persistent effect for the six-month RVTP over a longer horizon. The increase in RVTP and subsequent decline is qualitatively similar to what one would expect in response to a negative shock in a consumption-based asset pricing model (Eraker 2020). To illustrate this point, the Appendix derives the analogous IRFs in the Eraker and Wu (2017) equilibrium model and finds that the analogous term premia measures for realized and implied variance increase sharply after a VIX shock and then decline.

In contrast, the IRFs for IVTP exhibit a negative or insignificant response to VIX shocks over short horizons, consistent with the contemporaneous correlations documented in Table 8. The decrease in the implied volatility risk premium is most pronounced for short-dated maturities where the term premia estimates decline sharply at first and then start to rise, only becoming positive in the period two- to three-months after the shock to the VIX. These dynamics are driven by the estimated prices of risk in the term-structure model. Based on the decomposition from the previous section, the negative short-term response of one-month IVTP indicates that the slope factor dominates at first, but then over longer horizons the more persistent level factor drives the average response and the elevated values of the implied volatility risk premium.

[18] The negative correlation between IVTP and the VIX documented in Table 8 also holds for the SVIX^2. My estimate of the VIX and SVIX^2 are 89% correlated in levels and 87% correlated in weekly changes. The correlation of $\Delta IVTP_{t,1}$ and $\Delta SVIX_t^2$ in weekly changes is -.56*** (.05).

[19] The VAR is estimated separately for RVTP and IVTP for one-month and six-month maturities using weekly data with the optimal lag length selected by SBIC criterion. The IRFs are from a Cholesky decomposition with the VIX ordered first. The confidence intervals are block bootstrapped and the variables are standardized (z-scored) for ease of interpretation.

4.2 Relative Pricing of VIX Futures

The second application of the term-structure model investigates the relative pricing of VIX futures and index options. I define the VIX futures basis as the model pricing error,

$$Basis_{t,n} = Fut_{t,n} - Fut_{t,n}^{model}. \tag{28}$$

$Fut_{t,n}$ is the daily settlement price of the n-th futures contract. $Fut_{t,n}^{model}$ is the price from the term-structure model which is interpolated from the monthly grid to match the days to maturity of the n-th VIX futures contract. The analysis of VIX futures pricing errors complements Van Tassel (2020) which provides an in-depth investigation of the no-arbitrage violations across the VIX futures and index options markets. Recall that the term-structure model is estimated with option-implied variance swaps and realized variance data. The basis reveals how futures are valued relative to index options through the lens of the term-structure model.

Table 10 reports summary statistics for the VIX futures basis for the front six contracts. The bias is around -.50% on average across contracts, indicating that VIX futures are cheap relative to the option-implied model prices. The standard deviation of around 1% for the pricing errors is about double the variability of the variance swap pricing errors and large relative to typical VIX futures bid-ask spreads of .05%. Further, the minimum and maximum values of the basis show that VIX futures exhibit substantial deviations relative to the model prices. The front two contracts have errors that are negatively skewed and fat-tailed. The errors are less negatively skewed for the longer-dated contracts but the 25th percentile is still quite low at around -1.1% across contracts. Overall, the results indicate that VIX futures prices exhibit substantial deviations relative to the model and are lower than the model prices on average. The Appendix plots the VIX futures basis for each of the front six contracts and confirms that similar summary statistics hold for a post-financial crisis period starting in 2010. The magnitude and bias of the pricing errors do not decline as the VIX futures market becomes more established, instead they are pervasive throughout the sample.

4.2.1 Comparison to Non-Parametric Estimates

This section compares the term-structure model to a non-parametric approach for estimating VIX futures pricing errors. The purpose is two-fold. First, the non-parametric estimates provide external validity for the term-structure analysis that is model-free. Second, comparing the estimated pricing errors from the term-structure model and non-parametric approach allows for a better understanding of how VIX futures, index options, and VIX options are valued relative to each other.

The non-parametric approach is outlined in a recent set of papers that show how to value VIX futures from a synthetic variance swap forward and VIX options (Hülsbusch and Kraftschik (2018), Park (2020)). The no-arbitrage relationship is,

$$Fut_{t,T} = \sqrt{Fwd_{t,T}^2 - Var_t^{\mathbb{Q}}(VIX_T)}, \tag{29}$$

where $Fwd_{t,T}$ is a one-month forward variance swap rate starting at time T when the futures contract matures and $Var_t^{\mathbb{Q}}(VIX_T)$ is the risk-neutral variance of the VIX index at time T. To derive the non-parametric estimate for the n-th futures contract, I estimate Fwd_{t,T_n} as the synthetic variance swap forward rate between the futures expiration date and thirty calendar days later $[T_n,\, T_n + 30]$. Similarly, I apply the results from Park (2020) to estimate $Var_t^{\mathbb{Q}}(VIX_{T_n})$ from a portfolio of out-of-the-money VIX options whose maturity matches the futures expiration date. The non-parametric basis is then defined as the difference between the futures price and the replicating value from the formula above,

$$Basis_{t,n}^{NP} = Fut_{t,n} - \sqrt{Fwd_{t,T_n}^2 - Var_t^{\mathbb{Q}}(VIX_{T_n})}. \tag{30}$$

Table 11 reports summary statistics for the VIX futures basis estimated from the term-structure model and the non-parametric approach from 2007 to 2020 on contract-days when VIX options are available to estimate the non-parametric basis. As before, Panel A shows that VIX futures are cheap relative to the term-structure model by around .50% on average with a standard deviation of around 1% across contracts, showing that there is no discernible impact on the average level of the basis from the sample selection requirement that VIX options be available. In contrast, Panel B shows that VIX futures are expensive relative to index options and VIX options as measured by the non-parametric approach. The bias for the non-parametric basis is positive and increases across maturities from around .30% for the front contract to 1% for the sixth contract. Since VIX futures are cheap relative to index options in Panel A, but expensive relative to index options after the vol-of-vol adjustment ($Var_t^{\mathbb{Q}}(VIX_T)$) in Panel B, the analysis indicates that VIX options are expensive relative to futures and index options. In addition, the RMSE of the non-parametric basis is found to be increasing in contract maturity from .85% for the front contract to around 1.50%-2% for the fifth and sixth contracts. This contrasts the term-structure model which tracks VIX futures prices more accurately for longer-dated contracts.

Despite the different estimation approaches, the VIX futures basis from the term-structure model and non-parametric approach exhibit similar time-series variation even after for controlling for changes in the VIX. Table 12 illustrates this result by regressing weekly changes in the basis from the term-structure model onto weekly changes in the non-parametric basis and weekly changes in the VIX. The change in the basis from the term-structure model and non-parametric approach are positively and significantly correlated for each contract and on average across contracts. These results provide external validity for the term-structure analysis by showing that the non-parametric approach delivers similar time-series behavior for the VIX futures basis.

Comparing the two approaches for estimating the VIX futures basis, the advantage of the non-parametric estimates is the model-free nature of the approach. The disadvantage is the need for VIX options data which are not available before 2006 and whose liquidity is worse for longer maturities.[20] In contrast, the estimates of the basis from the term-structure model are not model free, but they do not require VIX options data and they readily extend to longer maturities.

4.2.2 Predicting VIX Futures Returns with IVTP and the Basis

The substantial variation of VIX futures prices relative to the term-structure model raises the question of whether the model pricing errors predict returns. A natural hypothesis is that the larger index options market should provide a measure of fair value for VIX futures. Since there isn't a direct no-arbitrage relationship between index options and VIX futures, the term-structure model is useful for making the comparison without needing to rely on additional data from another market such as VIX options data. According to the model, the pricing error and the implied volatility term premium should both predict VIX futures returns. To see this, note that expected returns for VIX futures can be decomposed into a pricing error and IVTP component as,

$$Fut_{t,n} - E_t[VIX_{T_n}] = \underbrace{(Fut_{t,n} - Fut_{t,n}^{model})}_{\text{Basis}_{t,n}} + \underbrace{(Fut_{t,n}^{model} - E_t[VIX_{T_n}])}_{\text{IVTP}_{t,n}}. \tag{31}$$

An increase in either the pricing error or the IVTP predicts a higher return from selling VIX futures.

Table 13 tests whether IVTP and the basis predict VIX futures returns. Panel A regresses weekly VIX futures excess returns for the n-month contract onto the standardized IVTP and basis for the corresponding contract. Both variables are economically and statistically significant at predicting returns. A one-standard deviation increase in IVTP increases expected returns by around .30 to

[20]For VIX options, the number of out-of-the-money options used to estimate $Var_t^{\mathbb{Q}}(VIX_T)$ and their volume is decreasing in time-to-maturity while the difference in estimates of $Var_t^{\mathbb{Q}}(VIX_T)$ between ask and bid quotes is increasing in time-to-maturity.

.40 futures points for shorter-dated contracts and .20 for longer-dated contracts. A one-standard deviation increase in the basis increases expected returns by around .40 to .50 for shorter-dated contracts and by .10 to .30 for longer-dated contracts. The economic significance can be seen by comparing the coefficient magnitudes to typical bid-ask spreads of .05 for VIX futures. Panel B shows that the results are robust to including the non-parametric estimate of the basis, the VIX index, and realized variance as additional predictors. While the coefficients on the non-parametric estimate of the basis are positive as expected (similar to the model estimate), the coefficients on the non-parametric estimate are lower in magnitude and less significant than the coefficients for the estimate of the basis from the term-structure model. The return predictability analysis suggests that, in addition to tracking VIX futures prices more closely, the estimates of the basis from the term-structure model are more significant at predicting returns than the non-parametric estimate.

4.2.3 Relationship to the VIX Premium

Cheng (2018) defines the VIX premium as a statistical measure of the expected return from selling VIX futures. The low premium response puzzle refers to the anomaly that the VIX premium tends to decline when measures of systematic risk increase. The response of IVTP to a VIX shock and the cheapening of VIX futures relative to the model when the VIX increases are both consistent with the low premium response puzzle. Figure 7 illustrates this result by plotting estimates of the implied volatility risk premium from the term-structure model alongside reduced form estimates from Cheng (2018). The VIX premium and IVTP in the plot are for a one-month maturity and the futures pricing error is the one-month weighted average pricing error across the front two contracts. The plot shows that the VIX premium is 62% correlated with the IVTP estimate and 85% correlated with the IVTP estimate plus the basis or weighted average pricing error.

The significant relationship between the VIX premium and the estimates from the term-structure model is notable because it questions the mechanism driving the dynamics of the implied volatility risk premium. Cheng (2018) argues that the low premium response puzzle is related to hedging demand, as the VIX premium and dealer net positions in futures contracts from the CFTC's Commitment of Traders Report tend to fall together when risk increases. This demand channel is consistent with the decline in VIX futures pricing errors in the term-structure model when the VIX increases and with the higher correlation of the VIX premium and the IVTP estimate when it is combined with the weighted average futures pricing error. At the same time, the demand channel does not directly explain the behavior of the IVTP estimate as the term-structure model is estimated with option-implied variance swap rates and realized variance data, not VIX futures data. In addition, the IVTP estimates extend back in time before the start of VIX futures trading, when demand for VIX futures could not have been driving the negative or muted response of IVTP to increases in risk. Instead, the term-structure analysis suggests that the negative risk-return tradeoff for implied volatility may be a more general feature of volatility markets, rather than being specific to the VIX futures market, emphasizing the importance of studying the risk premia for realized variance and implied volatility separately.

5 Conclusion

This paper estimates term premia for realized variance and implied volatility in a dynamic term-structure model. By modeling the logarithm of realized variance, the paper develops a new approach for pricing variance swaps and VIX futures in closed-form. The paper constructs a detailed dataset of synthetic variance swap rates from index options and realized variance estimates from high frequency S&P 500 data to estimate the model. The model is estimated over a 25-year sample period from 1996 to 2020. The paper then considers two empirical applications of the model. In the first application, the paper finds that realized variance term premia increase when there is a shock to the VIX, but implied volatility term premia decrease or exhibit a muted response. The behavior of implied volatility term premia poses a challenge to theoretical models in which term premia for realized and implied volatility term are both increasing or decreasing in the economy's systematic risk factors. In the second application, the paper studies the relative pricing of VIX futures and option-implied variance swaps. While the model tracks the futures price quite well, it also highlights periods with large dislocations in the prices of relative claims across volatility markets. The futures price is around .50% below the option-implied model price on average and cheapens even further when systematic risk increases. Estimates of futures pricing errors from the term-structure model are highly correlated with estimates of futures pricing errors from a non-parametric, model-free approach. Compared to the non-parametric estimates, the pricing errors from the term-structure model are smaller in magnitude for longer-maturity contracts and deliver more significant forecasts of VIX futures returns. The results highlight the behavior of the model's volatility term premia estimates and illustrate how the model can be used to relate the prices of VIX futures to index options.

References

Adrian, T., R. K. Crump, and E. Moench (2013). Pricing the term structure with linear regressions. *Journal of Financial Economics 110*(1), 110–138.

Adrian, T., R. K. Crump, and E. Moench (2015). Efficient, regression-based estimation of dynamic asset pricing models. *Journal of Financial Economics 118*(2), 211–244.

Aït-Sahalia, Y. and J. Jacod (2014). *High-frequency financial econometrics*. Princeton University Press.

Aït-Sahalia, Y., M. Karaman, and L. Mancini (2020). The term structure of equity and variance risk premia. *Journal of Econometrics 219*(2), 204–230.

Amengual, D. and D. Xiu (2018). Resolution of policy uncertainty and sudden declines in volatility. *Journal of Econometrics 203*(2), 297–315.

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.

Bekaert, G. and M. Hoerova (2014). The vix, the variance premium and stock market volatility. *Journal of Econometrics 183*(2), 181–192.

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

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.

CBOE (2019). The cboe volatility index. *https://www.cboe.com/micro/vix/vixwhite.pdf*.

Cheng, I.-H. (2018). The vix premium. *The Review of Financial Studies 32*(1), 180–227.

Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics 7*(2), 174–196.

Crump, R. K. and N. Gospodinov (2019). Deconstructing the yield curve. *Working Paper*.

Dai, Q. and K. Singleton (2003). Term structure dynamics in theory and reality. *The Review of financial studies 16*(3), 631–678.

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

Drechsler, I. and A. Yaron (2011). What's vol got to do with it. *Review of Financial Studies 24*(1), 1–45.

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. (2020). Predictability puzzles. *Working Paper*.

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.

French, K. R., G. W. Schwert, and R. F. Stambaugh (1987). Expected stock returns and volatility. *Journal of Financial Economics 19*(1), 3.

Fusari, N. and T. M. Gonzalez-Perez (2013). Volatility dynamics and the term structure of the variance risk premium. *Working Paper*.

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

Gonçalves, S. and N. Meddahi (2011). Box–cox transforms for realized volatility. *Journal of Econometrics 160*(1), 129–144.

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

Hodrick, R. J. (1992). Dividend yields and expected stock returns: Alternative procedures for inference and measurement. *Review of Financial studies 5*(3), 357–386.

Hülsbusch, H. and A. Kraftschik (2018). Consistency between s&p500 and vix derivatives: Insights from model-free vix futures pricing. *Journal of Futures Markets 38*(8), 977–995.

Johnson, T. L. (2019). A fresh look at return predictability using a more efficient estimator. *The Review of Asset Pricing Studies 9*(1), 1–46.

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.

Kalnina, I. and D. Xiu (2017). Nonparametric estimation of the leverage effect: A trade-off between robustness and efficiency. *Journal of the American Statistical Association 112*(517), 384–396.

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.

Lochstoer, L. A. and T. Muir (2022). Volatility expectations and returns. *The Journal of Finance 77*, 1055–1096.

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.

Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica: Journal of the Econometric Society*, 347–370.

Park, Y.-H. (2020). Variance disparity and market frictions. *Journal of Econometrics 214*(2), 326–348.

Piazzesi, M. (2010). Affine term structure models. In *Handbook of financial econometrics: Tools and Techniques*, pp. 691–766. Elsevier.

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

Stambaugh, R. F. (1999). Predictive regressions. *Journal of Financial Economics 54*(3), 375–421.

Van Tassel, P. (2020). The law of one price in equity volatility markets. *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.

Zhu, S.-P. and G.-H. Lian (2012). An analytical formula for vix futures and its applications. *Journal of Futures Markets 32*(2), 166–190.

Realized Variance and Variance Swap Summary Statistics

Table 1: This table reports summary statistics for the realized variance and synthetic variance swap rate data that is used to estimate the term-structure model. Realized variance is estimated from high frequency S&P 500 index data. The synthetic variance swap rates are derived from SPX index option quotes. The sample is monthly from 1996 to 2020.

RV and VS summary statistics (annualized volatility units)

Maturity	RV	VS_1	VS_2	VS_3	VS_6	VS_9	VS_{12}	VS_{18}	VS_{24}
Mean	15.00	20.75	21.05	21.40	22.04	22.31	22.51	22.93	23.22
Standard Deviation	9.04	8.25	7.68	7.27	6.59	6.15	5.85	5.59	5.45
Skewness	3.57	1.66	1.58	1.52	1.35	1.25	1.15	1.00	0.92
Kurtosis	23.63	6.98	6.62	6.31	5.44	5.03	4.72	4.22	4.02
Median	12.68	18.92	19.27	20.07	20.79	21.03	21.34	21.87	22.23
Autocorrelation 1mn	0.66	0.81	0.84	0.86	0.89	0.91	0.92	0.93	0.93

Model Parameter Estimates

Table 2: This table reports maximum likelihood estimates of the model parameters for the baseline VAR(1) three-factor model using monthly data from 1996 to 2020. The standard errors in parentheses are from the asymptotic robust covariance matrix. *p<.1, **p<.05, ***p<.01.

Panel A: Physical Parameters

	μ	$\Phi_{k,1}$	$\Phi_{k,2}$	$\Phi_{k,3}$
$\ln(RV)$	0.00	0.22*	0.19***	0.38***
	(0.06)	(0.13)	(0.04)	(0.13)
$PC1$	0.03	0.07	0.88***	0.04
	(0.07)	(0.27)	(0.09)	(0.28)
$PC2$	-0.01	-0.01	-0.03	0.80***
	(0.07)	(0.08)	(0.02)	(0.07)

Panel B: Prices of Risk

	Λ_0	$\Lambda_{1,k,1}$	$\Lambda_{1,k,2}$	$\Lambda_{1,k,3}$
$\ln(RV)$	-0.65***	0.23*	-0.07*	-0.04
	(0.06)	(0.13)	(0.04)	(0.13)
$PC1$	-0.04	0.06	-0.05	0.57**
	(0.08)	(0.27)	(0.09)	(0.28)
$PC2$	-0.04	-0.03	-0.00	0.11
	(0.07)	(0.08)	(0.02)	(0.07)

Panel C: Pricing error and VAR innovation parameters

ρ	$\sigma_{e,1}$	$\sigma_{e,2}$	$\sigma_{e,3}$	$\sigma_{e,6}$	$\sigma_{e,9}$	$\sigma_{e,12}$	$\sigma_{e,18}$	$\sigma_{e,24}$
0.61	0.71	0.26	0.41	0.36	0.33	0.39	0.26	0.35
(0.07)	(0.05)	(0.02)	(0.03)	(0.03)	(0.02)	(0.04)	(0.02)	0.03
$L_{v,11}$	$L_{v,21}$	$L_{v,22}$	$L_{v,31}$	$L_{v,32}$	$L_{v,33}$			
0.62	0.75	0.85	0.17	0.15	0.30			
(0.04)	(0.08)	(0.04)	(0.03)	(0.02)	(0.02)			

Variance Swap Pricing Errors

Table 3: This table reports the mean and standard deviation of the model's variance swap pricing errors in annualized volatility units for the baseline three-factor VAR(1) model using monthly and daily data from 1996 to 2020. The daily estimates are obtained by estimating the model as a monthly overlapping VAR.

Maturity	1	2	3	6	9	12	18	24	Avg.
Monthly Data ($T = 300$)									
Mean	-0.04	-0.04	0.02	0.07	-0.02	-0.07	-0.00	0.01	-0.01
Standard Deviation	0.78	0.26	0.43	0.56	0.49	0.44	0.30	0.51	0.47
Daily Data ($T = 6,200$)									
Mean	-0.10	-0.06	0.01	0.05	-0.02	-0.05	-0.00	0.01	-0.02
Standard Deviation	0.80	0.29	0.46	0.60	0.53	0.47	0.32	0.52	0.50

Variance Swap Return Predictability

Table 4: This table reports variance swap return predictability regressions. Variance swap excess returns are regressed onto model expected returns using overlapping returns from daily data for each forecast horizon h and variance swap maturity n. The table reports Newey-West standard errors with $3 \cdot h \cdot 21$ lags. R^2_{oos} is the out-of-sample explanatory power relative to a historical mean model using expanding window estimation from 2000 on. The sample period is 1996 to 2020. *p<.1, **p<.05, ***p<.01.

$$R_{t+h,n} = \beta_0 + \beta_1 E_t[R_{t+h,n}] + \epsilon_{t+h,n}$$

Maturity (n)	1	2	3	6	9	12	18	24
One-month horizon ($h = 1$, $T = 6,179$)								
β_0	-0.00	-0.08*	-0.10	-0.15	-0.17	-0.17	-0.19	-0.20
	(0.01)	(0.05)	(0.07)	(0.10)	(0.12)	(0.13)	(0.16)	(0.18)
β_1	0.91***	1.28***	1.39***	1.51***	1.53***	1.58***	1.65***	1.74***
	(0.14)	(0.20)	(0.25)	(0.25)	(0.25)	(0.23)	(0.26)	(0.29)
R^2_{adj}	0.18	0.13	0.12	0.11	0.10	0.11	0.10	0.10
R^2_{oos}	0.12	0.08	0.07	0.05	0.05	0.04	0.04	0.04
Three-month horizon ($h = 3$, $T = 6,137$)								
β_0			-0.21	-0.25	-0.30	-0.30	-0.31	-0.33
			(0.13)	(0.21)	(0.26)	(0.30)	(0.36)	(0.44)
β_1			1.25***	1.24***	1.23***	1.23***	1.28***	1.35***
			(0.17)	(0.17)	(0.16)	(0.14)	(0.14)	(0.14)
R^2_{adj}			0.27	0.20	0.17	0.17	0.17	0.17
R^2_{oos}			0.23	0.16	0.13	0.12	0.11	0.11
Six-month horizon ($h = 6$, $T = 6,074$)								
β_0				-0.40	-0.46	-0.51	-0.50	-0.56
				(0.31)	(0.43)	(0.52)	(0.66)	(0.80)
β_1				1.20***	1.17***	1.17***	1.22***	1.27***
				(0.14)	(0.13)	(0.11)	(0.11)	(0.12)
R^2_{adj}				0.34	0.30	0.27	0.28	0.27
R^2_{oos}				0.29	0.24	0.21	0.20	0.19

Variance Swap Return Predictability Robustness

Table 5: This table investigates the robustness of the variance swap return predictability results following Johnson (2019). For each horizon h and maturity n, daily returns are regressed onto the model expected return, $X_t = E_t[R_{t+h,n}]$, summed h days into the past following the Hodrick (1992) approach for estimating overlapping return predictability regressions under the null hypothesis of no predictability. The point estimate and standard error are scaled by $Var(\sum_{s=0}^{h-1} X_{t-s})/Var(X_t)$. The table adjusts for Stambaugh bias using a simulation procedure and computes asymptotic and bootstrap SEs and p-values following Johnson (2019). The model continues to significantly predict returns. The Appendix reports the corresponding weighted least squares estimates that account for the conditional volatility of returns. The sample period is 1996 to 2020. *p<.1, **p<.05, ***p<.01.

$$\hat{\beta} = \arg\min_{\beta} \sum_{t=1}^{T} (R_{t+1} - \left(\sum_{s=0}^{h-1} X_{t-s}\right) \cdot \beta)^2$$

Maturity	1	2	3	6	9	12	18	24
One-month horizon								
Stambaugh β_{adj}	1.17	1.36	1.41	1.42	1.43	1.47	1.52	1.61
Unadjusted β	1.23	1.44	1.50	1.55	1.58	1.62	1.68	1.76
SE(Asym)	(0.25)	(0.38)	(0.42)	(0.40)	(0.39)	(0.37)	(0.38)	(0.39)
p-val(Asym %)	0.00***	0.03***	0.08***	0.04***	0.02***	0.01***	0.01***	0.00***
SE(Boot)	(0.29)	(0.35)	(0.36)	(0.36)	(0.36)	(0.37)	(0.38)	(0.39)
p-val(Boot %)	0.00***	0.00***	0.00***	0.00***	0.00***	0.00***	0.00***	0.00***
Three-month horizon								
Stambaugh β_{adj}			1.52	1.31	1.23	1.23	1.23	1.29
Unadjusted β			1.62	1.39	1.32	1.33	1.35	1.42
SE(Asym)			(0.43)	(0.35)	(0.32)	(0.31)	(0.31)	(0.32)
p-val(Asym %)			0.04***	0.02***	0.01***	0.01***	0.01***	0.01***
SE(Boot)			(0.39)	(0.34)	(0.34)	(0.34)	(0.36)	(0.37)
p-val(Boot %)			0.00***	0.00***	0.00***	0.00***	0.30***	0.20***
Six-month horizon								
Stambaugh β_{adj}				1.63	1.35	1.33	1.28	1.34
Unadjusted β				1.68	1.39	1.37	1.33	1.39
SE(Asym)				(0.34)	(0.29)	(0.28)	(0.30)	(0.32)
p-val(Asym %)				0.00***	0.00***	0.00***	0.00***	0.00***
SE(Boot)				(0.39)	(0.35)	(0.34)	(0.36)	(0.37)
p-val(Boot %)				0.00***	0.00***	0.00***	0.10***	0.00***

VIX Futures Return Predictability

Table 6: This table reports regressions of VIX futures holding period returns onto the implied volatility term premia estimates from 2007 to 2020. VIX_{T_n} is the special opening quotation (SOQ) used to calculate the expiration day value of VIX futures contracts. R^2_{oos} is the out-of-sample explanatory power relative to a historical mean model using expanding window estimation from 2007 on. The table reports Newey-West standard errors.

Return Predictability: $Fut_{t,n} - VIX_{T_n} = \beta_0 + \beta_1 IVTP_{t,n} + \epsilon_{t,n}$

Contract (n)	1	2	3	4	5	6
β_0	-0.03	-0.12	-0.37	-0.51	-0.68	-0.84
	(0.47)	(0.91)	(1.09)	(1.30)	(1.44)	(1.66)
β_1	1.68***	1.22***	1.17***	1.21***	1.21***	1.28***
	(0.62)	(0.31)	(0.21)	(0.21)	(0.19)	(0.18)
R^2_{adj}	0.08	0.08	0.11	0.15	0.16	0.19
R^2_{oos}	0.01	-0.02	0.02	0.01	0.11	0.19
N	3502	3502	3502	3502	3502	3502

VIX Futures Return Predictability Robustness

Table 7: This table investigates the robustness of the VIX futures return predictability regressions. The approach follows Johnson (2019) and reports OLS and WLS Hodrick-style return predictability regressions for an n-month horizon for contract n. The predictor variable is the IVTP for each contract. The sample period is 2007 to 2020. *p<.1, **p<.05, ***p<.01.

Predicting VIX Futures Returns with IVTP

Contract (n)	1	2	3	4	5	6
Panel A: OLS estimate for IVTP						
Stambaugh β_{adj}	1.62	1.28	1.05	1.01	0.98	0.93
Unadjusted β	1.69	1.42	1.07	0.99	0.94	0.88
SE(Asym)	(0.73)	(0.49)	(0.35)	(0.30)	(0.28)	(0.27)
p-val(Asym %)	2.66**	0.92***	0.31***	0.07***	0.05***	0.07***
SE(Boot)	(0.45)	(0.35)	(0.29)	(0.27)	(0.26)	(0.26)
p-val(Boot %)	0.00***	0.00***	0.00***	0.00***	0.00***	0.00***
Panel B: WLS estimate for IVTP						
Stambaugh β_{adj}	1.22	0.96	0.83	0.82	0.83	0.81
Unadjusted β	1.29	1.09	0.85	0.81	0.78	0.76
SE(Asym)	(0.45)	(0.29)	(0.24)	(0.22)	(0.23)	(0.23)
p-val(Asym %)	0.72***	0.10***	0.04***	0.02***	0.03***	0.05***
SE(Boot)	(0.41)	(0.32)	(0.27)	(0.25)	(0.25)	(0.25)
p-val(Boot %)	0.10***	0.10***	0.20***	0.00***	0.00***	0.00***
Panel C: WLS first stage conditional variance: $R^2_{t+1} = \gamma \cdot VIX_t + e_{t+1}$						
γ	14.08***	7.72***	4.88***	3.41***	2.56***	1.95***
	(1.77)	(1.09)	(0.80)	(0.55)	(0.35)	(0.19)
R^2_{adj}	0.05	0.05	0.05	0.06	0.07	0.09

Variance Swap Rate and VIX Futures Factor Loadings

Table 8: Panel A regresses monthly changes in realized variance term premia onto z-scored changes in realized variance and the first two principal components of variance swap rates. Panel B reports the analogous results for VIX futures. The loadings from the regressions are qualitatively similar to the time-series average of the exact sensitivities of term premia to the state variables reported in the Appendix.

Variance Swap and VIX Futures Decompositions from 1996-2020

Maturity in months	1	2	3	6	9	12	18	24
Panel A: Realized Variance Term Premia Decomposition: $\Delta RVTP_{t,n} = \beta' \Delta f_t + \epsilon_{t,n}$								
ΔRV	-0.80***	-0.43***	-0.29***	-0.14***	-0.09***	-0.06***	-0.04**	-0.02*
	(0.07)	(0.04)	(0.03)	(0.03)	(0.02)	(0.02)	(0.02)	(0.01)
$\Delta PC1_{lvl}$	1.67***	1.48***	1.42***	1.37***	1.37***	1.36***	1.32***	1.26***
	(0.03)	(0.04)	(0.05)	(0.06)	(0.05)	(0.04)	(0.03)	(0.02)
$\Delta PC2_{lvl}$	1.05***	0.38***	0.03	-0.40***	-0.55***	-0.61***	-0.64***	-0.63***
	(0.02)	(0.03)	(0.03)	(0.02)	(0.02)	(0.02)	(0.03)	(0.03)
R^2_{adj}	0.98	0.99	0.98	0.96	0.95	0.95	0.96	0.97
Panel B: Implied Volatility Term Premia Decomposition: $\Delta IVTP_{t,n} = \beta' \Delta f_t + \epsilon_{t,n}$								
ΔRV	0.19***	0.16**	0.16**	0.14**	0.11**	0.09**	0.07**	0.06**
	(0.06)	(0.07)	(0.08)	(0.06)	(0.05)	(0.04)	(0.03)	(0.03)
$\Delta PC1_{lvl}$	0.26***	0.60***	0.87***	1.41***	1.66***	1.74***	1.69***	1.53***
	(0.10)	(0.12)	(0.12)	(0.08)	(0.06)	(0.07)	(0.08)	(0.09)
$\Delta PC2_{lvl}$	-0.96***	-1.26***	-1.39***	-1.40***	-1.28***	-1.17***	-1.00***	-0.86***
	(0.06)	(0.09)	(0.11)	(0.12)	(0.10)	(0.09)	(0.08)	(0.07)
R^2_{adj}	0.67	0.62	0.62	0.72	0.83	0.89	0.91	0.91

RVTP and IVTP versus the VIX and Stock Market Returns

Table 9: This table reports the correlation of weekly changes in RVTP and IVTP with the VIX and CRSP value-weighted stock market returns. Panels A and B regress z-scored changes in the term premia estimate onto z-scored changes in the risk factors. The table reports Newey-West standard errors. The sample period is 1996 to 2020 using overlapping daily data.

Correlation of Weekly Changes: $\Delta Y_t = \beta \Delta X_t + e_t$

Term Premia (ΔY_t)	$\Delta RVTP_{t,n}$				$\Delta IVTP_{t,n}$			
Maturity	1	3	6	12	1	3	6	12
Panel A: Weekly Change in the VIX Index (ΔX_t)								
ΔVIX_t	0.85***	0.79***	0.65***	0.55***	-0.71***	-0.48***	-0.13	0.33***
	(0.08)	(0.08)	(0.08)	(0.08)	(0.05)	(0.08)	(0.08)	(0.07)
R^2_{adj}	0.73	0.62	0.42	0.30	0.50	0.23	0.02	0.11
Panel B: Weekly CRSP Value-Weighted Stock Market Return (ΔX_t)								
$RMRF_t$	-0.66***	-0.65***	-0.58***	-0.52***	0.49***	0.28***	-0.02	-0.39***
	(0.07)	(0.07)	(0.07)	(0.06)	(0.05)	(0.06)	(0.06)	(0.05)
R^2_{adj}	0.44	0.42	0.33	0.27	0.24	0.08	0.00	0.15

VIX Futures Basis Summary Statistics

Table 10: This table reports summary statistics for the VIX futures basis or model pricing errors for the front six futures contracts from 2007 to 2020 ($T = 3502,\ N = 6$).

$\text{Basis}_{t,n} = Fut_{t,n} - Fut_{t,n}^{model}$

Contract	1	2	3	4	5	6	Avg.
Mean	-0.34	-0.33	-0.49	-0.61	-0.62	-0.57	-0.49
Standard Deviation	0.98	1.12	1.12	1.07	1.02	0.97	1.05
Skewness	-3.85	-0.25	0.12	0.65	0.51	-0.07	-0.48
Kurtosis	44.25	10.07	5.08	4.77	4.22	4.73	12.19
Minimum	-15.43	-8.58	-6.40	-5.78	-5.70	-7.71	-8.27
25th-Percentile	-0.66	-0.97	-1.24	-1.38	-1.35	-1.27	-1.15
Median	-0.25	-0.42	-0.69	-0.82	-0.80	-0.63	-0.60
75th-Percentile	0.13	0.34	0.33	0.21	0.12	0.12	0.21
Maximum	4.60	7.72	5.54	4.98	4.56	3.04	5.07
RMSE	1.04	1.17	1.22	1.23	1.19	1.13	1.16
t-statistic	-5.31	-1.94	-2.33	-2.82	-3.15	-3.30	-3.14

Term-Structure Model versus Non-Parametric VIX Futures Basis Estimates

Table 11: This table compares estimates of the VIX futures basis from the term-structure model to a non-parametric approach that estimates VIX futures prices from synthetic variance swap rates and VIX options. The sample period is from 2007 to 2020 on days when VIX options are available for the non-parametric estimate. The last column averages across the contracts.

VIX Futures Basis Estimates

Contract	1	2	3	4	5	6	Avg.
Panel A: Term-Structure Model: $\text{Fut}_{t,n} - \text{Fut}_{t,n}^{model}$							
Mean	-0.37	-0.33	-0.50	-0.63	-0.69	-0.66	-0.53
Standard Deviation	0.98	1.13	1.11	1.07	0.98	0.90	1.03
RMSE	1.05	1.17	1.22	1.24	1.20	1.12	1.17
Panel B: Non-Parametric: $\text{Fut}_{t,n} - \sqrt{\text{Fwd}_{t,n} - Var_t^{\mathbb{Q}}(VIX_{T_n})}$							
Mean	0.28	0.41	0.68	0.57	0.78	0.94	0.61
Standard Deviation	0.80	0.89	1.11	1.08	1.24	1.75	1.15
RMSE	0.85	0.98	1.30	1.22	1.46	1.99	1.30
Sample Size	2538	3443	3468	3407	3296	3196	3224

Time-Series Relationships with the VIX Futures Basis

Table 12: This table reports regressions of the weekly change in the VIX futures basis onto the weekly change in the VIX index and the non-parametric estimate of the basis from 2007 to 2020 by contract. The last column uses the weekly change in the average basis across the front six contracts. The variables are z-scored in each regression. The table reports Newey-West standard errors.

$$\Delta \text{Basis}_t = \beta \cdot \Delta X_t + \epsilon_t$$

Contract	1	2	3	4	5	6	Avg.
ΔVIX	-0.39***	-0.40***	-0.34***	-0.32***	-0.20**	-0.11	-0.40***
	(0.05)	(0.07)	(0.06)	(0.09)	(0.10)	(0.12)	(0.06)
ΔNon-Parametric Basis	0.38**	0.31***	0.20***	0.25***	0.33***	0.26***	0.60***
	(0.15)	(0.04)	(0.07)	(0.05)	(0.06)	(0.09)	(0.09)
R^2_{adj}	0.32	0.20	0.13	0.12	0.13	0.09	0.41
T	2534	3439	3464	3403	3292	3192	3496

Predicting VIX Futures Returns with IVTP and the Basis

Table 13: This table reports VIX futures return predictability regressions. Panel A regresses weekly VIX futures returns onto the IVTP and VIX futures basis estimates from the term-structure model. Panel B adds the non-parametric basis, VIX, and RV as additional predictors. The predictor variables are z-scored for ease of interpretation. The table reports Newey-West standard errors to account for the overlapping weekly returns. The variation in sample size reflects the availability of longer-dated VIX futures and VIX options contracts to estimate the non-parametric basis. The sample period is from 2007 to 2020.

$$Fut_{t,n} - Fut_{t+h,n} = \beta \cdot X_t + e_{t+h,n}$$

Panel A: Predicting returns with IVTP and Basis

Contract (n)	1	2	3	4	5	6
Constant	0.19**	0.14*	0.06	0.03	0.03	0.03
	(0.09)	(0.07)	(0.06)	(0.05)	(0.05)	(0.04)
IVTP	0.35**	0.41***	0.34***	0.28***	0.20***	0.15***
	(0.15)	(0.15)	(0.12)	(0.10)	(0.07)	(0.05)
Basis	0.44***	0.51***	0.30***	0.17**	0.10**	0.07
	(0.16)	(0.13)	(0.11)	(0.07)	(0.04)	(0.04)
R^2_{adj}	0.05	0.09	0.06	0.04	0.02	0.02
N	3497	3497	3497	3497	3497	3497

Panel B: Robustness to additional predictor variables

Contract (n)	1	2	3	4	5	6
Constant	0.22*	0.14*	0.05	0.03	0.03	0.05
	(0.12)	(0.07)	(0.05)	(0.05)	(0.04)	(0.04)
IVTP	0.32	0.44**	0.43***	0.40***	0.35***	0.27***
	(0.21)	(0.18)	(0.15)	(0.15)	(0.12)	(0.09)
Basis	0.56**	0.48***	0.28***	0.20**	0.14**	0.13***
	(0.22)	(0.15)	(0.10)	(0.09)	(0.06)	(0.04)
Non-Parametric Basis	0.26	0.12	0.02	0.09	0.07	0.02
	(0.19)	(0.09)	(0.07)	(0.06)	(0.05)	(0.04)
VIX	0.12	-0.26	-0.20	-0.25	-0.24	-0.16
	(0.36)	(0.30)	(0.33)	(0.26)	(0.21)	(0.16)
RV	0.04	0.08	-0.19	-0.18	-0.15	-0.18
	(0.45)	(0.37)	(0.29)	(0.19)	(0.16)	(0.13)
R^2_{adj}	0.07	0.09	0.09	0.10	0.09	0.08
N	2533	3438	3463	3402	3291	3191

Realized Variance Term Premia

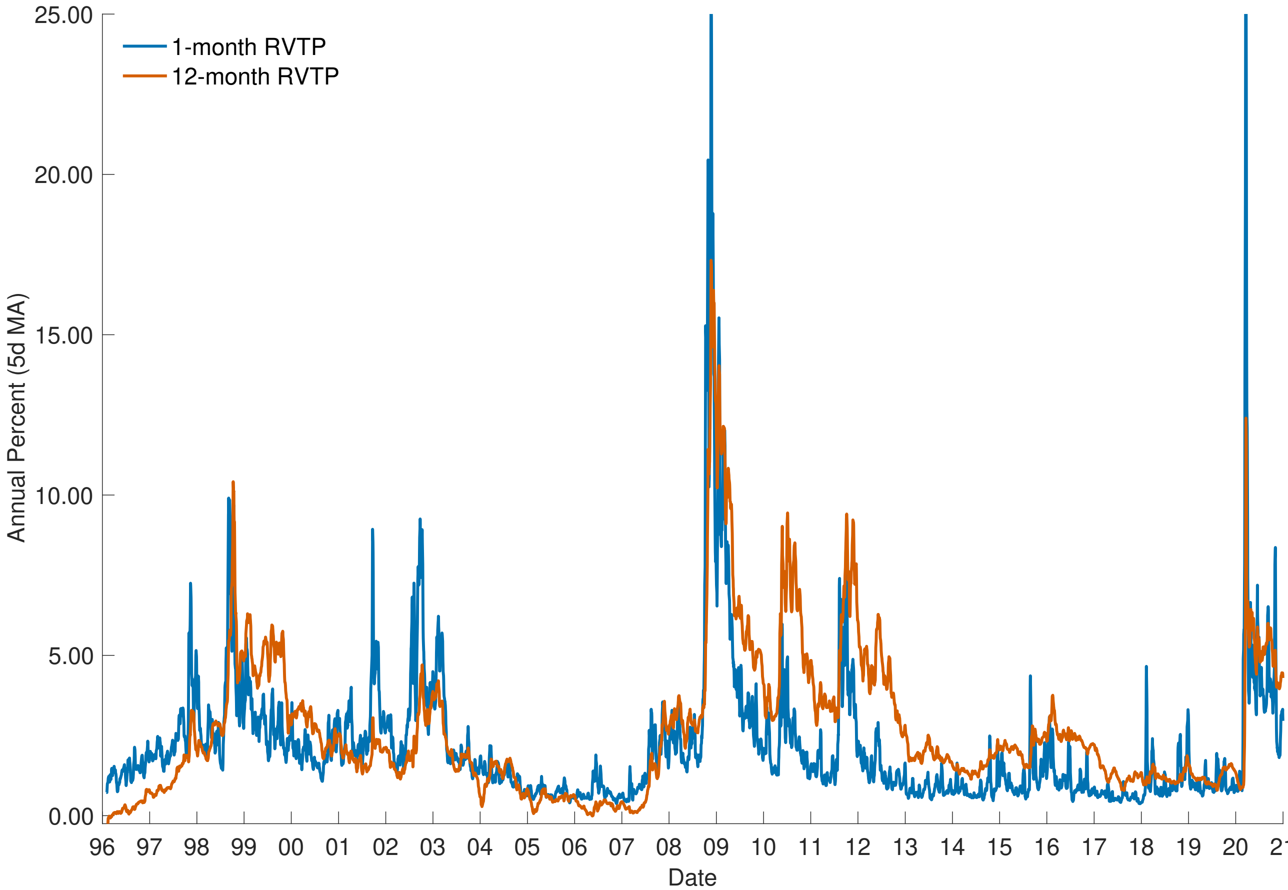

Figure 1: This figure plots the realized variance term premia (RVTP) estimates for a one-month and twelve-month horizons in annualized variance units. RVTP represent the expected holding period returns from receiving fixed in variance swaps of different maturities. High levels of RVTP predict high returns from selling volatility by receiving fixed in variance swaps. The RVTP estimates can be interpreted as the cost of insuring against realized variance shocks over different horizons.

Implied Volatility Term Premia

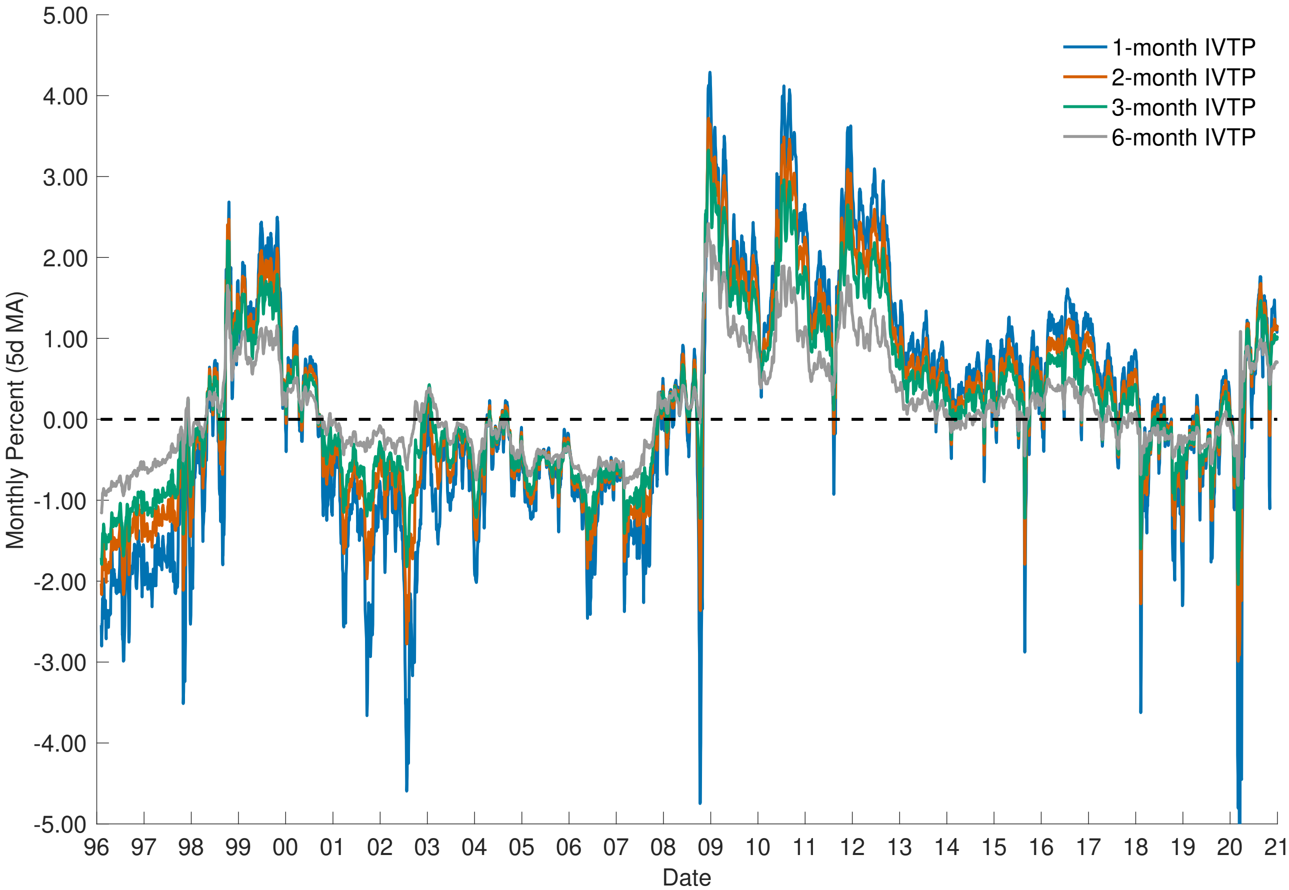

Figure 2: This figure plots the implied volatility term premia (IVTP) estimates over different horizons in monthly volatility units. IVTP represent the holding period return from selling VIX futures contracts for different maturities. High levels of IVTP predict high returns from selling VIX futures. The IVTP estimates can be interpreted as the cost of insuring against implied volatility shocks over different horizons.

Relative Pricing of VIX Futures versus Option-Implied Model Price

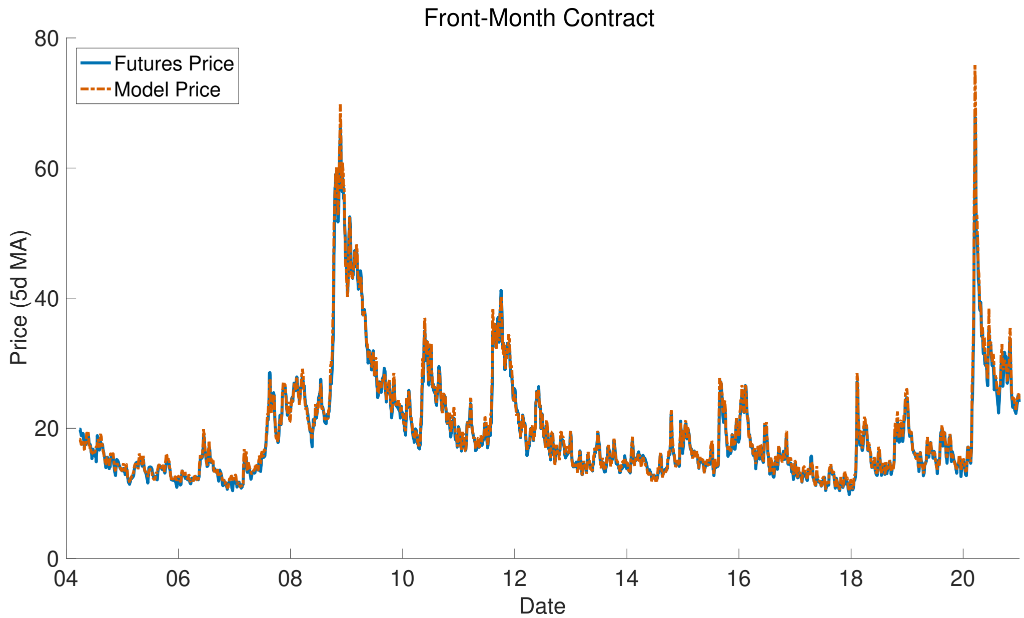

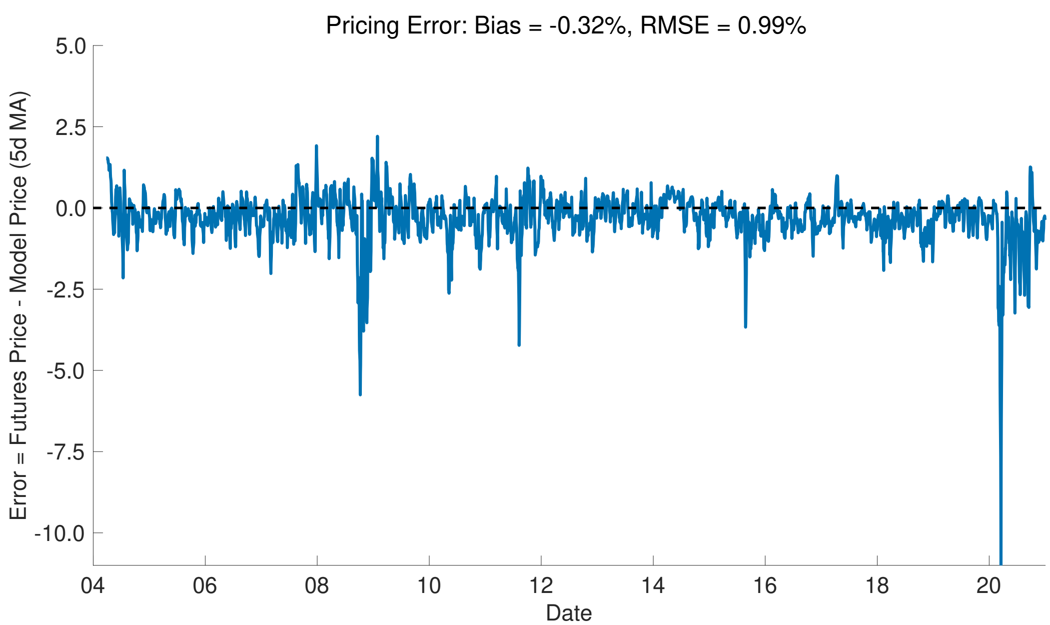

Figure 3: The top figure plots the futures price for the front-month contract against the estimated model price. The bottom figure plots the difference between the futures price and the model price. Negative values in the bottom plot indicate that VIX futures are lower than the model price. Large pricing errors reveal dislocations between the prices of index options and VIX futures as measured by the term-structure model. The sample period is March 2004 to December 2020.

State Variables in Three-Factor Logarithmic Model

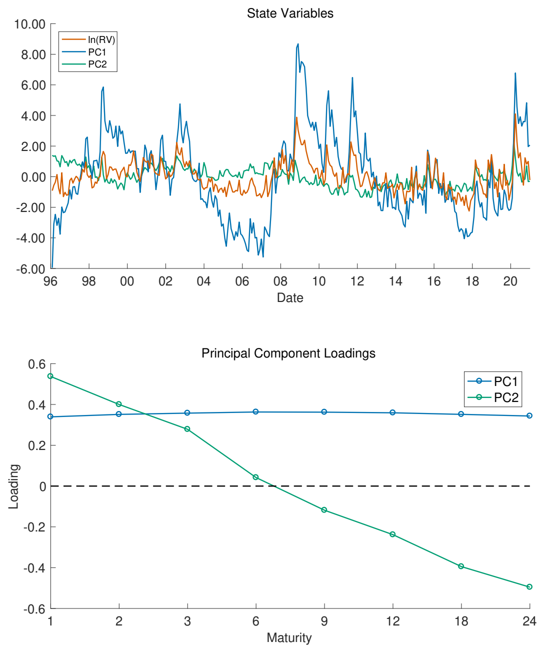

Figure 4: This figure plots the state variables, $X_t = [\ln RV_t \, PC1_t \, PC2_t]$, and the loadings for the principal component factors. $\ln RV_t$ is the logarithm of realized variance. $PC1_t$ and $PC2_t$ are the first and second principal components from log variance swap rates which can be interpreted as level and slope factors.

Model Fit for Variance Swap Rates

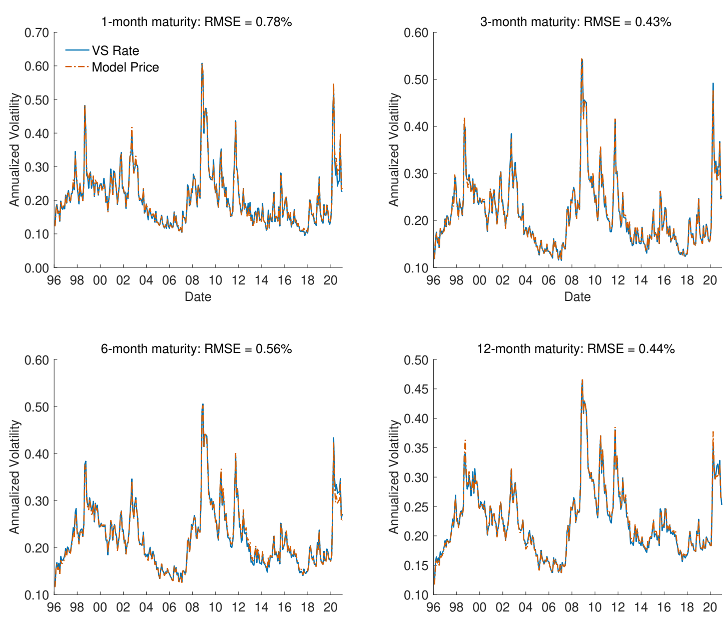

Figure 5: This figure plots the variance swap rate data against the estimated model prices for one-, three-, six-, and twelve-month maturities. A three-factor logarithmic model with two principal components fits the cross section of variance swap rates with small pricing errors as measured by the root-mean-squared-errors (RMSEs) in the titles of the subplots.

Impulse Response Functions of RVTP and IVTP to a VIX Shock

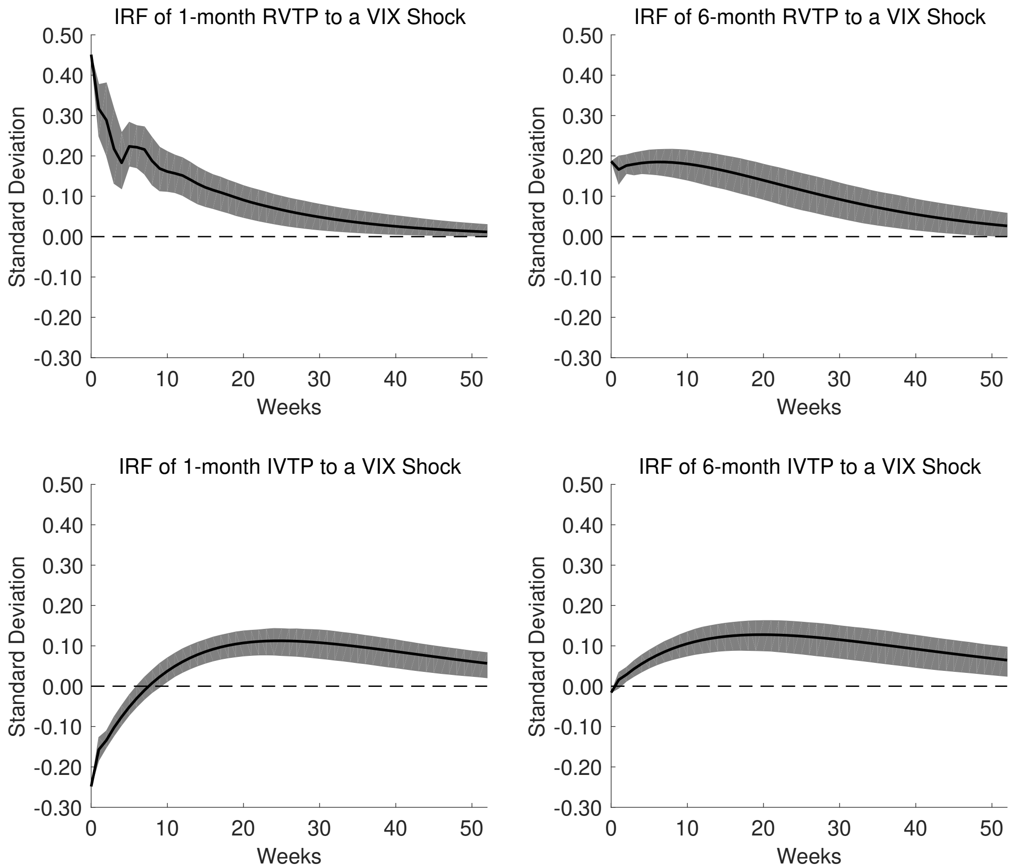

Figure 6: This figure plots the impulse response functions (IRFs) of the one-month and six-month RVTP and IVTP estimates to a one-standard deviation VIX shock in a bivariate VAR using weekly data. The IRFs are from a Cholesky decomposition with the VIX ordered first. The sample period is 1996 to 2020.

IVTP and VIX Futures Basis versus the VIX Premium

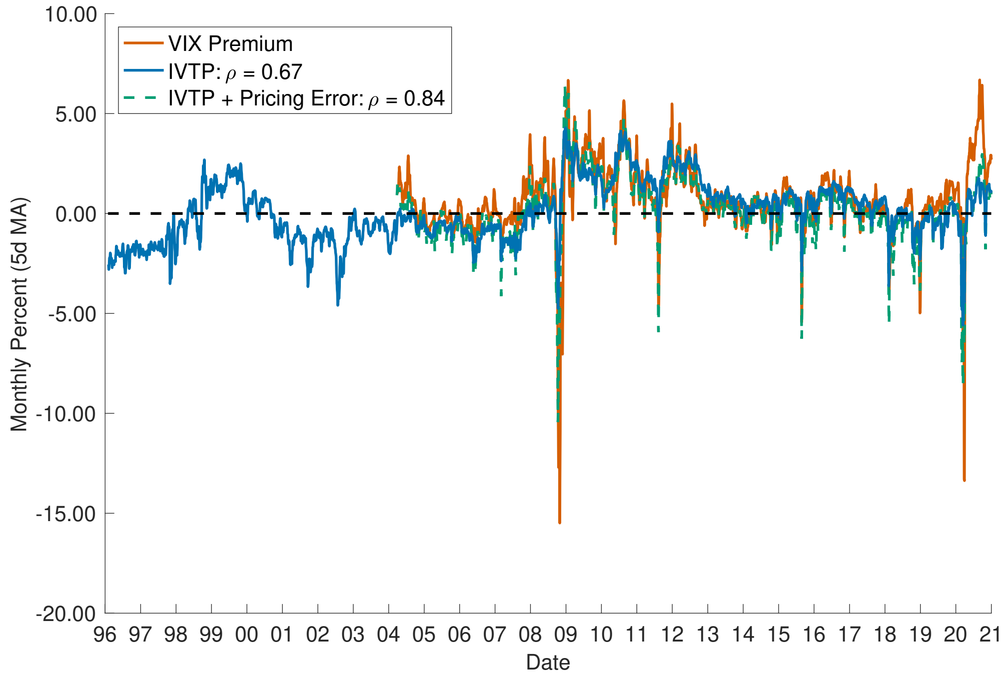

Figure 7: This figure plots the VIX premium from Cheng (2018) against the one-month $IVTP$ estimate from the term-structure model and the $IVTP$ estimate plus the model pricing error. The model estimates are 67% and 84% correlated with the VIX premium from March 2004 to December 2020. The $IVTP$ estimate from the model extends back before the start of VIX futures trading.

Appendix for "Equity Volatility Term Premia"

Peter Van Tassel

A Appendix

A.1 Synthetic Variance Swap Rates

I compute synthetic variance swap rates from the price of a replicating portfolio of options following Carr and Wu (2009). I perform this computation every day for expirations between ten days and three years to maturity with at least three quotes for out-of-the-money (OTM) puts and calls with non-zero bids.[21] For each maturity and date, I fit a flexible implied volatility function by local linear regression to OTM options and compute synthetic variance swap rates as the weighted average of fitted option prices.[22] To obtain a constant maturity term-structure, 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 indices. When necessary, I use nearest neighbor extrapolation to estimate the one-month and two-year synthetic rates.

Figure A.1 provides an example of the estimation procedure on December 20, 2019. The top plot reports the fitted implied volatility functions against log-moneyness for different maturities. The close fit indicates that the implied volatility functions provide an accurate estimate of the risk-neutral distribution. The bottom plot reports the estimated synthetic variance swap rates against the CBOE volatility indices and Bloomberg synthetic variance swap rates.[23] The estimated rates closely align with the CBOE indices. The Bloomberg rates are somewhat lower but within typical bid-ask spreads.

The empirical choice of a two-year term-structure balances the earlier years in the sample when some extrapolation is required versus the latter years in the sample when longer maturities are available so no extrapolation is needed. Figure A.2 illustrates this point by plotting the maturities of the estimated synthetic rates each day. Starting in December 2005, an option with a three-year maturity is introduced at the end of each calendar year that rolls down to a two-year maturity

[21] I determine which options are OTM from the forward rate implied by put-call parity. 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 function. I estimate the forward rate as the median implied forward rate from as many as ten strike prices that are closest to the strike price that minimizes the absolute difference between call and put prices.

[22] Following the CBOE construction of the VIX index, I include options with strike prices that are successively OTM until either all options are used or there are two consecutive strike prices with zero bids. As a result, the strike price range varies as volatility changes over time and across maturities. To extrapolate beyond the strike price range, I use the implied volatility of the put and call that are furthest OTM to append log-Normal tails. I compute synthetic variance swap rates for "standard" 3rd Friday expiration dates and begin to include SPX "weekly" expirations in 2014 on Fridays other than the 3rd Friday of the month following the VIX index methodology.

[23] In addition to the VIX which tracks one-month (30-day) implied volatility of SPX index options, the CBOE also tracks three-month and six-month implied volatility with the VIX3M (formerly VXV) and VIX6M (formerly VXMT) indices.

throughout the year. As a result, there is almost no extrapolation at the long-end of the curve from 2006 to 2020. Prior to 2006 the maximum maturity with an estimated synthetic variance swap rate is typically between 18-months and two-years.

Table A.1 compares the estimated synthetic variance swap rates to CBOE and Bloomberg data for external validation. Panel A shows that the average mean and median difference between the estimated synthetic rates and CBOE indices is .17% and .11% across maturities with an average standard deviation of .41%. Panel B shows that similar results hold when the estimated rates are compared to Bloomberg data. Panel C reports estimates of bid-ask spreads for the synthetic rates from bid and ask quotes for SPX options using the CBOE VIX construction methodology. The mean differences in Panels A and B are within the average bid-ask spreads across maturities. Finally, in terms of comovement, the table shows that the estimated rates are highly correlated with the CBOE and Bloomberg data in monthly, weekly, and daily changes.

Figure A.3 illustrates the similarity in the time-series dynamics results by plotting the one-month synthetic variance swap rate versus the CBOE Volatility index (VIX) at month-end dates along with a scatter plot of daily changes. Similar to the summary statistics, the daily changes are highly correlated and the RMSE is low. Overall, the results indicate that the estimated rates closely track the alternative data sources. The estimated rates are used in this paper because they are available over the full sample period for a wide range of maturities.

A.2 Two-Scale Realized Variance Estimation

I estimate realized variance for the S&P 500 index following the two-scale approach from Zhang et al. (2005). In the first step, I use a sparse five-minute sampling frequency to compute realized variance estimates from one-minute data for five different subsamples. The two-scale realized variance estimate for each trading day, or second stage estimate, is the average of the first stage estimates across subsamples to reduce sampling variability. I apply standard data cleaning techniques for high frequency data when implementing the two-scale estimator. I use one-minute intraday prices during regular market hours from 9:30am to 4:00pm from Thomson Reuters Tick History (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 Model Estimation

A.3.1 Specification Analysis

Table A.2 reports the model specification analysis. Panel A provides information criteria for selecting the VAR(p) lag length and the number of principal component factors (K_{PC}) in the state vector. Since the level and slope factors capture over 99% of the variation in log variance swap rates, this study only considers using up to two PC factors to avoid overfitting. The objective is to choose a model that minimizes the information criteria which include the Schwarz (SBIC), and Hannan and Quinn (HQIC), and Akaike (AIC) measures. The best model according to the IC measures is bolded and has one lag $p = 1$ and two principal component factors $K_{PC} = 2$. Panel B reports the sample autocorrelations of the VAR residuals in the selected three-factor VAR(1) model out to six lags. The autocorrelations are close to zero and insignificant at the 5% level in univariate tests. Ljung-Box tests of the null hypothesis that the autocorrelations are jointly equal to zero are not rejected at the 5% level for any of the variables. Figure A.4 illustrates the autocorrelation tests. The left panels plot the residual autocorrelations alongside 95% pointwise confidence intervals. The right panel plots the time-series of the VAR residuals. The low and primarily insignificant values of the residual autocorrelations provide evidence against modeling higher order MA processes in the VAR. Reverting back to Table A.2, Panel C reports augmented Dickey-Fuller tests. The null hypothesis of a unit-root process is rejected at the 1% level for each of the variables in favor of the alternative that the data is generated by a stationary process. The maximum eigenvalue of $\hat{\Phi}$ is .85 (.03) with a 95% bootstrapped confidence interval from (.81, .93) which does not reject the hypothesis that the estimated VAR is stationary.

[24]On most days I observe a price every minute so no interpolation is required. The mean (median) number of observations per day is 389.7 (391) out of $6.5 \times 60 + 1 = 391$ possible one-minute observations from 9:30am to 4:00pm. Out of 6220 trading days there are only 242 days with fewer than 391 observations and most of these days occur during the early years in the sample. On these days, the mean (median) number of observations is 357.2 (388.5).

A.3.2 Likelihood Function

To compute the MLE, I assume the pricing errors are conditionally Normal. The likelihood function from the forecast error decomposition is,

$$\begin{aligned} f(Y_t, X_t|X_{t-1}, \Theta) &= f(Y_t|X_t, \Theta) f(X_t|X_{t-1}, \Theta) \\ &= f(Y_t|X_t, \mu^{\mathbb{Q}}, \Phi^{\mathbb{Q}}, L_v, \rho, \sigma_e) f(X_t|X_{t-1}, \mu, \Phi, L_v). \end{aligned} \tag{32}$$

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

$$\begin{aligned} LL &= \textstyle\sum_{t=1}^{T} \ln f(Y_t|X_t, \Theta) + \sum_{t=2}^{T} \ln f(X_t|X_{t-1}, \Theta) \\ &= -\tfrac{T \cdot N_\tau}{2} \ln(2\pi\sigma_e^2) - \tfrac{1}{2} \textstyle\sum_{t=1}^{T} \sum_{n\in\tau} ((Y_{t,n} - g_n(X_t, \mu^{\mathbb{Q}}, \Phi^{\mathbb{Q}}, \Sigma_v) - \rho \cdot e_{t-1,n})/\sigma_{e,n})^2 \\ &\quad -\tfrac{T \cdot K}{2} \ln(2\pi) - \tfrac{T}{2} \ln|\Sigma_v| - \tfrac{1}{2} \textstyle\sum_{t=1}^{T} (X_t - \mu - \Phi X_{t-1})' \Sigma_v^{-1} (X_t - \mu - \Phi X_{t-1}). \end{aligned} \tag{33}$$

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 simplifies maximum likelihood estimation as the likelihood function only needs to be maximized over the remaining parameters $(\mu^{\mathbb{Q}}, \Phi^{\mathbb{Q}}, L_v, \rho, \sigma_e)$.

A.3.3 RVTP and IVTP Sensitivities to State Variables

In addition to the linear approximations from the regressions in Table 8, the exact sensitivities or partial derivatives of volatility term premia with respect to the state variables are also available. For example, the response of RVTP to changes in the state variables is,

$$\nabla RVTP_{t,n} = \frac{12}{n} \left(\sum_{i=1}^{n} B_i \cdot e^{A_i + B_i' X_t} - \sum_{i=1}^{n} B_i^P \cdot e^{A_i^P + (B_i^P)' X_t} \right). \tag{34}$$

Figure A.5 reports sensitivities of RVTP and IVTP with respect to changes in the state vector over time. The top left subplots report the time-series average of the partial derivatives. The top right subplots show the sensitivities on October 31, 2007 when the state vector is close to its unconditional mean. The bottom right subplots show the sensitivities during the financial crisis on November 30, 2008, a period with high realized and implied volatility and an inverted variance swap curve. The magnitude of these sensitivities is much larger for RVTP than average, which is illustrated by the gray box highlighting the scale of the other plots. The bottom right subplots report the partial derivatives on December 31, 2016, a low volatility period with an upward sloping variance swap curve. The analysis highlights how the sensitivity of RVTP and IVTP with respect to the state vector can change over time due to the nonlinear nature of the model. The average partial derivatives and regression coefficients in Table 8 provide a sense for the typical relationships.

A.3.4 Five-Factor Model with Daily and Weekly RV Factors

The state vector in the baseline three-factor model includes $\ln RV_t$ which is estimated over an entire month and two principal component factors from log variance swap rates that are observed at the end of the month. The asynchronous observation of $\ln RV_t$ and the PC factors may lead to concerns that the model's volatility forecasts could be improved by using higher frequency realized variance factors, thus impacting the term premia estimates and empirical results. In particular, the heterogeneous autoregressive (HAR) model of Corsi (2009) and the extensions in Bekaert and Hoerova (2014) show that volatility forecasts can be improved by using realized variance estimates over different horizons.

Embellishing the model setup, I follow Corsi (2009) by adding daily $\ln RVD_t$ and weekly $\ln RVW_t$ log realized variance factors to the state vector and estimate the five-factor model,

$$X_t = [\ln RV_t \ \ln RVD_t \ \ln RVW_t \ PC1_t \ PC2_t], \tag{35}$$

with $K_{PC} = 2$ principal component factors. Figure A.6 compares the estimates from the five-factor VAR(1) model to the baseline three-factor VAR(1) model and a three-factor VAR(2) specification. The model prices and volatility term premia estimates are economically very similar across the different model specifications, showing that the results are robust to the concerns regarding asynchronicity and the lag length of the VAR.

A.3.5 Time-Varying Volatility-of-Volatility

The covariance matrix of the VAR(1) residuals for X_t is constant in the baseline model. While this is a standard assumption in many term-structure models, one may wonder how allowing for a time-varying covariance matrix would impact the estimation results.

Before extending the analysis, it is important to recall that the state vector already models second moments directly. In this sense, the baseline model is already a nonlinear stochastic volatility model that features time-varying volatility-of-volatility. In the baseline model, vol-of-vol is increasing in the $\ln RV_t$ and $PC1_t$ factors. In addition, since the model prices realized variance exactly, the model already matches the realized vol-of-vol for the stock market.

Another observation before considering an extension is that the log transformation attenuates much of the excess kurtosis and vol-of-vol that is present in the level of realized variance. From 1996 to 2020, the kurtosis of realized variance is 87 in levels versus 4 in logs. There is also limited vol-of-vol in the VAR residuals for the log model. In unreported results that are available upon request, I estimate OLS regressions of the squared VAR(1) residuals onto the lagged state variables for log and linear specifications of the state vector. The regressions test for evidence of time-varying vol-of-vol in the residuals and are motivated by the usual affine setup for stochastic volatility models (Piazzesi (2010)). In the log specification, the coefficients are insignificant and the explanatory power is low. In the linear or affine specification, the coefficients are significant and the explanatory power is higher. The results suggest that modeling a time-varying covariance matrix for the VAR residuals is more of a concern when realized variance is modeled in levels, not in logs as in this paper.

Nonetheless, one can extend the baseline model and examine the robustness of the estimates. To that end, I consider the following modification,

$$\begin{aligned}
X_{t+1} &= \mu + \Phi X_t + \left(L_v S_t^{1/2}\right) v_{t+1},\ \epsilon_{t+1}|\mathcal{F}_t \sim N(0, I) \\
X_{t+1} &= \mu^{\mathbb{Q}} + \Phi^{\mathbb{Q}} X_t + \left(L_v S_t^{1/2}\right) v_{t+1}^{\mathbb{Q}},\ \epsilon_{t+1}^{\mathbb{Q}}|\mathcal{F}_t \sim N(0, I) \\
S_t &= \begin{pmatrix} \alpha_1 + \beta_1' X_t & & \\ & \ddots & \\ & & \alpha_K + \beta_K' X_t \end{pmatrix}.
\end{aligned} \tag{36}$$

The derivation of model prices proceeds as before with an adjustment for the time-varying covariance matrix. Variance swap forwards and VIX futures prices continue to be exponential affine in the state vector. The recursive pricing equations for the extended model are,

$$\begin{aligned}
A_n &= A_{n-1} + B_{n-1}'\mu^{\mathbb{Q}} + \tfrac{1}{2} B_{n-1}' L_v \Sigma_\alpha L_v' B_{n-1} \\
B_n' &= B_{n-1}'\Phi^{\mathbb{Q}} + \tfrac{1}{2}\tilde{\beta}_{n-1}',
\end{aligned} \tag{37}$$

where,

$$\tilde{\beta}_{n-1}' \equiv \sum_{i=1}^{K} W_{ii,n-1}\beta_i', \quad W_{n-1} \equiv L_v' B_{n-1} B_{n-1}' L_v, \tag{38}$$

and $S_t \equiv \Sigma_\alpha + \Sigma_{\beta,X_t}$ with $\Sigma_\alpha = Diag(\alpha)$ and $\Sigma_{ii,\beta,X_t} = \beta_i' X_t$. I estimate the model by maximum likelihood and impose the constraint that the time-varying covariance matrix is positive definite at each point in time. Figure A.7 reports the estimated variance swap rates and volatility term premia in the baseline model versus the extension with a time-varying VAR(1) covariance matrix. The model variance swap rates closely match the data and the volatility term premia estimates are broadly similar. Overall, the plots show that the empirical results are robust to allowing for a time-varying VAR covariance matrix.

A.3.6 Estimation with Overlapping Daily Data

Adjusting the notation slightly, the model can accommodate daily data by allowing the state vector to follow a monthly vector autoregression with overlapping observations and a horizon of $h = 21$ trading days,

$$\begin{aligned} X_{t+h} &= \mu + \Phi X_t + v_{t+h}, & v_{t+h}|\mathcal{F}_t \sim N(0, \Sigma_v) \\ Y_{t,n} &= g_n(X_t, \mu^{\mathbb{Q}}, \Phi^{\mathbb{Q}}, \Sigma_v) + e_{t,n}, & e_{t,n}|\mathcal{F}_t \sim (\rho \cdot e_{t-h,n}, \sigma_{e,n}^2). \end{aligned} \tag{39}$$

The nonlinear least squares estimation proceeds as before after adjusting the standard errors in the first-step for the overlapping nature of the VAR. Estimating the model with daily data increases the sample size which improves precision and is more demanding because it requires fitting the more extreme observations of variance swap rates that occur within month.[25] In addition, estimating the model with daily data allows for a study of the relative pricing of variance swaps and VIX futures and for the estimation of the VIX futures convexity adjustment at a daily frequency.

A.4 Equity Volatility Term Premia in an Equilibrium Model

This section derives RVTP and IVTP in the consumption-based asset pricing model from Eraker and Wu (2017). The stochastic discount factor (SDF) is,

$$M_t = E_t x_T^{-\gamma} e^{-r_f t} = e^{\alpha(u_{-\gamma}, t, T) + \beta'(u_{-\gamma}, t, T) X_t - r_f t}. \tag{40}$$

In the two-factor volatility model and $T \to \infty$ limit it follows from Ito's lemma that,

$$\frac{dM_t}{M_t^-} = -r_f dt - \gamma \sigma_t dB_t^x - \eta \sigma_v \sigma_t dB_t^v - \phi \sigma_\theta \sqrt{\theta_t} dB_t^\theta + (e^{-\eta \xi_t} - 1) dN_t - l_0(\rho(-\eta) - 1) dt. \tag{41}$$

[25] 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.

The price of variance futures is derived in Appendix A.4 as $F_{t,t+\tau} = Var_t^Q(\ln P_{t+\tau}) = a(\tau) + b(\tau)\sigma_t^2 + c(\tau)\theta_t$ from the conditional cumulant generating function for $\ln P_{t+\tau}$. The instantaneous risk premium for variance futures is thus,

$$\begin{aligned} -Cov_t(\tfrac{dM_t}{M_t^-}, dF_{t,t+\tau}) &= -Cov_t\left(\tfrac{dM_t}{M_t^-}, a(\tau) + b(\tau)d\sigma_t^2 + c(\tau)d\theta_t\right) \\ &= \left(b(\tau)l_0\mu_\xi(1 - \tfrac{1}{(1+\eta\mu_\xi)^2}) + b(\tau)\eta\sigma_v^2\sigma_t^2 + c(\tau)\sigma_\theta^2\phi\theta_t\right)dt. \end{aligned} \tag{42}$$

The risk premium has three factors that are all negative. The second and third-term are time-varying and increasing in risk as measured by either spot volatility σ_t or the long-run volatility factor θ_t.

In addition to the instantaneous risk premium, it is also possible to compute the realized variance and implied variance term premia for any horizon τ. The stock price dynamics under $\mathbb{P}$ are,

$$d\ln P_t = \left(\frac{\partial\alpha(u_{1-\gamma}, t, T)}{\partial t} - \frac{\partial\alpha(u_{-\gamma}, t, T)}{\partial t}\right)dt + r_f dt + d\ln x_t + \lambda_\sigma d\sigma_t^2 + \lambda_\theta d\theta_t. \tag{43}$$

It follows that realized variance is an affine function of σ_t^2 and θ_t,

$$\begin{aligned} d\ln P_t^2 &= (\sigma_t dB_t^x + \lambda_\sigma(\sigma_v\sigma_t dB_t^v + \xi_t dN_t) + \lambda_\theta\sigma_\theta\sqrt{\theta_t}dB_t^\theta)^2 \\ &= \left((1+\lambda_\sigma^2\sigma_v^2)\sigma_t^2 + \lambda_\theta^2\sigma_\theta^2\theta_t\right)dt + \lambda_\sigma^2\xi_t^2 dN_t. \end{aligned} \tag{44}$$

The expected realized variance from time t to time $t+\tau$ is then,

$$E_t\left[\int_t^{t+\tau} d\ln P_s^2\right] = 2\lambda_\sigma^2 l_0\mu_\xi^2\tau + (1+\lambda_\sigma^2\sigma_v^2)\int_t^{t+\tau} E_t[\sigma_s^2]ds + \lambda_\theta^2\sigma_\theta^2\int_t^{t+\tau} E_t[\theta_s]ds. \tag{45}$$

We can compute this expectation directly. Define $Y_t = (\sigma_t^2,\ \theta_t)$. For this analysis note that we can drop $\ln x_t$ as a state variable because $(dB_t^x, dB_t^v, dB_t^\theta, dN_t)$ are independent and because $\ln x_t$ does not feed back into X_t The dynamics for Y_t are thus,

$$dY_t = \left(\begin{pmatrix} 0 \\ k_\theta\theta \end{pmatrix} + \begin{pmatrix} -\kappa & \kappa \\ 0 & -\kappa_\theta \end{pmatrix}\begin{pmatrix} \sigma_t^2 \\ \theta_t \end{pmatrix}\right)dt + \begin{pmatrix} \sigma_v\sigma_t & 0 \\ 0 & \sigma_\theta\sqrt{\theta_t} \end{pmatrix}\begin{pmatrix} dB_t^v \\ dB_t^\theta \end{pmatrix} + \begin{pmatrix} \xi_t dN_t \\ 0 \end{pmatrix}. \tag{46}$$

We can rewrite dY_t as,

$$dY_t = K(\mu - Y_t)dt + \sigma(Y_t)dB_t + dJ_t, \tag{47}$$

where $E[Y_t] = \mu$ and dB_t and dJ_t are martingale components. The conditional expectation of Y_t is then,

$$E_t[Y_s] = f(t, Y_t) = \mu + e^{-K(s-t)}(Y_t - \mu), \tag{48}$$

where $e^{-K(s-t)} \equiv \sum_{n=0}^{\infty} \frac{1}{n!}(-K(s-t))^n$ is the matrix exponential. Rewriting Y_t in this way,

$$\begin{aligned} dY_t &\equiv \left(K_0^{Y,P} + K_1^{Y,P} Y_t\right) dt + \sigma(Y_t) dB_t + dJ_t \\ &= (-K_1^{Y,P})((-K_1^{Y,P})^{-1} K_0^{Y,P} - Y_t)dt + \sigma(Y_t)dB_t + dJt \\ &\equiv K^P(\mu^P - Y_t)dt + \sigma(Y_t)dB_t + dJ_t. \end{aligned} \tag{49}$$

where,

$$K_0^{Y,P} = \begin{pmatrix} l_0 \mu_\xi \\ \kappa_\theta \theta \end{pmatrix}, \; K_1^{Y,P} = \begin{pmatrix} -\kappa & \kappa \\ 0 & -\kappa_\theta \end{pmatrix}. \tag{50}$$

The realized variance forecast is then,

$$\begin{aligned} RVF_{t,t+\tau} &\equiv E_t\left[\int_t^{t+\tau} d\ln P_s^2\right] \\ &= a\tau + \int_t^{t+\tau} (b\; c) \cdot E_t[Y_s] ds \\ &= a\tau + (b\; c) \cdot \int_t^{t+\tau} \left(\mu^P + e^{-K^P(s-t)}(Y_t - \mu^P)\right) ds \\ &= a\tau + (b\; c) \cdot \left(\mu^P \tau + (K^P)^{-1}(I - e^{-K^P \tau})(Y_t - \mu^P)\right) \end{aligned} \tag{51}$$

where $(a, b, c) = (2\lambda_\sigma^2 l_0 \mu_\xi^2, 1 + \lambda_\sigma^2 \sigma_v^2, \lambda_\theta^2 \sigma_\theta^2)$. Variance swap rates follow from the same computation under the risk-neutral measure with $a^Q = 2\lambda_\sigma^2 l_0^Q (\mu_\xi^Q)^2$,

$$VS_{t,t+\tau} \equiv E_t^Q\left[\int_t^{t+\tau} d\ln P_s^2\right] = a^Q \tau + (b\; c) \cdot \left(\mu^Q \tau + (K^Q)^{-1}(I - e^{-K^Q \tau})(Y_t - \mu^Q)\right), \tag{52}$$

To compute this expectation we need the risk-neutral dynamics which are,

$$\begin{aligned} dY_t &= \left(K_0^{Y,Q} + K_1^{Y,Q} Y_t\right) dt + \sigma(Y_t) dB_t^Q + dJ_t^Q \\ &= (-K_1^{Y,Q})((-K_1^{Y,Q})^{-1} K_0^{Y,Q} - Y_t)dt + \sigma(Y_t)dB_t^Q + dJ_t^Q \\ &\equiv K^{Y,Q}(\mu^Q - Y_t)dt + \sigma(Y_t)dB_t^Q + dJ_t^Q, \end{aligned} \tag{53}$$

where,

$$K_0^{Y,Q} = \begin{pmatrix} l_0^Q \mu_\xi^Q \\ \kappa_\theta \theta \end{pmatrix}, \; K_1^{Y,Q} = \begin{pmatrix} -\kappa & \kappa \\ 0 & -\kappa_\theta \end{pmatrix} \tag{54}$$

and $\eta = -\beta_2(u_{-\gamma})$ and $\phi = -\beta_3(\mu_{-\gamma})$ are the prices of risk from the SDF. The realized variance term premium $(Q - P)$ is thus, $RVTP_{t,t+\tau} = VS_{t,t+\tau} - RVF_{t,t+\tau}$. Similarly, the implied variance term premium is, $IVTP_{t,T_1,T_2} = VS_{t,T_2} - VS_{t,T_1} - E_t^P[VS_{T_1,T_2}]$, which can be derived from the observation that,

$$E_t^P[VS_{T_1,T_2}] = a^Q\tau + (b\,c)\cdot\left(\mu^Q\tau + (K^Q)^{-1}(I - e^{-K^Q\tau})(\mu^P + e^{-K(T_1-t)}(Y_t - \mu^P) - \mu^Q)\right). \tag{55}$$

where $\tau = T_2 - T_1$. This risk premium measures the expected return for bearing exposure to implied volatility risk. The $a^Q\tau$ term stemming from jump risk cancels out in the $IVRP_{t,T_1,T_2}$ measure which reduces the $IVRP$ risk premium relative to the forward variance risk premium $RVTP_{t,T_2} - RVTP_{t,T_1}$.

Estimated Synthetic Variance Swaps versus CBOE and Bloomberg Data

Table A.1: This table compares the estimated synthetic variance swap rates to CBOE and Bloomberg data. Panels A and B report the mean, standard deviation, and median of the difference between the estimated synthetic variance swap rates and the alternative data sources in annualized volatility units along with the correlation between the estimated rate and the alternative rate in monthly, weekly, and daily changes. Panel C reports estimates of synthetic variance swap bid-ask spreads from bid and ask quotes for SPX options using the CBOE VIX construction methodology. The table shows that the estimated synthetic rates used in the paper are highly correlated with the alternative data sources and within typical bid-ask spreads of the alternative rates on average. The advantage of the estimated synthetic rates is their availability across maturities and during the full sample period from 1996 to 2020. In contrast, the Bloomberg data is only available starting in November 2008 and the CBOE VIX3M and VIX6M indices are only available starting in December 2007 and January 2008 respectively.

Panel A: Synthetic rates vs. CBOE volatility indices

Maturity	1	3	6	Avg.
Mean Difference	0.34	0.13	0.05	0.17
Standard Deviation	0.57	0.37	0.30	0.41
Median Difference	0.27	0.07	-0.00	0.11
Correlation of monthly changes	0.99	1.00	1.00	1.00
Correlation of weekly changes	0.98	0.99	0.99	0.99
Correlation of daily changes	0.96	0.96	0.97	0.96

Panel B: Synthetic rates vs. Bloomberg synthetic rates

Maturity	1	2	3	6	9	12	18	24	Avg.
Mean Difference	0.73	0.38	0.29	0.17	0.19	0.16	0.15	0.08	0.27
Standard Deviation	1.01	0.60	0.54	0.46	0.51	0.49	0.45	0.43	0.56
Median Difference	0.48	0.26	0.18	0.08	0.10	0.08	0.10	0.05	0.17
Correlation of monthly changes	0.99	0.99	0.99	0.99	0.99	0.99	0.98	0.98	0.99
Correlation of weekly changes	0.97	0.98	0.98	0.98	0.97	0.96	0.94	0.94	0.97
Correlation of daily changes	0.91	0.94	0.93	0.93	0.92	0.88	0.84	0.80	0.89

Panel C: Synthetic variance swap rate bid-ask spread estimates

Maturity	1	2	3	6	9	12	18	24	Avg.
Mean	1.41	1.28	1.19	1.05	1.03	1.15	1.34	1.31	1.22
Median	1.29	1.14	1.05	0.90	0.85	0.93	0.97	0.94	1.01

Model Specification Analysis

Table A.2: This table provides model specification analysis. Panel A reports information criteria for selecting the VAR(p) lag length and the number of principal component factors (K_{PC}) in the state vector. The objective is to choose a model that minimizes the information criteria which include the Schwarz (SBIC), and Hannan and Quinn (HQIC), and Akaike (AIC) measures. The best model according to the IC measures is bolded and has one lag $p = 1$ and two principal component factors $K_{PC} = 2$. Panel B reports the sample autocorrelations for the VAR residuals in the selected three-factor VAR(1) model out to six lags. The autocorrelations are close to zero and insignificant at the 5% level in univariate tests. Ljung-Box tests of the null hypothesis that the autocorrelations are jointly equal to zero are not rejected at the 5% level for any of the state variables. Panel C reports augmented Dickey-Fuller tests. The null hypothesis of a unit-root process is rejected for each of the variables in favor of the alternative that the data is generated by a stationary process.

Panel A: Model Lag Length Selection

Model	p	K_{PC}	SBIC	HQIC	AIC
1	1	1	-8.60	-8.62	-8.64
2	1	2	**-10.20**	**-10.25**	**-10.28**
3	2	1	-8.75	-8.81	-8.85
4	2	2	-9.87	-10.01	-10.09
5	3	1	-8.66	-8.80	-8.88
6	3	2	-9.59	-9.88	-10.04
7	4	1	-8.49	-8.72	-8.85
8	4	2	-9.34	-9.84	-10.12

Panel B: VAR(1) residuals

Variable	$\hat{v}_{\ln RV}$	$\hat{v}_{PC1}$	$\hat{v}_{PC2}$
Panel B.I: Autocorrelation			
ρ_1	0.11	0.01	-0.11*
ρ_2	0.06	-0.06	-0.02
ρ_3	0.08	-0.04	0.02
ρ_4	-0.09*	-0.01	0.04
ρ_5	0.08	0.05	0.10*
ρ_6	-0.01	-0.06	0.07
Panel B.II: Ljung-Box tests			
Variable	$\hat{v}_{\ln RV}$	$\hat{v}_{PC1}$	$\hat{v}_{PC2}$
Q_{LB}	11.35	3.34	9.06
p-val	7.8*	76.6	17.0

Panel C: ADF tests

Variable	$X_{\ln RV}$	X_{PC1}	X_{PC2}
H_0 : Unit root			
t-statistic	-6.82	-4.02	-5.88
p-value	0.00***	0.13***	0.00***
H_0 : Unit root with trend and lag			
t-statistic	-5.85	-3.75	-5.52
p-value	0.00***	1.91**	0.00**

Variance Swap Return Predictability Robustness: WLS Estimates

Table A.3: This table reports in-sample weighted-least-squares (WLS) estimates of variance swap return predictability regressions following Johnson (2019). Panel A reports the first stage estimate of conditional variance σ_t^2 from a regression of squared daily returns for each maturity onto realized variance over the past month (a constant is omitted so the conditional variance estimates are positive; the coefficient and SE are multiplied by 1e4 for readability). Panel B reports the WLS estimates for the return predictability regressions. The model continues to significantly predict returns in the WLS regressions.

Panel A: First Stage Conditional Variance Regressions: $R_{t+1}^2 = \gamma \cdot RV_t + e_{t+1}$

Maturity (n)	1	2	3	6	9	12	18	24
γ	12.38***	26.69***	38.78***	69.22***	100.73***	133.66***	156.74***	194.89***
	(3.19)	(6.70)	(9.51)	(15.15)	(24.27)	(43.07)	(31.79)	(41.25)
R^2_{adj}	0.07	0.08	0.08	0.11	0.10	0.09	0.13	0.11

Panel B: Second Stage WLS Regressions: $\hat{\beta} = \arg\min_\beta \sum_{t=1}^{T} \left(\left(R_{t+1} - \left(\sum_{s=0}^{h-1} X_{t-s} \right) \cdot \beta \right) / \hat{\sigma}_t \right)^2$

Maturity (n)	1	2	3	6	9	12	18	24
One-month horizon								
Stambaugh β_{adj}	1.03	1.08	1.04	0.97	0.95	0.99	0.94	0.99
Unadjusted β	1.09	1.17	1.14	1.10	1.09	1.14	1.10	1.14
SE(Asym)	(0.12)	(0.16)	(0.17)	(0.19)	(0.21)	(0.22)	(0.25)	(0.27)
p-val(Asym %)	0.00***	0.00***	0.00***	0.00***	0.00***	0.00***	0.01***	0.02***
SE(Boot)	(0.15)	(0.20)	(0.21)	(0.21)	(0.21)	(0.22)	(0.23)	(0.23)
p-val(Boot %)	0.00***	0.00***	0.00***	0.00***	0.00***	0.00***	0.00***	0.00***
Three-month horizon								
Stambaugh β_{adj}			1.10	0.92	0.89	0.93	0.87	0.92
Unadjusted β			1.20	1.01	0.98	1.03	0.98	1.04
SE(Asym)			(0.19)	(0.18)	(0.20)	(0.21)	(0.24)	(0.26)
p-val(Asym %)			0.00***	0.00***	0.00***	0.00***	0.03***	0.05***
SE(Boot)			(0.24)	(0.22)	(0.22)	(0.22)	(0.23)	(0.24)
p-val(Boot %)			0.00***	0.00***	0.00***	0.00***	0.00***	0.00***
Six-month horizon								
Stambaugh β_{adj}				1.09	0.94	0.97	0.90	0.97
Unadjusted β				1.14	0.98	1.02	0.95	1.01
SE(Asym)				(0.22)	(0.21)	(0.22)	(0.25)	(0.28)
p-val(Asym %)				0.00***	0.00***	0.00***	0.03***	0.05***
SE(Boot)				(0.26)	(0.23)	(0.23)	(0.24)	(0.25)
p-val(Boot %)				0.00***	0.00***	0.00***	0.00***	0.00***

VIX Futures Basis Summary Statistics: Full Sample and Post-Crisis Sample

Table A.4: This table reports summary statistics for the model's VIX futures pricing errors for the front three contracts from 3/26/04 and for the front six contracts from 1/4/10.

$\text{Basis}_{t,n} = Fut_{t,n} - Fut_{t,n}^{model}$

Sample Period	2004 to 2020				2010 to 2020						
Contract	1	2	3	Avg.	1	2	3	4	5	6	Avg.
Mean	-0.32	-0.21	-0.28	-0.27	-0.36	-0.53	-0.75	-0.88	-0.88	-0.80	-0.70
Standard Deviation	0.94	1.08	1.15	1.06	0.93	0.95	0.97	0.96	0.90	0.86	0.93
Skewness	-3.79	-0.42	-0.10	-1.43	-4.92	0.81	0.87	1.39	1.07	0.00	-0.13
Kurtosis	45.65	10.01	4.26	19.97	63.07	12.37	6.59	8.81	7.65	7.22	17.62
Minimum	-15.43	-8.58	-6.40	-10.14	-15.43	-7.26	-5.42	-5.78	-5.70	-7.71	-7.88
25th-Percentile	-0.64	-0.86	-1.14	-0.88	-0.65	-1.07	-1.33	-1.48	-1.44	-1.38	-1.22
Median	-0.24	-0.23	-0.34	-0.27	-0.27	-0.62	-0.90	-1.07	-1.02	-0.88	-0.79
75th-Percentile	0.13	0.43	0.56	0.37	0.09	-0.06	-0.24	-0.36	-0.37	-0.27	-0.20
Maximum	4.60	7.72	5.54	5.95	4.60	7.72	5.54	4.98	4.56	3.04	5.07
RMSE	0.99	1.10	1.18	1.09	0.99	1.09	1.22	1.30	1.25	1.17	1.17
t-statistic	-5.87	-1.35	-1.33	-2.85	-4.71	-3.33	-3.72	-4.34	-4.95	-5.16	-4.37

Synthetic Variance Swap Rate Estimation

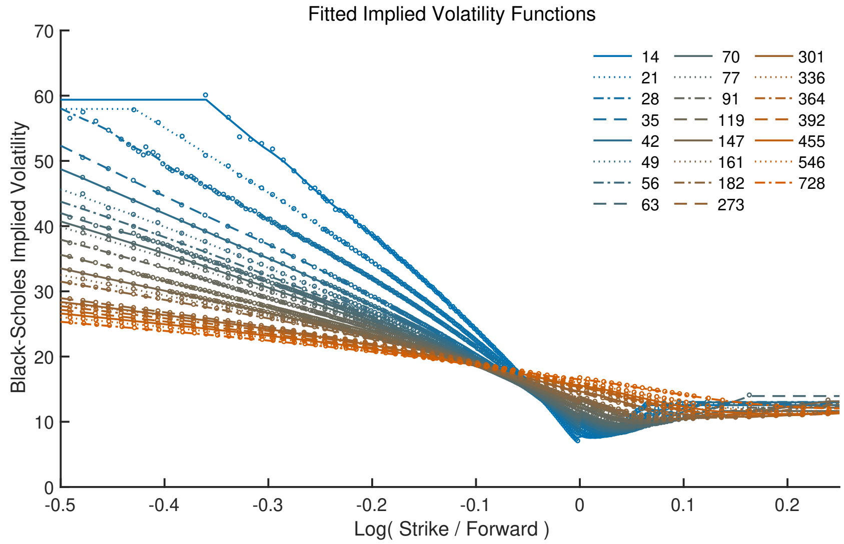

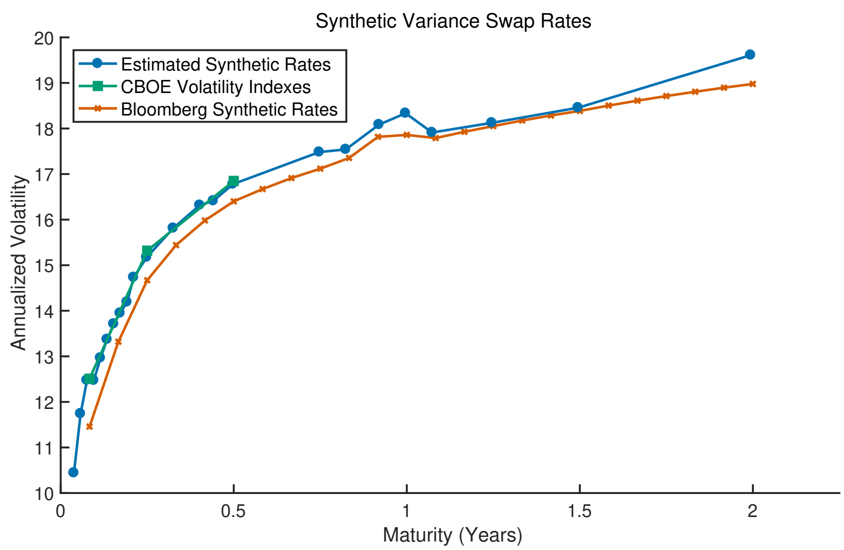

Figure A.1: This figure illustrates the synthetic variance swap rate estimation procedure for the S&P 500 index on December 20, 2019. The top plot reports the fitted implied volatility functions against log-moneyness for different maturities with the legend indicating the calendar days until expiration. I estimate the synthetic variance swap rates as a weighted average of out-of-the money option prices following Carr and Wu (2009) for each maturity. As the plot indicates, I extrapolate the strike price range with non-zero bids by appending log-Normal tails with flat implied volatility functions. The bottom plot reports the estimated synthetic variance swap rates against the CBOE volatility indices and synthetic variance swap rates from Bloomberg. In the paper, I interpolate the estimated synthetic rates onto a monthly grid from one-month to two-years. The hump in the term-structure for maturities around one-year reflects market pricing of the 2020 election at the end of 2019.

SPX Index Option Maturities with Estimated Synthetic Variance Swap Rates

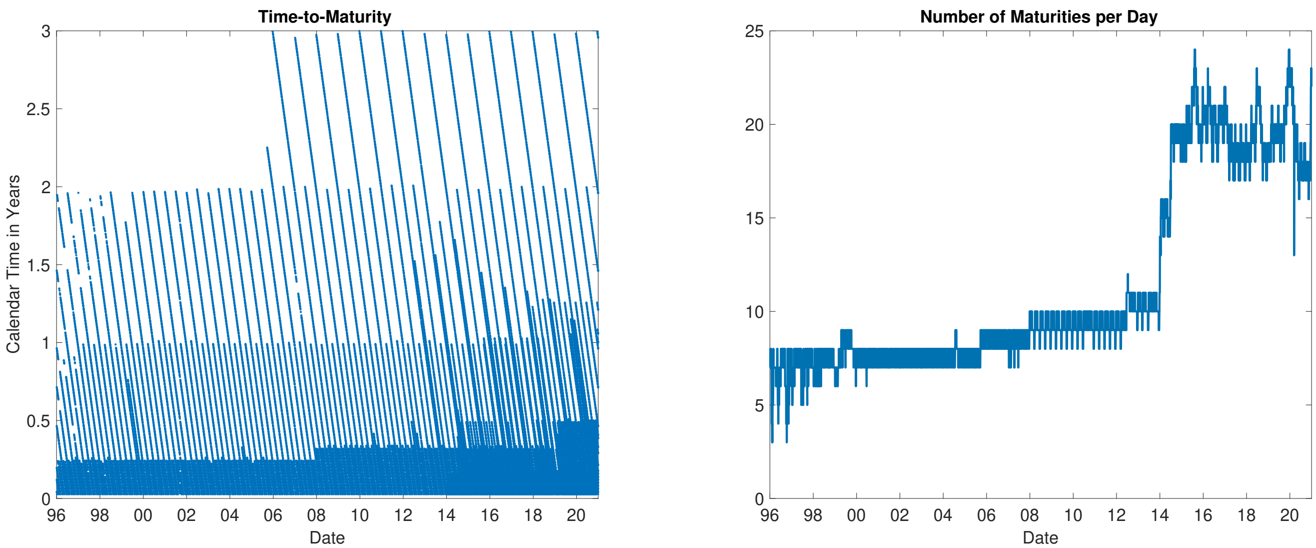

Figure A.2: This figure plots the time-to-maturity of the SPX index option maturities for which synthetic variance swap rates were estimated and the number of maturities for each day in the sample.

One-Month Synthetic Variance Swap Rate versus the VIX

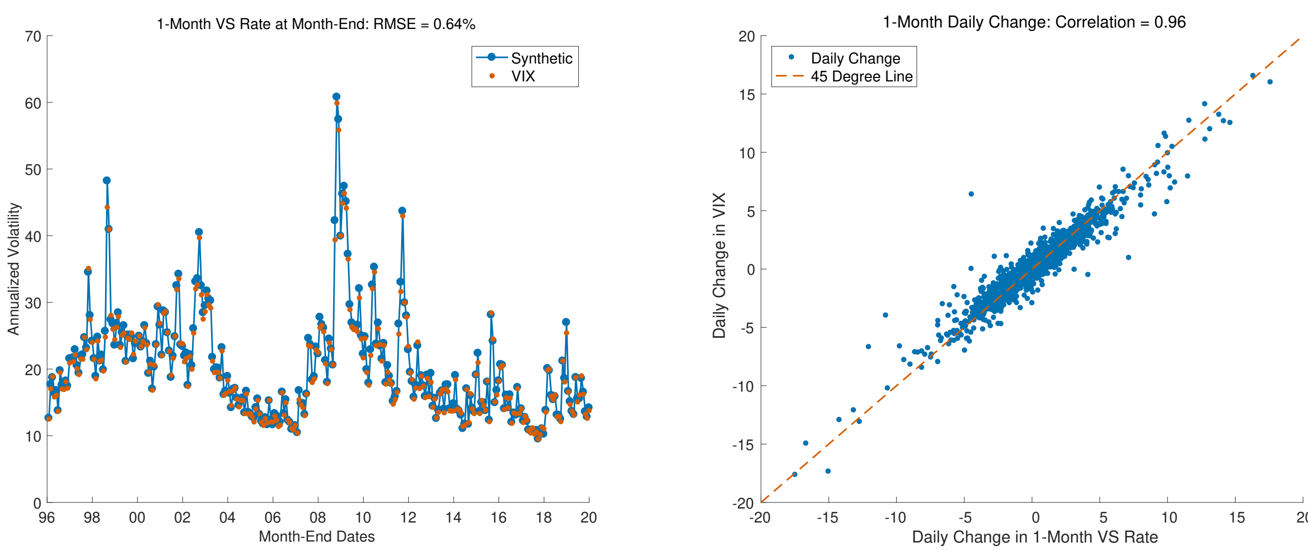

Figure A.3: This figure plots the estimated one-month synthetic variance swap rate against the VIX index from 1996 to 2020. The left plots report the time series at month-end dates. The right plots report the daily changes which are 96% correlated.

VAR Residuals in the Baseline Model and Autocorrelation Tests

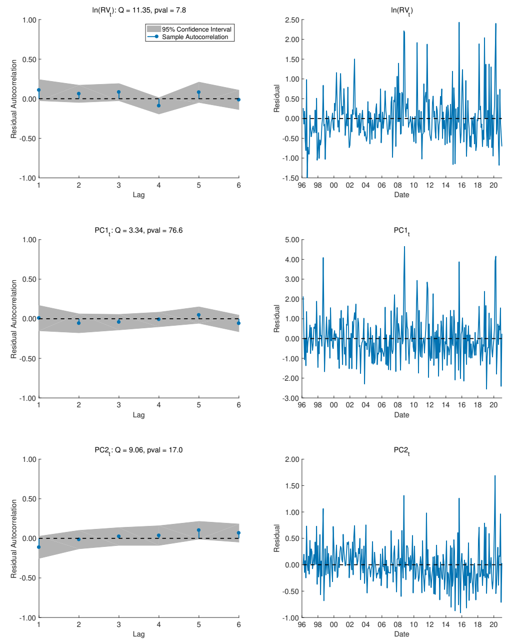

Figure A.4: The left plots report the autocorrelation estimates of the VAR residuals in the baseline three-factor model alongside 95% pointwise confidence intervals for the first six lags. The autocorrelation estimates are close to zero and insignificant in univariate tests. The title reports the Ljung-Box test-statistic and p-value for the null hypothesis that the autocorrelations are jointly equal to zero which is not rejected at the 5% level for any of the state variables. The right plots report the time-series of the VAR residuals.

RVTP and IVTP Sensitivites to the State Vector

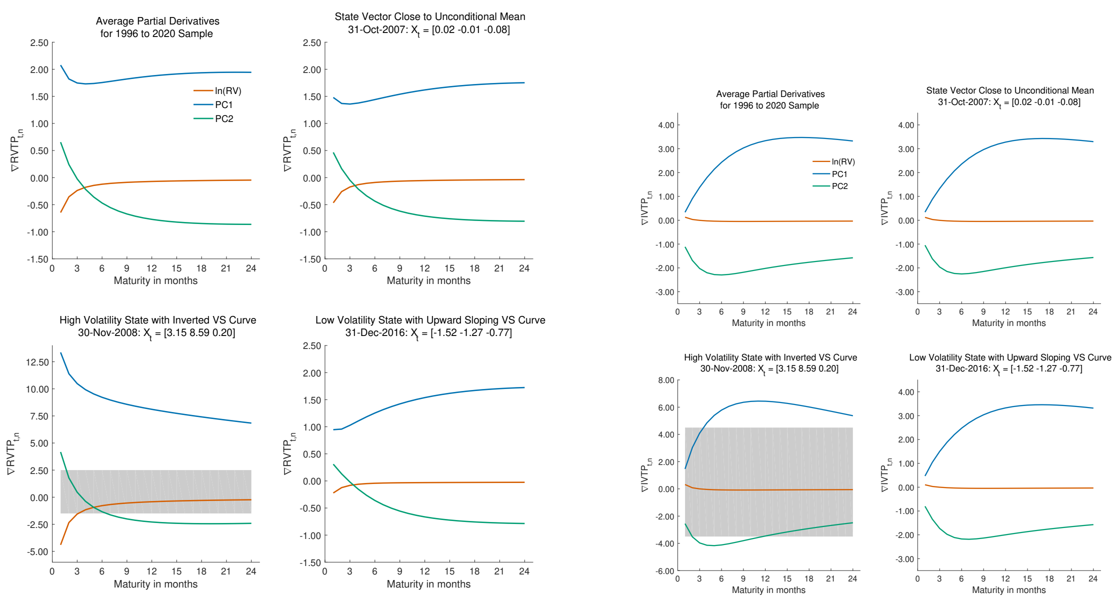

Figure A.5: The left (right) figure plots the sensitivity of realized variance term premia $\nabla RVTP_{t,n}$ (implied volatility term premia $\nabla IVTP_{t,n}$) to a one standard deviation increase in the state variables on average and at different points in time.

Robustness to Model Specification

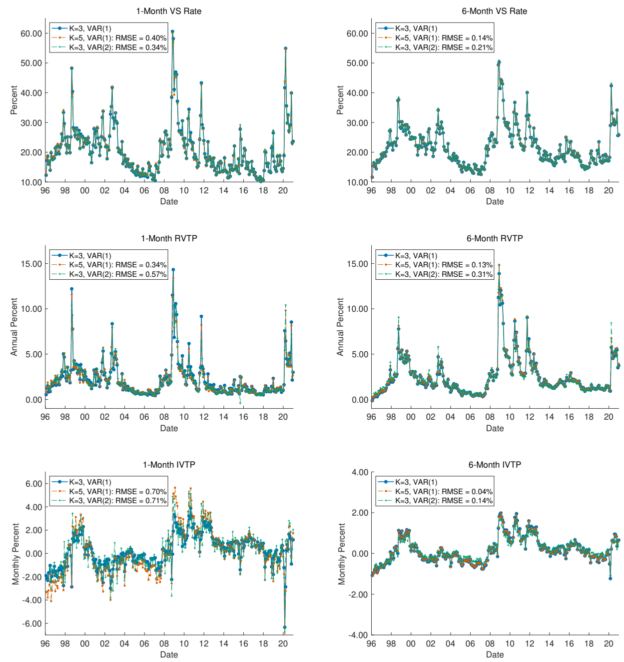

Figure A.6: This figure plots the model estimates of variance swap rates and realized variance term premia for the baseline three-factor VAR(1) model against a five-factor VAR(1) model with daily and weekly log realized variance factors and a three-factor VAR(2) model.

Robustness to Time-Varying VAR Covariance Matrix

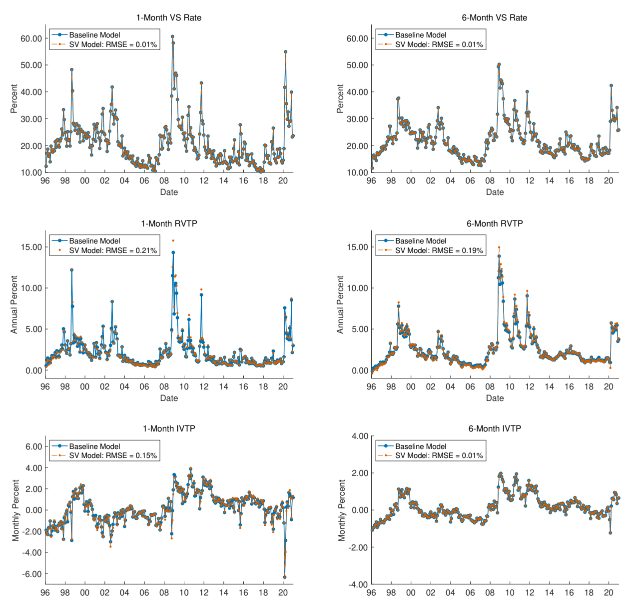

Figure A.7: This figure plots the model estimates of variance swap rates and RVTP for the baseline model and for a model with a time-varying covariance matrix for the VAR(1) residuals.

IRFs of RVTP and IVTP to a VIX Shock in the Eraker and Wu Model

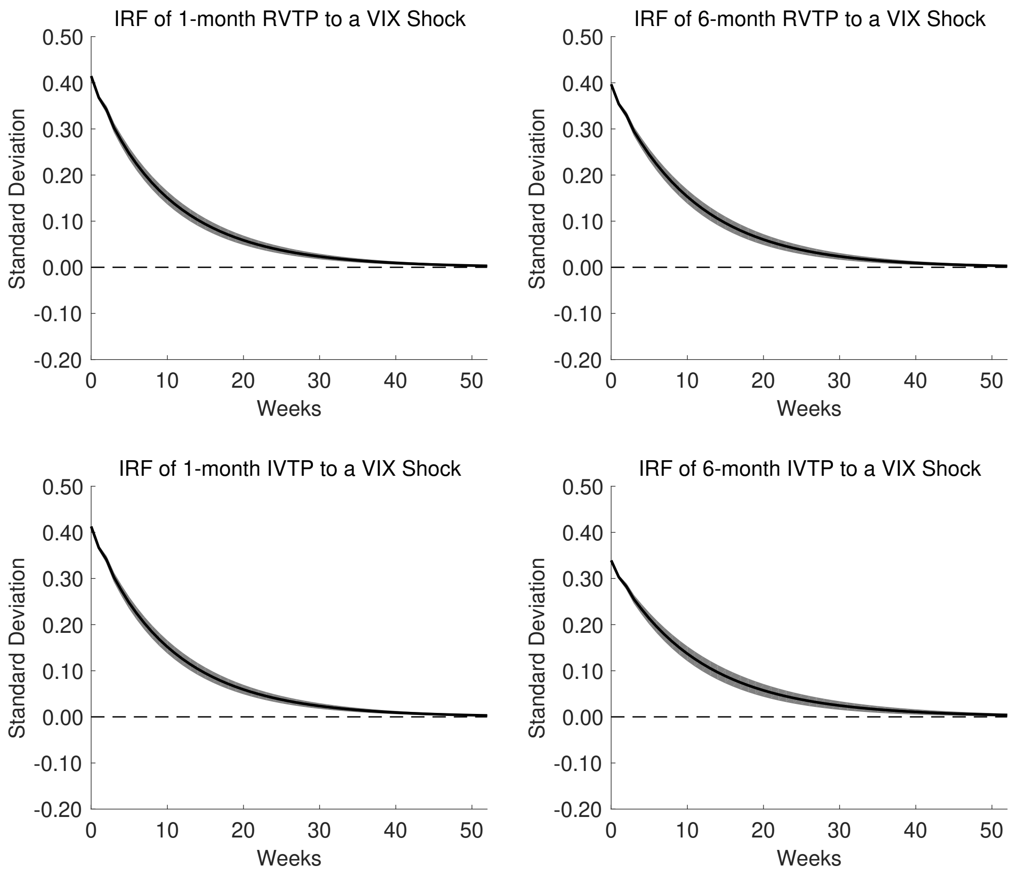

Figure A.8: This figure plots the IRFs of RVTP and IVTP for one-month and six-month horizons to a one standard deviation VIX shock from a bivariate VAR for 100 years of data simulated from the two-factor Eraker and Wu (2017) equilibrium model. The VAR is estimated with weekly data using four lags and the IRFs are from a Cholesky decomposition with the VIX ordered first.

Relative Pricing of VIX Futures Across Contracts

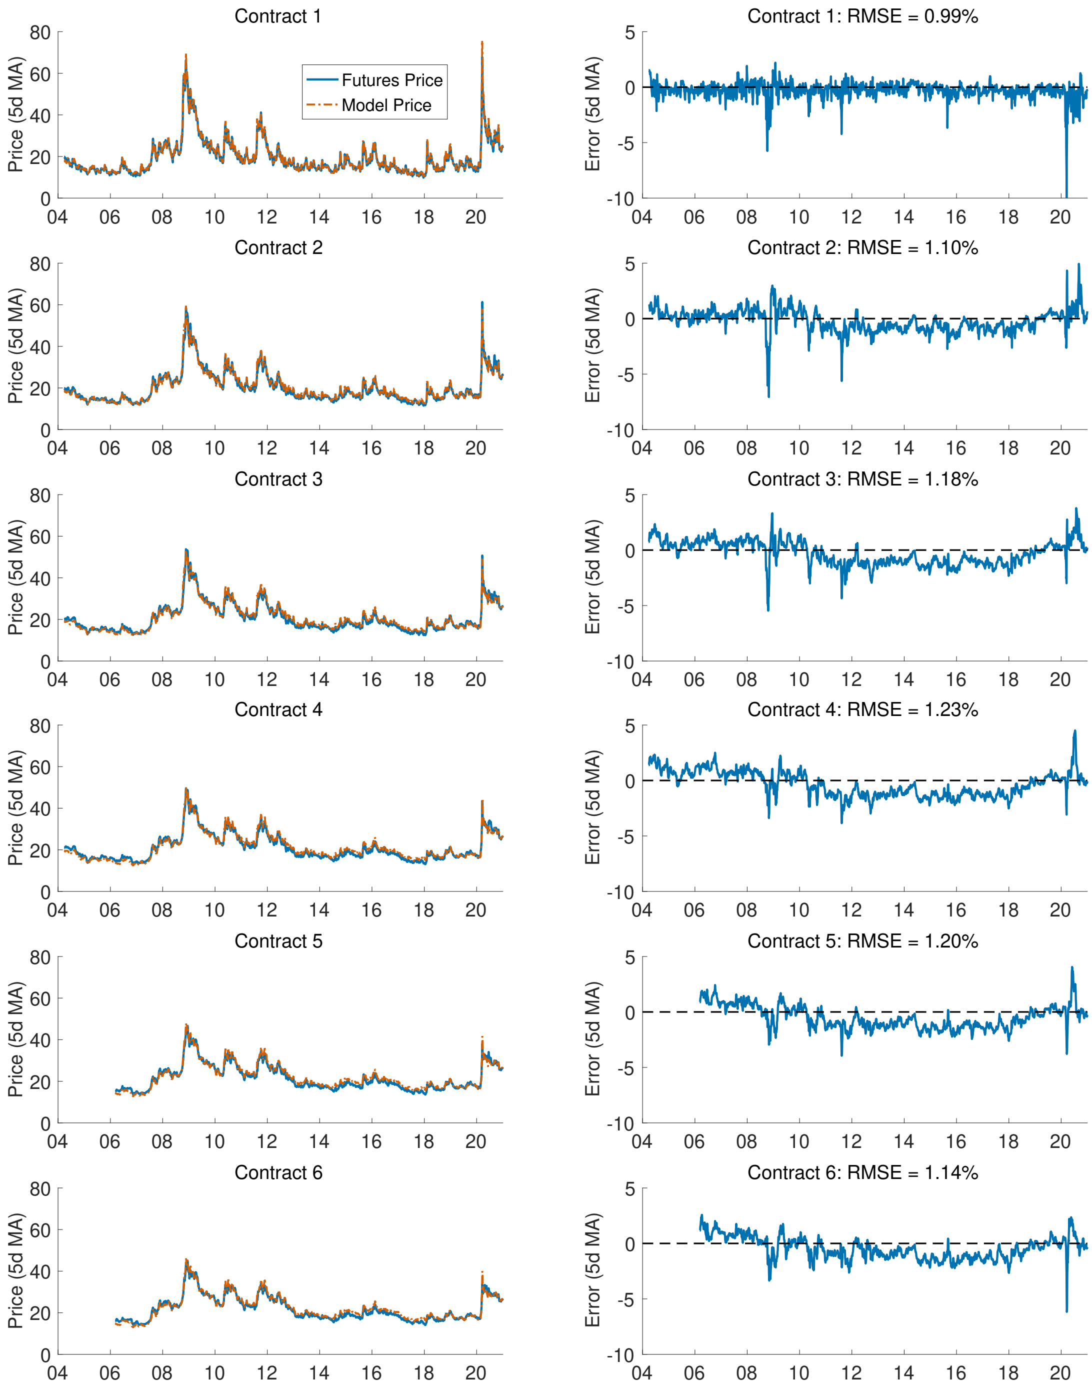

Figure A.9: The left figures plot the model price against the futures price for the front six contracts. The right figures plot the difference between the futures price and model price. Negative values in the right plots indicate that VIX futures are cheap relative to the option-implied model price. The plots report five-day moving averages with the root mean-squared-error (RMSE) of the pricing error in the right plot titles.