JOURNAL OF FINANCIAL AND QUANTITATIVE ANALYSIS Vol. 52, No. 6, Dec. 2017, pp. 2461–2490

doi:10.1017/S0022109017000825

Risk Premia and the VIX Term Structure

Travis L. Johnson*

Abstract

The shape of the Chicago Board Options Exchange Volatility Index (VIX) term structure conveys information about the price of variance risk rather than expected changes in the VIX, a rejection of the expectations hypothesis. The second principal component, SLOPE, summarizes nearly all this information, predicting the excess returns of synthetic Standard & Poor's (S&P) 500 variance swaps, VIX futures, and S&P 500 straddles for all maturities and to the exclusion of the rest of the term structure. SLOPE's predictability is incremental to other proxies for the conditional variance risk premia, economically significant, and inconsistent with standard asset pricing models.

I. Introduction

The Chicago Board Options Exchange (CBOE) Volatility Index (VIX), the most widely followed index of market volatility, is an estimate of Standard & Poor's (S&P) 500 return volatility over the next month derived from S&P 500 option prices. This estimate reflects both the conditional expectation of future S&P 500 volatility and a risk premium inherited from the options it is based on. Previous papers show that options are priced as if volatility were higher than it actually is, indicating a negative variance risk premium (see Coval and Shumway (2001), Bakshi and Kapadia (2003), and Bakshi and Madan (2006)). As a result, the VIX systematically overestimates realized volatility, and assets with positive variance risk exposure earn negative abnormal returns.

I follow the methodology used to compute the VIX to form the VIX term structure, estimates of annualized S&P 500 return volatility over the next 1, 2, 3, 6, 9, and 12 months. Just as the VIX is composed of both conditional

*Johnson (corresponding author), travis.johnson@mccombs.utexas.edu, University of Texas at Austin McCombs School of Business. I thank Anat Admati, Mary Barth, Hendrik Bessembinder (the editor), Bjorn Eraker, Sebastian Infante, Bryan Kelly (the referee), Arthur Korteweg, Kristoffer Laursen, Ian Martin, Stefan Nagel, Paul Pfleiderer, Monika Piazzesi, Jan Schneider, Ken Singleton, Eric So, Suhas Sridharan, Mitch Towner, and seminar participants at Boston College, Dartmouth College, Rice University, Stanford University, University of California–Berkeley, University of Houston, University of Maryland, University of Pennsylvania, University of Rochester, University of Texas at Austin, and University of Wisconsin–Madison for their helpful comments. This paper is based on my dissertation at Stanford University titled "Essays on Information in Options Markets."

volatility expectations and a risk premium, the shape of the VIX term structure reflects both the expected path of future return volatility and different risk premia associated with variance risk at different maturities. For example, there are two potentially complementary explanations for a downward-sloping VIX term structure: Markets expect return variance to decline, and exposure to short-term variance risk commands a larger risk premium than exposure to long-term variance risk.

In this article, I estimate the extent to which time variations in the shape of the VIX term structure reflect changes in the expected path of future VIX (the "expectations hypothesis") and, conversely, the extent to which they reflect changes in variance risk premia. Across 10 specifications with forecast horizons of 1 month and 1 quarter, I strongly reject the expectations hypothesis. This implies changes in the premium investors pay for variance assets with different maturities drive much of the variation in the shape of the VIX term structure, meaning this shape should predict excess returns of variance assets.

I find that a single factor, the second principal component (PC) SLOPE, summarizes nearly all information about variance risk premia in the VIX term structure. SLOPE negatively predicts future returns of 18 variance assets: 6 maturities each for synthetic S&P 500 variance swaps, VIX futures, and S&P 500 straddles. More surprising, the rest of the VIX term structure adds almost no predictive power for returns incremental to SLOPE, meaning that although many factors are required to describe movements in the VIX and its term structure, only SLOPE is related to movements in variance risk premia.

My methodology and results are similar to those in Cochrane and Piazzesi (2005), who show that although many factors are required to describe movements in bond yields and their term structure, only one of these factors is related to movements in bond risk premia. In addition to studying a different set of assets, the primary difference between this article and Cochrane and Piazzesi is that I do not use a first-stage regression to find the single linear combination of the term structure that best predicts average asset returns. Instead, I use the second PC as the single factor because it predicts variance asset returns better than do the other PCs. My informal selection approach is more conservative than Cochrane and Piazzesi because it is not designed to find the best possible single linear factor.

In addition to summarizing nearly all information about variance risk premia in the VIX term structure, SLOPE is an economically significant and robust predictor of variance asset returns. As an illustration of its economic significance, the difference in next-day (next-month) returns across extreme SLOPE quintiles for the 18 variance assets ranges from 29 basis points (bps) to 181 bps (7.4% to 36.4%). This predictive relation is robust to alternate forecast horizons, removing extreme SLOPE events from the sample, and alternate definitions of SLOPE. Furthermore, SLOPE predicts returns incrementally to other indicators that the prior literature suggests are related to variance risk premia, including estimates of implied minus expected variance.

To measure variance risk premia, I use future returns of variance-sensitive investments studied in the literature: variance swaps as in Dew-Becker, Giglio, Le, and Rodriguez (2017), VIX futures as in Eraker and Wu (2014), and S&P

500 straddles as in Coval and Shumway (2001). Although I observe returns for VIX futures and S&P 500 straddles, I do not observe variance swap returns and therefore use returns of option portfolios designed to replicate variance swaps ("synthetic variance swaps"). Variance asset returns are better suited to this study than differences between option-implied and expected or realized variance (used in Todorov (2010), Carr and Wu (2009), and elsewhere) for the reasons detailed in Section III, the most important of which is they allow me to examine the next-day and next-month risk premia associated with changes in variance at different maturities. Differences between option-implied and expected or realized variance, by contrast, can be estimated for different maturities but doing so results in estimates of risk premia over the entire time to maturity. Therefore, any differences in estimated risk premia could be due to differences across maturities in next-day or next-month risk premia, or differences across future horizons in risk premia. By using variance asset returns, I rule out the latter and focus on the former.

My results provide three puzzling empirical patterns for future work on variance risk premia to explain. The first is the insignificant relation between the first PC (LEVEL) of the VIX term structure and variance asset returns. Most option pricing models (e.g., Heston (1993)), models of variance risk premia (e.g., Bakshi and Madan (2006)), and asset pricing models (e.g., Merton (1973), Martin (2013), and Campbell, Giglio, Polk, and Turley (2017)) predict risk premia are high when the LEVEL (not SLOPE) of volatility is high. The second is the positive relation between SLOPE and conditional variance risk premia together with the negative relation between maturity and unconditional variance risk premia. Explaining these facts together requires investors to be more averse to increases in short-term variance than long-term variance, although long-term variance asset prices increase more in times with large variance risk premia.

A third puzzling empirical pattern is that when SLOPE is low, future variance risk premia are not just smaller, they actually change sign and become positive for 17 of 18 variance assets. For example, 12-month S&P 500 straddles have average returns of 30 bps per day above the risk-free rate when SLOPE is in its lowest quintile. This indicates the correlation between variance fluctuations and marginal utility changes sign over time, meaning investors who normally pay large premia to protect against variance increases occasionally worry about variance decreases and therefore price variance assets at a discount.

Taken together, the results in my article have important implications for both researchers in financial economics and investors in variance assets. For researchers, my results provide surprising patterns for new theories of variance risk premia to explain and allow future empirical work to easily summarize all variance risk premia information in the VIX term structure using SLOPE alone. For investors or traders using variance-sensitive assets such as VIX exchange-traded notes (ETNs) or S&P 500 options, my results show SLOPE is an economically significant and timely indicator of expected returns.

The remainder of this article is organized as follows: Section II describes my article's relation to prior research, Section III details the construction of the VIX term structure and variance asset returns, Section IV presents my empirical results, and Section V concludes.

II. Relation to Prior Research

This article builds on prior research studying the unconditional and conditional risk premium associated with innovations in marketwide variance. The unconditional variance risk premium is negative (see, e.g., Coval and Shumway (2001), Bakshi and Kapadia (2003), Broadie, Chernov, and Johannes (2009), and Carr and Wu (2009)), meaning that assets whose value is increasing in market volatility earn negative risk premia and option-implied volatility is higher than average realized volatility. Furthermore, Aıt-Sahalia, Karaman, and Mancini (2015), Dew-Becker et al. (2017), and Eraker and Wu (2014) show unconditional variance risk premia are downward sloping, meaning risk premia are largest for variance assets with shorter maturities. In fact, long-dated variance assets have an unconditional risk premium close to zero. As discussed in Dew-Becker et al., this is difficult to reconcile with most neo-classical asset pricing models, which predict that unconditional variance risk premia are upward sloping. My results add to this challenge by showing the conditional variance risk premia are larger for all maturities when the price of short-dated variance assets is abnormally low relative to the price of long-dated variance assets.

I contribute most directly to the literature studying determinants of conditional variance risk premia. Many models and empirical analysis (e.g., Heston (1993), Bakshi and Kapadia (2003), and Bakshi and Madan (2006)) suggest variance risk premia are larger when volatility is high. Todorov (2010) and Aıt-Sahalia et al. (2015) show variance risk premia are larger following downward jumps in equity prices. Relatedly, Corradi, Distaso, and Mele (2013) find variance risk premia are larger in times with recent stock market declines and high volatility. Barras and Malkhozov (2016) show that variance risk premia are related to the risk-bearing capacity of broker-dealers, as proxied by their aggregate leverage ratio. Finally, Feunou, Fontaine, Taamouti, and Tédongap (2014) show that 2 factors from the variance term structure predict future excess variance at forecast horizons of 1 to 12 months. I add to this literature by showing conditional variance risk premia information in the VIX term structure, a natural indicator that includes the level of volatility, is summarized by SLOPE. Furthermore, unlike many of the papers in this area, I assess conditional variance risk premia using variance assets with many different maturities and show that SLOPE predicts their future returns incrementally to existing indicators.

My single-factor results are particularly surprising given that multiple volatility factors are necessary in many other settings. Pricing the cross section of equity returns (Adrian and Rosenberg (2008)), explaining the dynamics of the VIX term structure (Egloff, Leippold, and Wu (2010)), pricing the S&P 500 volatility surface (Christoffersen, Heston, and Jacobs (2009), Christoffersen, Jacobs, Ornthanalai, and Wang (2008)), and pricing VIX options (Mencía and Sentana (2013)) are all dramatically improved by adding a second volatility factor. Although I also find that the dynamics of the VIX term structure are explained well by two factors (LEVEL and SLOPE), only SLOPE is consistently related to variance risk premia.

Bollerslev, Tauchen, and Zhou (2009) and Drechsler and Yaron (2011) study the relation between variance risk premia and equity risk premia. Both papers show that a proxy for conditional variance risk premia, the difference between VIX^2 and an estimate of statistical-measure variance, positively predicts equity returns. In untabulated tests, I find that SLOPE does not predict equity returns despite being positively correlated with implied minus statistical variance. Instead, the equity return predictability afforded by the VIX term structure in Bakshi, Panayotov, and Skoulakis (2011) and Feunou et al. (2014) is attributable to other PCs of the VIX term structure, none of which predicts variance asset returns. This indicates that either SLOPE predicts variance asset returns for nonrisk reasons such as mispricing or demand-based price impacts (as in Garleanu, Pedersen, and Poteshman (2009)) or SLOPE represents a type of variance risk premia outside the Bollerslev et al. (2009) and Drechsler and Yaron (2011) models.

My evidence on the returns of a SLOPE-based dynamic straddle strategy builds on the extant work studying dynamic variance asset portfolios in Aıt-Sahalia et al. (2015), Filipović, Gourier, and Mancini (2016), and Egloff et al. (2010). Unlike these papers, my goal is not to compute an optimal portfolio strategy, but rather to document that the predictability afforded by the VIX term structure is summarized by a single factor.

III. Constructing the VIX Term Structure and Variance Asset Returns

A. VIX Term Structure

A key construct in my analysis is the VIX term structure, which I compute by replicating the CBOE's VIX calculation, but with target maturities longer than 1 month. The VIX calculation is an estimate of the model-free implied volatility measure originating in Breeden and Litzenberger (1978). If options are available for every strike price, the VIX equals:

$$\text{VIX}^2_{T,t} \equiv \frac{2e^{rT}}{T}\left\{\int_0^{F_t}\frac{1}{K^2}\text{PUT}_t(K;t+T)dK + \int_{F_t}^{\infty}\frac{1}{K^2}\text{CALL}_t(K;t+T)dK\right\}, \tag{1}$$

where F_t is the time t forward price of the S&P 500 at time $t+T$, and $\text{PUT}_t(K;t+T)$ and $\text{CALL}_t(K;t+T)$ are the prices at time t of puts and calls expiring at time $t+T$ with strike price K. As shown in Neuberger (1994) and Carr and Madan (1998), if the S&P 500 follows a diffusion process $dS_t/S_t = rdt+\sigma_t dZ_t$ under the risk-neutral measure, $\text{VIX}^2_{T,t}$ equals the risk-neutral expectation of average future instantaneous variance:

$$\text{VIX}^2_{T,t} = \frac{1}{T}\mathbb{E}^{\mathbb{Q}}_t\left[\int_t^{t+T}\sigma_s^2 ds\right]. \tag{2}$$

The standard approach to estimating the VIX equation (1) empirically, used by the CBOE to compute the VIX, discretizes the integral at the available strike

prices and truncates it at the smallest and largest available strike prices, making the expression:

$$(3) \qquad \hat{\mathrm{VIX}}^2_{T,t} \equiv \frac{2e^{rT}}{T}\sum_{K_i}\frac{1}{K_i^2}\mathrm{OPTION}_t(K_i;t+T)\Delta K_i,$$

where $\mathrm{OPTION}_t(K_i;t+T)$ is the price of the out-of-the-money option for strike K_i at time t with expiration date $t+T$. The VIX calculation (see www.cboe.com/micro/vix/vixwhite.pdf for details) further specifies how to determine which option is out of the money and provides additional corrections, all of which I follow.

Using closing quotes for S&P 500 index options and risk-free rates from 1996 through 2013 via OptionMetrics, I compute $\mathrm{VIX}_{T,t}$ for $T=1$, 2, 3, 6, 9, and 12 months at the close of each day t. These maturities represent the approximate times to expiration typically available for index options. Together, they form the VIX term structure at t.

Table 1 presents descriptive statistics for the VIX term structure. In both medians and means, long-term VIX are higher than short-term VIX, indicating that the average term structure is upward sloping. There is substantial variability in the VIX at all horizons, though short-term VIX are more volatile than long-term VIX. One potential reason is that because return volatility is mean reverting, times with high (low) VIX have not so high (low) long-term VIX. Another potential reason is that risk premia change more over time for short-dated variance risk than for long-dated variance. In the analysis that follows, I provide evidence that

TABLE 1

Summary Statistics for the VIX Term Structure and Its PCs

Table 1 presents summary statistics for the VIX term structure and its principal components (PCs). The VIX term structure is an annualized model-free estimate of option-implied volatility for the Standard & Poor's (S&P) 500 index 1, 2, 3, 6, 9, and 12 months into the future. Panel A presents the summary statistics for the term structure. Panel B presents both the definitions and variances of the PCs of the implied variance (VIX^2) term structure. The sample contains 4,445 daily observations from 1996 through 2013.

Panel A. VIX Term Structure

Statistic	VIX_1	VIX_2	VIX_3	VIX_6	VIX_9	VIX_{12}
Mean	21.7%	22.0%	22.2%	22.6%	22.6%	22.7%
Standard dev.	8.5%	7.9%	7.5%	6.7%	6.3%	6.1%
1st percentile	10.6%	11.3%	11.8%	12.9%	13.1%	13.3%
10th percentile	13.0%	13.5%	14.0%	14.8%	15.1%	15.4%
25th percentile	16.0%	16.6%	17.0%	17.9%	18.1%	18.3%
Median	20.2%	20.7%	21.2%	21.9%	22.1%	22.1%
75th percentile	24.9%	25.3%	25.5%	26.1%	25.7%	26.2%
90th percentile	31.5%	30.8%	30.5%	30.5%	30.2%	30.4%
99th percentile	54.5%	52.0%	50.7%	45.7%	43.5%	41.5%

Panel B. Principal Component Definitions and Variance

	LEVEL PC1	SLOPE PC2	CURVE PC3	PC4	PC5	PC6
VIX_1^2	0.52	−0.57	−0.55	0.16	0.04	−0.28
VIX_2^2	0.48	−0.24	0.24	−0.25	0.04	0.77
VIX_3^2	0.44	−0.01	0.62	−0.33	−0.01	−0.57
VIX_6^2	0.36	0.30	0.15	0.65	−0.58	0.07
VIX_9^2	0.32	0.44	−0.02	0.32	0.78	0.02
VIX_{12}^2	0.29	0.58	−0.48	−0.53	−0.25	−0.01
Variance ($\times 10^5$)	100.73	4.75	0.32	0.16	0.13	0.07
% of total	94.89	4.47	0.30	0.15	0.12	0.07

both potential reasons contribute to the relative movements of long- and short-dated volatility.

Given the strong correlations between VIX at different horizons, a natural way to study variations in the shape of the term structure is to rotate it into 6 orthogonal PCs. I apply this linear rotation to option-implied variances (VIX^2) rather than to volatilities because return variances combine linearly across maturity (assuming no autocorrelation in returns) whereas volatilities do not. Panel B of Table 1 shows definitions of and summary statistics for the resulting PCs, scaled so that their variances equal the 6 eigenvalues of the VIX term structure's covariance matrix. The first PC loads positively on all 6 VIX^2, and therefore reflects the LEVEL of the term structure. The second PC loads negatively on short-horizon VIX^2 but positively on long-horizon VIX^2 and therefore reflects the slope of the term structure. When SLOPE is low (high), the term structure is downward (upward) sloping. Note that the positive coefficients in the definition of SLOPE have larger magnitudes than do the negative coefficients.[1] As a result, it is possible for SLOPE to be positive on a day when the VIX term structure is strictly decreasing. For this reason, my analyses compare high-SLOPE periods to low-SLOPE periods, making the average SLOPE irrelevant.

Figure 1 plots the standardized LEVEL and SLOPE PCs of the VIX term structure. LEVEL follows the familiar pattern of the VIX, remaining low and stable during normal times and spiking upward during market downturns. In normal times, SLOPE is high, indicating an upward-sloping term structure. Furthermore, in low-volatility times there is a clear positive correlation between SLOPE and LEVEL. When LEVEL spikes upward, however, SLOPE spikes downward, indicating a downward-sloping term structure and negative correlation between SLOPE and LEVEL. This time-varying correlation averages out to 0, by construction, in the full sample. In Section IV.D, I further discuss these changes in correlation and address the non-normality of SLOPE apparent in Figure 1.

B. Variance-Sensitive Asset Returns

The most common definition of the conditional variance risk premia is the difference between conditional variance under the risk-neutral and physical measures. However, this quantity is inherently unobservable because asset prices reflect risk-neutral rather than physical variance. Therefore, to investigate the information in the VIX term structure about variance risk premia at different maturities, I need to estimate variance risk premia for each day t and each maturity T. Other papers estimate this quantity using:

i) $VIX^2_{T,t} - \hat{\mathbb{E}}_t(RV^2_{t+1,t+T})$, the difference between option-implied and expected future realized variance, where the physical-measure expectation is based on a statistical model (Todorov (2010), Bekaert and Hoerova (2014), and others), and

[1] SLOPE needs to have positive average loadings to be uncorrelated with LEVEL because the VIX term structure tends to be downward sloping (in a geometric sense) when volatility is high due to mean reversion, meaning a zero-sum definition of SLOPE would be negatively correlated with LEVEL.

ii) $\text{RV}^2_{t+1,t+T} - \text{VIX}^2_{T,t}$, the difference between future realized variance and option-implied variance (Carr and Wu (2009), Feunou et al. (2014), and others).

The problem with these measures in this setting is that they do not allow me to examine the premia associated with variance risk at many maturities while holding the forecast horizon fixed. Both can be estimated for different maturities T, but doing so results in estimates of risk premia over the entire time to maturity. As a consequence, any differences in estimated risk premia could be due to differences across maturities in risk premia or differences across future horizons in risk premia. For this reason, I use returns of variance-sensitive assets with different maturities to proxy for variance risk premia. These assets offer different exposures to the (potentially) many variance risk factors reflected in the VIX term structure while allowing me to hold the forecast horizon and holding period fixed.

Variance asset returns offer two other advantages over measures based on comparisons of option-implied and realized or model-expected variance. The first is they directly relate to asset pricing models that study the risk premia associated with investable assets. The second is variance asset returns do not depend on which statistical model is used to estimate $\hat{\mathbb{E}}_t(\text{RV}^2_{t+1,t+T})$ or $\text{RV}^2_{t+1,t+T}$.

The first variance asset I use is an S&P 500 variance swap, a contract that swaps a fixed payment for a variable amount proportional to the realized variance of the S&P 500 index (as used in Dew-Becker et al. (2017)). Without over-the-counter swap pricing data, I proxy for variance asset returns using the returns of "synthetic variance swaps," option portfolios designed to replicate variance swaps. The key insight behind these replicating portfolios is that $\text{VIX}^2_{t,T}$ is the price of a particular portfolio of traded options that replicates a variance swap

FIGURE 1

Level and Slope of the VIX^2 Term Structure

Figure 1 presents the first two principal components (PCs) of the VIX^2 term structure. The VIX^2 term structure is an annualized model-free estimate of option-implied variance for the Standard & Poor's (S&P) 500 index 1, 2, 3, 6, 9, and 12 months into the future. I plot its first two PCs, which I call LEVEL and SLOPE, standardized so that both have mean of 0 and standard deviation of 1. The PCs are defined in Table 1. The sample contains 4,445 daily observations from 1996 through 2013.

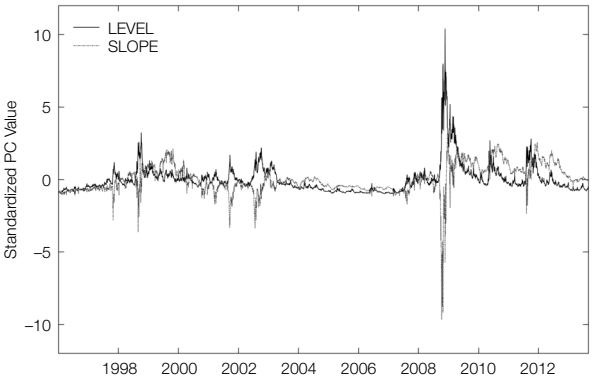

(Carr and Madan (1998)). As detailed in the Appendix, the daily synthetic variance swap returns I use are the day return of this replicating portfolio of options. To keep the maturity constant and the replicating portfolio as accurate as possible, the monthly synthetic variance swap returns I use are the daily returns compounded from $t+1$ through $t+21$. To match the VIX term structure, I construct returns for synthetic variance swaps with $T=1$, 2, 3, 6, 9, and 12 months to maturity.

Note that although the portfolio of options that replicates a variance swap has a value on day t that is proportional to $\text{VIX}^2_{T,t}$, its value on day $t+1$ is not proportional to $\text{VIX}^2_{T,t+1}$ or $\text{VIX}^2_{T-1,t+1}$. The reason is the portfolio of options used to replicate a variance swap at t has different weights and uses different options than the portfolio used at $t+1$. This difference is critical because the VIX index is not directly investable and has no drift driven by variance risk premium, whereas the returns of synthetic variance swaps I study are investable (up to transaction costs) and subject to risk-based drift.

On most days, there are no options expiring exactly T months later, and so the VIX calculation uses a linear combination of options with the two nearest expiration dates to $t+T$. As detailed in the Appendix, I use this same linear combination to form a portfolio at time t and compute its time $t+1$ returns, resulting in a "constant maturity" strategy. Because of the mismatch between T and available expiration dates, and because of the discreteness in strike prices, the portfolio returns I use are imperfect proxies for true variance swap returns. However, unlike changes in the VIX itself, these returns are tradable (the portfolio weights sum to 1) and not interpolated.

Although variance swaps are the most direct measure of variance risk, without data on over-the-counter pricing their empirical implementation requires computing the return of the option portfolio described above and is therefore subject to more illiquidity-driven noise than directly observed variance assets. For this reason, the second variance asset I study is VIX futures, promises to exchange a fixed payment for the prevailing VIX index value at a prespecified date (as used in Eraker and Wu (2014)). These contracts have traded since 2004, and historical end-of-day data are available on CBOE.com, meaning no replication is necessary. Following the approach detailed in the Appendix, I compute daily returns of constant maturity VIX futures strategies that use a mixture of the two maturity dates nearest to a target maturity date $T=1$, 2, 3, 4, 5, or 6 months from the current date t and compound them to compute monthly returns. I use these 6 target maturities because there are reliable data on contracts with approximately these times to maturity starting in 2004.

Like the synthetic variance swap returns, the VIX futures returns I construct are for portfolios formed using information available on day t, held for a day, and then rebalanced using information available on day $t+1$. Constant maturity strategies are common in the VIX futures market; for example, the popular VXX ETN uses a constant maturity strategy with target maturity of 1 month, and so my daily 1-month VIX futures returns are nearly identical (correlation 96%, result untabulated) to those of VXX.

Although VIX futures returns are more liquid than the out-of-the-money options used to compute synthetic variance swap returns, they only started trading in 2004, limiting the power of return predictability tests. As a middle ground, I compute at-the-money S&P 500 straddle returns (as used in Coval and Shumway (2001)). At-the-money options are the most liquid, and straddle portfolios have many fewer positions and therefore smaller transactions costs than option portfolios replicating variance swaps. Moreover, straddle returns are available since the beginning of the OptionMetrics data set in 1996. Following the approach detailed in the Appendix, I compute the daily returns of constant maturity straddle strategies that use a mixture of the 2 maturity dates nearest a target maturity date that is always $T = 1, 2, 3, 6, 9$, or 12 months from the current date and compound them to compute monthly returns.

Table 2 provides summary statistics for the returns of these 18 variance-sensitive assets. Because of the variance risk to which they are exposed, these assets have substantially negative abnormal returns, as low as -1% per day or -18% per month. They are also subject to extreme volatility, as high as 15% per day. Nevertheless, they offer significantly negative Sharpe ratios, most ranging from -0.25 to -0.75 on an annualized basis.

Table 2 also shows there are significant differences in average risk premia earned by these assets. Specifically, synthetic variance swaps tend to have more negative Sharpe ratios than VIX futures and S&P 500 straddles, and longer maturity variance assets tend to have less negative Sharpe ratios than shorter maturity assets, echoing the results in Dew-Becker et al. (2017) and Eraker and Wu (2014).

To help explain the negative risk premia earned by these assets, I appeal to the pricing model used in Ang, Hodrick, Xing, and Zhang (2006) and regress excess asset returns on contemporaneous excess market returns and innovations in the 1-month VIX index:

$$(4) \quad r_{i,t} - r_{f,t} = \alpha_i + \beta_{i,\text{MKT}}(r_{\text{MKT},t} - r_{f,t}) + \beta_{i,\Delta\text{VIX}}(\text{VIX}_{1,t} - \text{VIX}_{1,t-1}) + \epsilon_{i,t}.$$

For each asset, Table 2 reports estimates of α_i, $\beta_{i,\text{MKT}}$, and $\beta_{i,\Delta\text{VIX}}$, which represent the asset's sensitivity to changes in the market and VIX after controlling for changes in the VIX and market, respectively.[2] As a result, although the variance assets have large negative capital asset pricing model (CAPM) betas due to the negative correlation between variance changes and market returns, in the Ang et al. (2006) framework they have relatively small market betas. Reassuringly, all 18 test assets have positive $\beta_{i,\Delta\text{VIX}}$, which indicates they are exposed to variance risk. Moreover, those with larger $\beta_{i,\Delta\text{VIX}}$ tend to have more negative Sharpe ratios, indicating variance risk is an important factor in explaining the unconditional risk premia of these assets. Finally, all 18 test assets have negative daily and monthly α_i in the Ang et al. model.

The biggest takeaway from Table 2 is that although these 18 test assets are all positively exposed to variance risk, they have different exposure to variance

[2]The intercept α is not an excess return because the change in VIX is not a traded asset. However, assuming VIX is stationary, an estimate of α in equation (4) is an unbiased estimate of the CAPM alpha.

risk at different maturities. Several results in Table 2 support this takeaway. The first is that the relations between $\beta_{i,\Delta VIX}$ and both risk premia and Sharpe ratios are nonlinear, suggesting a linear single factor model is insufficient. The second is, as described previously, that longer maturities have smaller absolute Sharpe ratios, indicating long-term and short-term variance risk are priced differently. The third is, even holding maturity fixed, that the three different types of assets have different Sharpe ratios, indicating their risk exposures are not identical. The fourth is, as presented in Panel D, that although the test assets are all positively correlated

TABLE 2
Summary Statistics for Returns of Variance Assets

Table 2 presents summary statistics for daily returns of 18 variance assets in excess of the risk-free rate. The first group, presented in Panel A, are option portfolios that replicate variance swaps at 6 maturities. The second group, presented in Panel B, are 6 constant-maturity VIX futures strategies. The third group, presented in Panel C, are 6 constant-maturity at-the-money Standard & Poor's (S&P) 500 straddle strategies. For each asset, I compute α_i, $\beta_{i,MKT}$, and $\beta_{i,\Delta VIX}$, the coefficients in a time-series regression of excess asset returns on contemporaneous excess market returns and changes in the VIX:

$$r_{i,t} - r_{f,t} = \alpha_i + \beta_{i,MKT}\left(r_{MKT,t} - r_{f,t}\right) + \beta_{i,\Delta VIX}\left(VIX_{1,t} - VIX_{1,t-1}\right) + \epsilon_{i,t}.$$

Panel D presents the correlation matrix for daily excess returns of the 18 variance assets. The sample contains 4,445 daily observations from 1996 through 2013 for variance swaps and straddles, and 2,375 daily observations from 2004 through 2013 for VIX futures.

	Maturity (months)					
Statistics	1	2	3	6	9	12
Panel A. Excess Synthetic S&P 500 Variance Swap Returns						
Daily Returns						
Mean	−1.36%	−0.64%	−0.35%	−0.20%	−0.12%	−0.08%
Standard dev.	15.08%	9.85%	7.80%	5.28%	4.36%	4.19%
Sharpe ratio (ann.)	−1.44	−1.03	−0.71	−0.61	−0.45	−0.29
Skewness	4.18	3.64	2.66	2.23	1.42	1.39
α_i	−1.40%	−0.63%	−0.33%	−0.19%	−0.11%	−0.06%
$\beta_{i,MKT}$	0.99	−0.46	−0.86	−0.54	−0.57	−0.58
$\beta_{i,\Delta VIX}$	7.26	4.40	3.10	2.19	1.61	1.22
Monthly Returns						
Mean	−18.31%	−9.52%	−5.56%	−3.76%	−2.40%	−1.37%
Standard dev.	126.52%	73.56%	49.55%	27.98%	21.82%	21.25%
Sharpe ratio (ann.)	−0.50	−0.45	−0.39	−0.47	−0.38	−0.22
Skewness	10.17	7.89	4.99	2.63	1.86	1.83
α_i	−18.96%	−9.49%	−5.22%	−3.63%	−2.11%	−1.14%
$\beta_{i,MKT}$	0.99	−0.15	−0.66	−0.28	−0.54	−0.43
$\beta_{i,\Delta VIX}$	16.65	10.31	6.82	3.80	2.55	2.31
Panel B. Excess VIX Futures Returns						
Daily Returns						
Mean	−0.19%	−0.20%	−0.14%	−0.08%	−0.07%	−0.08%
Standard dev.	3.79%	3.63%	3.32%	3.22%	3.27%	3.73%
Sharpe ratio (ann.)	−0.81	−0.89	−0.67	−0.41	−0.34	−0.34
Skewness	0.97	0.81	0.57	0.49	0.62	0.44
α_i	−0.16%	−0.18%	−0.11%	−0.06%	−0.04%	−0.06%
$\beta_{i,MKT}$	−1.00	−1.00	−0.93	−0.82	−0.93	−0.73
$\beta_{i,\Delta VIX}$	1.00	0.87	0.74	0.71	0.56	0.57
Monthly Returns						
Mean	−3.78%	−4.02%	−2.92%	−1.76%	−1.59%	−2.28%
Standard dev.	20.51%	19.04%	16.08%	15.39%	14.69%	13.55%
Sharpe ratio (ann.)	−0.64	−0.73	−0.63	−0.40	−0.37	−0.58
Skewness	2.90	2.23	1.73	1.79	1.67	0.71
α_i	−2.90%	−3.11%	−2.09%	−0.94%	−0.83%	−1.64%
$\beta_{i,MKT}$	−1.52	−1.58	−1.44	−1.43	−1.32	−1.12
$\beta_{i,\Delta VIX}$	1.78	1.41	0.96	0.78	0.86	0.49

(continued on next page)

TABLE 2 (continued)
Summary Statistics for Returns of Variance Assets

Statistics	Maturity (months) 1	2	3	6	9	12
Panel C. Excess S&P 500 Straddle Returns						
Daily Returns						
Mean	−0.32%	−0.18%	−0.10%	−0.03%	−0.02%	0.00%
Standard dev.	5.86%	3.41%	2.63%	1.75%	1.44%	1.31%
Sharpe ratio (ann.)	−0.88	−0.84	−0.62	−0.30	−0.17	0.00
Skewness	3.65	3.35	2.95	1.14	1.43	1.08
α_i	−0.39%	−0.22%	−0.14%	−0.06%	−0.04%	−0.02%
$\beta_{i,\text{MKT}}$	2.33	1.45	1.13	0.85	0.74	0.67
$\beta_{i,\Delta\text{VIX}}$	3.14	1.94	1.48	0.96	0.76	0.64
Monthly Returns						
Mean	−5.29%	−3.10%	−1.70%	−0.42%	−0.13%	0.13%
Standard dev.	35.74%	21.03%	16.35%	11.45%	9.27%	8.13%
Sharpe ratio (ann.)	−0.51	−0.51	−0.36	−0.13	−0.05	0.05
Skewness	3.40	2.64	2.24	1.54	1.26	0.99
α_i	−6.52%	−4.06%	−2.52%	−1.08%	−0.72%	−0.40%
$\beta_{i,\text{MKT}}$	2.11	1.65	1.42	1.15	1.03	0.92
$\beta_{i,\Delta\text{VIX}}$	5.51	3.56	2.83	1.94	1.56	1.27

Panel D. Correlations Among Daily Returns of Variance Assets

	Maturity	Variance Swaps 1	3	6	12	VIX Futures 1	3	6	S&P 500 Straddles 1	3	6	12
Variance Swaps	1	1.00	0.85	0.85	0.62	0.76	0.69	0.43	0.85	0.77	0.63	0.46
	3		1.00	0.92	0.76	0.87	0.80	0.52	0.66	0.73	0.64	0.48
	6			1.00	0.79	0.86	0.81	0.55	0.64	0.71	0.66	0.54
	12				1.00	0.79	0.75	0.51	0.42	0.53	0.54	0.51
VIX Futures	1					1.00	0.92	0.59	0.53	0.63	0.58	0.50
	3						1.00	0.56	0.46	0.58	0.55	0.48
	6							1.00	0.26	0.35	0.35	0.31
S&P 500 Straddles	1								1.00	0.83	0.70	0.54
	3									1.00	0.86	0.71
	6										1.00	0.79
	12											1.00

with each other, their correlations are mostly between 50% and 80%, indicating they are exposed to different risk factors. Together, these results make it unlikely that the conditional risk premia of all 18 tests assets are related in the same direction to a single factor in the VIX term structure. However, in Section IV, I show that this is indeed the case.

IV. Empirical Results

A. Expectations Hypothesis

A natural hypothesis is that the shape of the VIX term structure reflects expectations about future changes in return variance and not differences in variance risk premia. For example, this "expectations hypothesis" states an upward-sloping VIX term structure reflects market's expectation that the VIX will increase over the next year rather than higher variance risk premia in longer term options. The expectations hypothesis for the VIX term structure is directly comparable to the expectations hypothesis for bond markets, which states the shape of the Treasury yield curve reflects expectations about future changes in Treasury yields and not differences in bond risk premia.

The motivation for the expectations hypothesis can be seen from manipulating equation (2), which assumes only that the underlying index has no jumps:

$$\begin{aligned}(5)\qquad \mathrm{VIX}^2_{k+m,t} &= \frac{1}{k+m}\mathbb{E}^{\mathbb{Q}}_t\left[\int_t^{t+k+m}\sigma_s^2 ds\right]\\ &= \frac{1}{k+m}\left(\mathbb{E}^{\mathbb{Q}}_t\left[\int_t^{t+k}\sigma_s^2 ds+\int_{t+k}^{t+k+m}\sigma_s^2 ds\right]\right)\\ &= \frac{k}{k+m}\mathrm{VIX}^2_{k,t}+\frac{m}{k+m}\mathbb{E}^{\mathbb{Q}}_t\left(\mathrm{VIX}^2_{m,t+k}\right),\end{aligned}$$

where m is a VIX maturity and k is a forecast horizon. In this case, equation (5) says that the current long-term $(k+m)$ VIX^2 is a weighted average of the current short-term (k) VIX^2 and risk-neutral expected short-term (m) VIX^2 k periods into the future. I rearrange equation (5) to find the market's risk-neutral expected future VIX^2:

$$(6)\qquad \mathbb{E}^{\mathbb{Q}}_t\left(\mathrm{VIX}^2_{m,t+k}\right) = \mathrm{VIX}^2_{k+m,t}+\frac{k}{m}\left(\mathrm{VIX}^2_{k+m,t}-\mathrm{VIX}^2_{k,t}\right).$$

The expectations hypothesis takes equation (6) a step further by assuming the risk premium $\mathbb{E}^{\mathbb{P}}_t\left(\mathrm{VIX}^2_{m,t+k}\right)-\mathbb{E}^{\mathbb{Q}}_t\left(\mathrm{VIX}^2_{m,t+k}\right)$ is constant and equal to a. This implies the current shape of the term structure reflects statistical-measure expectations about future changes in the VIX. Specifically, substituting into equation (6), the expectations hypothesis implies:

$$\begin{aligned}(7)\qquad \mathbb{E}^{\mathbb{P}}_t\left(\mathrm{VIX}^2_{m,t+k}\right) &= a+\mathbb{E}^{\mathbb{Q}}_t\left(\mathrm{VIX}^2_{m,t+k}\right)\\ &= a+\mathrm{VIX}^2_{k+m,t}+\frac{k}{m}\left(\mathrm{VIX}^2_{k+m,t}-\mathrm{VIX}^2_{k,t}\right).\end{aligned}$$

I test equation (7) using regressions of the form:

$$(8)\qquad \mathrm{VIX}^2_{m,t+k}-\mathrm{VIX}^2_{k+m,t} = a+b\left(\mathbb{E}^{\text{EXP\_HYP}}_t\left(\mathrm{VIX}^2_{m,t+k}\right)-\mathrm{VIX}^2_{k+m,t}\right)+\epsilon_{m,t+k},$$

$$(9)\qquad \mathrm{VIX}^2_{m,t+k}-\mathrm{VIX}^2_{m,t} = a+b\left(\mathbb{E}^{\text{EXP\_HYP}}_t\left(\mathrm{VIX}^2_{m,t+k}\right)-\mathrm{VIX}^2_{m,t}\right)+\epsilon_{m,t+k},$$

$$(10)\qquad \mathbb{E}^{\text{EXP\_HYP}}_t\left(\mathrm{VIX}^2_{m,t+k}\right) \equiv a+\mathrm{VIX}^2_{k+m,t}+\frac{k}{m}\left(\mathrm{VIX}^2_{k+m,t}-\mathrm{VIX}^2_{k,t}\right),$$

where the expectations hypothesis predicts $b=1$. The first specification (8) tests the expectation hypothesis prediction that VIX "decays" $\mathrm{VIX}^2_{m,t+k}-\mathrm{VIX}^2_{k+m,t}$, whereas the second specification (8) tests the prediction that VIX "changes" $\mathrm{VIX}^2_{m,t+k}-\mathrm{VIX}^2_{m,t}$.

To estimate specifications (8) and (9), I require that m, k, and $k+m$ are maturities contained in the VIX term structure (1, 2, 3, 6, 9, and 12 months). I therefore test the expectations hypothesis using maturities $m=1$ and $m=2$ for forecast horizon $k=1$, and maturities $m=3$, $m=6$, and $m=9$ for forecast horizon $k=3$. This yields 10 tests of the expectations hypothesis, 5 for predicting VIX decays and 5 for predicting VIX changes.

Table 3 presents the results of these 10 tests of the expectations hypothesis. In every case $\hat{b}<1$, and I strongly reject the expectations hypothesis null

TABLE 3
Expectations Hypothesis for the VIX Term Structure

Table 3 presents tests of the expectations hypothesis for the VIX term structure, which states:

$$\mathbb{E}_t^{\text{EXP\_HYP}}\left(\text{VIX}^2_{m,t+k}\right) = a + \text{VIX}^2_{k+m,t} + \frac{k}{m}\left(\text{VIX}^2_{k+m,t} - \text{VIX}^2_{k,t}\right),$$

where VIX_T, t^2 is an annualized modelfree estimate of the option-implied variance for the Standard & Poor's (S&P) 500 index T months into the future measured on day t. Specifically, in Panel A, I test the implications of the expectations hypothesis for predicting the decay in VIX using regressions of the form:

$$\text{VIX}^2_{m,t+k} - \text{VIX}^2_{m+k,t} = a + b \times \left(\mathbb{E}_t^{\text{EXP\_HYP}}\left(\text{VIX}^2_{m,t+k}\right) - \text{VIX}^2_{m+k,t}\right) + c \times \text{VIX}^2_{m,t} + \epsilon_{m,t+k}$$

for a variety of k and m. In Panel B, I test the implications of the expectations hypothesis for predicting changes in VIX using regressions of the form:

$$\text{VIX}^2_{m,t+k} - \text{VIX}^2_{m,t} = a + b \times \left(\mathbb{E}_t^{\text{EXP\_HYP}}\left(\text{VIX}^2_{m,t+k}\right) - \text{VIX}^2_{m,t}\right) + c \times \text{VIX}^2_{m,t} + \epsilon_{m,t+k}.$$

Standard errors are in parentheses, computed using Newey and West (1987) with lags equal to 1.5 times the number of overlapping days. I also present p-values for the expectations hypothesis null $b = 1$.

	$k=1$: Predicting Next-Month Decay				$k=3$: Predicting Next-Quarter Decay					
	$m=1$		$m=2$		$m=3$		$m=6$		$m=9$	
Panel A. Predicting VIX Decay ($VIX^2_{m,t+k} - VIX^2_{m+k,t}$)										
$\hat{b}$	−0.299	−1.263	−0.562	−1.430	0.088	−1.251	0.319	−0.745	−0.017	−0.947
	(0.442)	(0.416)	(0.577)	(0.468)	(0.341)	(0.327)	(0.342)	(0.263)	(0.380)	(0.295)
$\hat{c}$	—	−0.285	—	−0.250	—	−0.508	—	−0.394	—	−0.382
	—	(0.080)	—	(0.061)	—	(0.066)	—	(0.058)	—	(0.049)
Exp. hyp. p-value	0.3%	0.0%	0.7%	0.0%	0.8%	0.0%	4.6%	0.0%	0.7%	0.0%
R^2	0.6%	10.5%	2.3%	12.5%	0.0%	16.0%	0.7%	14.2%	0.0%	16.0%
Panel B. Predicting VIX Change ($VIX^2_{m,t+k} - VIX^2_{m,t}$)										
$\hat{b}$	0.350	−0.132	0.232	−0.250	0.544	−0.125	0.650	0.064	0.418	0.001
	(0.221)	(0.208)	(0.341)	(0.265)	(0.171)	(0.164)	(0.170)	(0.135)	(0.198)	(0.188)
$\hat{c}$	—	−0.285	—	−0.252	—	−0.508	—	−0.387	—	−0.354
	—	(0.080)	—	(0.057)	—	(0.066)	—	(0.053)	—	(0.047)
Exp. hyp. p-value	0.3%	0.0%	2.4%	0.0%	0.8%	0.0%	4.0%	0.0%	0.3%	0.0%
R^2	3.4%	13.0%	1.3%	11.1%	8.9%	23.4%	8.5%	20.3%	3.7%	17.9%

($b = 1$), echoing the conclusion Mixon (2007) reaches using the term structure of Black–Scholes (1973) implied volatilities. Given the size of the variance risk premium, the failure of the expectations hypothesis in other settings, and the evidence in other papers of variance asset return predictability, the failure of the expectations hypothesis for the VIX term structure is not surprising. But what is surprising is that there appears to be little or no relation between future VIX movements and movements predicted by the expectations hypothesis. All 5 of the decay-predicting tests in Panel A have negative and insignificant $\hat{b}$. The VIX change regressions in Panel B have positive $\hat{b}$, though only the next-quarter predictions are statistically different from 0.

Furthermore, the positive values of $\hat{b}$ in Panel B of Table 3 are due entirely to mean reversion in VIX^2_m. If the expectations hypothesis is correct, the optimal forecast of mean reversion should be captured perfectly using the shape of the VIX term structure, leaving no room to incrementally predict VIX changes using the current VIX. However, as I show in Table 3, when I add the current $\text{VIX}^2_{m,t}$ to the right-hand side of equations (8) and (9), not only does the current VIX load negatively, reflecting mean reversion, but it drowns out any predictability afforded

by the expectations hypothesis term. This implies the shape of the term structure predicts next-quarter VIX changes because it is a noisy proxy for expected mean reversion, which is more precisely measured by VIX_m^2 alone. After controlling for risk aversion, all 10 of the $\hat{b}$ in Table 3 are either negative or positive but insignificant.

Put more broadly, Table 3 shows that, contrary most models and intuition, the VIX term structure does not reliably increase (decrease) after the VIX term structure is upward (downward) sloping. To the extent it does, the VIX term structure contains no information other than the simple mean reversion already captured by the current VIX.

B. Single-Factor Tests

Given the failure of the expectations hypothesis, it must be that time variation in the shape of the VIX term structure is driven by changes in variance risk premia embedded in options used to compute VIX. In this section, I show that the variations in variance risk premia across different maturities are driven almost entirely by different exposures to variations in a single factor: the second PC of the term structure (SLOPE).

My approach and conclusions mimic those in Cochrane and Piazzesi (2005), who find that a single factor summarizes nearly all information about bond risk premia in the Treasury term structure, though with two key differences. The first is that the Cochrane–Piazzesi factor is tent-shaped, whereas SLOPE is monotonic. The second, and more important, difference is that Cochrane and Piazzesi estimate their single factor using a first-stage regression of average returns across all test assets on the full term structure, effectively choosing the best possible single factor. By contrast, I use the second PC as the single factor because it predicts variance asset returns better than do the other PCs. My informal factor selection approach is more conservative because it limits the scope of potential linear combinations to the factors resulting from principal components analysis and is not designed to compute the optimal single factor.

As discussed in Section III, I use returns for 18 variance assets as proxies for variance risk premia. Given the evidence in Table 2, it seems likely that each of these 18 assets has different loadings on multiple factors in the VIX term structure, resulting in conditional expected returns of the form:

$$\mathbb{E}_t(r_{i,t+1}) - r_{f,t+1} = a_i + \mathbf{VIX}_t^2 \times \boldsymbol{\gamma}_i, \tag{11}$$

where $\mathbf{VIX}_t^2$ is a vector of the 6 $VIX_{m,t}^2$ and $\boldsymbol{\gamma}_i$ is a vector of the 6 loadings for asset i on the 6 VIX. By rotating $\mathbf{VIX}_t^2$ using the PC definitions $\boldsymbol{\Gamma}$ given in Table 1, equation (11) can be rewritten as:

$$\mathbb{E}_t(r_{i,t+1}) - r_{f,t+1} = a_i + \mathbf{PC}_t \times \boldsymbol{\lambda}_i, \tag{12}$$

where $\mathbf{PC}_t \equiv \mathbf{VIX}_t^2 \times \boldsymbol{\Gamma}$ is a vector of the 6 PCs of $\mathbf{VIX}_t^2$ and $\boldsymbol{\lambda}_i = \boldsymbol{\Gamma}^{-1}\boldsymbol{\gamma}_i$ is a vector of the 6 loadings for asset i on the 6 VIX.

A much more restrictive asset pricing model is that all information about the risk premia of these 18 test assets in the VIX term structure can be summarized by the second PC, $SLOPE_t$. This hypothesis implies expected returns take the form:

$$\mathbb{E}_t(r_{i,t+1}) - r_{f,t+1} = a_i + b_i \mathrm{SLOPE}_t. \tag{13}$$

Note that in equation (13), all time-series variation in variance risk premia come from variations in SLOPE_t, whereas all cross-sectional differences in variance risk premia come from the constant factor loading b_i and intercept a_i.

I test whether the restricted model (13) holds empirically using the fact that model (13) is equivalent to model (12) when the factor loadings $\lambda_{i,j}$ are 0 except for SLOPE $j=2$. I therefore estimate the unrestricted model using regressions of the form:

$$r_{i,t+1} - r_{f,t+1} = a_i + \mathbf{PC}_t \times \boldsymbol{\lambda}_i + \epsilon_{i,t+1}, \tag{14}$$

and test the single-factor hypothesis using a χ^2 test for the hypothesis that λ_i are jointly 0 for all factors except SLOPE.[3]

The results in Table 4 largely support the single-factor hypothesis for both next-day and next-month variance asset returns. With 3 test assets, each with 6 maturities, and 2 forecast horizons, Table 4 presents 36 variations of my test of the single-factor hypothesis. In all 36 cases, SLOPE negatively predicts variance asset returns, 33 of which are statistically significant. No other PC predicts variance asset returns with nearly such consistency. As discussed further below, the failure of the LEVEL factor is particularly surprising because most models predict variance risk premia should be large during high-variance times.

I reject the single-factor hypothesis using the χ^2 test in only 8 of the 36 cases. Even in these cases, however, SLOPE delivers most of the predictability. These rejections are mostly due to the fifth PC, PC5, which significantly predicts returns in 9 cases. However, I discount the importance of this predictability for four reasons. The first is the inherent difficulty in interpreting PC5, which explains only 0.12% of the total variance in the VIX term structure. The second is that the relation between PC and future returns is positive in some cases and negative in others. The third is that Table 4 presents the adjusted R^2 for SLOPE-only regressions as well as the full 6-factor regressions, and even in cases where the single-factor hypothesis is rejected, for example, for next-month S&P straddle returns in Panel F of Table 4, the R^2 afforded by SLOPE alone is almost as large as the unrestricted R^2.

The final reason I discount the predictability offered by other PCs that reject of the single-factor hypothesis is that regressions with SLOPE alone outperform unrestricted regressions in out-of-sample (OOS) tests. For both restricted and unrestricted models, I compute OOS R^2 by estimating an OOS predicted return for each t:

$$\hat{r}_{i,t+1} - r_{f,t+1} = \hat{a}_{i,t} + \mathbf{PC}_t \times \hat{\boldsymbol{\lambda}}_{i,t}, \tag{15}$$

where the coefficients $\hat{a}_{i,t}$ and $\hat{\boldsymbol{\lambda}}_{i,t}$ are estimated using only past and future observations where the left-hand side does not overlap with observation t.[4] Because they use past and future data, these OOS regressions assess how much of the predictive relation is due to small-sample overfitting but do not represent the economic value of the predictor to a real-time investor, an issue I revisit in later in

[3] To make the economic magnitudes of $\boldsymbol{\lambda}_i$ easier to interpret, I scale the PCs to have a standard deviation of 1.

[4] For next-day returns in Panels A–C of Table 4, I use all observations but t. For next-month returns in Panels D–F, I use all observations but $t-20$ through $t+20$.

Figure 2. In 33 of 36 regressions, the OOS R^2 for SLOPE alone is higher than the unrestricted OOS R^2. Even in the three exceptions, the OOS performance is very close.

Taken together, the evidence in Table 4 indicates SLOPE summarizes all economically meaningful information about variance risk premia in the VIX term structure.

TABLE 4
Single-Factor Tests for Conditional Variance Risk Premia

Table 4 presents tests of the single-factor hypothesis that all variance risk premium information in the VIX term structure is contained in the second principal component (PC), SLOPE. For 18 variance assets, I regress future excess returns on the 6 PCs of the VIX term structure, each scaled to have a standard deviation of 1. For each regression, I present two R^2 measures for the 6 PCs combined and for SLOPE alone: adjusted R^2 and an out-of-sample (OOS) R^2 based on fitted values $\hat{r}_{i,t+1}$ estimated using all observations except those overlapping with $r_{i,t+1}$. I test the single-factor null that the coefficients on all PCs except SLOPE are 0, using a χ^2 hypothesis test for their joint significance. Panel A tests the single-factor hypothesis for daily synthetic Standard & Poor's (S&P) 500 variance swap returns, Panel B for daily VIX futures returns, and Panel C for daily at-the-money S&P 500 straddle returns, all net of the risk-free rate. Panels D–F repeat the exercise using overlapping observations of next-month returns. Daily returns are in basis points, monthly returns are in percentages, and standard errors for the coefficients and p-values for the single-factor hypothesis tests are in parentheses. For monthly returns, standard errors are computed using Newey and West (1987) with 32 lags. * and ** indicate significance at the 5% and 1% levels, respectively.

	Maturity (months)					
	1	2	3	6	9	12
Panel A. Predicting Next-Day S&P 500 Variance Swap Returns						
$LEVEL_t$	−15.31	−11.02	−2.76	−1.98	−0.30	−1.14
	(24.87)	(20.11)	(16.12)	(12.18)	(10.02)	(7.98)
$SLOPE_t$	−55.37*	−59.60**	−56.49**	−31.85**	−24.76**	−19.35*
	(26.59)	(20.27)	(16.17)	(11.87)	(9.32)	(7.61)
$CURVE_t$	−23.82	−22.76	−21.41	−3.62	1.12	3.89
	(25.05)	(19.81)	(16.37)	(11.34)	(8.97)	(7.79)
$PC4_t$	11.78	11.41	14.85	11.40	11.59	5.02
	(29.33)	(21.56)	(16.82)	(11.45)	(9.11)	(7.46)
$PC5_t$	13.13	3.12	−4.01	−3.87	−12.25*	−8.95
	(18.99)	(12.12)	(9.42)	(7.14)	(5.11)	(4.61)
$PC6_t$	−6.66	−27.41	−3.54	−8.92	−5.33	−2.57
	(23.29)	(17.17)	(13.54)	(10.71)	(8.91)	(7.46)
Adj. R^2	0.05%	0.39%	0.51%	0.32%	0.35%	0.15%
SLOPE adj. R^2	0.11%	0.34%	0.50%	0.34%	0.30%	0.19%
OOS R^2	−0.19%	0.03%	0.14%	−0.13%	−0.06%	−0.13%
SLOPE OOS R^2	0.03%	0.24%	0.39%	0.22%	0.18%	0.10%
Single-factor χ^2	2.63	5.13	2.98	2.51	9.29	4.75
p-value	(75.7%)	(40.0%)	(70.3%)	(77.5%)	(9.8%)	(44.7%)
Panel B. Predicting Next-Day VIX Futures Returns						
$LEVEL_t$	6.10	8.21	9.29	9.08	11.78	10.24
	(10.22)	(9.36)	(8.72)	(9.06)	(10.68)	(13.85)
$SLOPE_t$	−30.94**	−28.70**	−17.27	−19.93*	−18.35	−9.56
	(10.89)	(9.47)	(8.94)	(9.62)	(11.13)	(11.81)
$CURVE_t$	−9.63	−5.43	2.90	8.28	5.54	24.54
	(11.41)	(10.37)	(9.89)	(10.32)	(11.06)	(12.95)
$PC4_t$	−4.76	−4.65	−3.88	−5.44	−8.18	6.45
	(11.04)	(10.37)	(9.75)	(9.96)	(12.03)	(15.19)
$PC5_t$	−10.25	−8.23	−8.81	−6.33	−1.69	−14.43
	(9.66)	(9.49)	(9.19)	(9.28)	(9.07)	(17.57)
$PC6_t$	−11.32	−12.38	−13.73	−13.48	−7.00	−4.58
	(8.92)	(8.36)	(7.67)	(8.03)	(8.13)	(11.28)
Adj. R^2	0.69%	0.63%	0.36%	0.52%	0.33%	0.52%
SLOPE adj. R^2	0.63%	0.58%	0.23%	0.34%	0.27%	0.02%
OOS R^2	−0.11%	−0.09%	−0.43%	−0.41%	−0.79%	−1.21%
SLOPE OOS R^2	0.41%	0.40%	0.03%	0.11%	−0.01%	−0.24%
Single-factor χ^2	3.18	3.79	6.06	5.71	2.99	6.07
p-value	(67.3%)	(58.0%)	(30.0%)	(33.5%)	(70.2%)	(30.0%)

(continued on next page)

TABLE 4 (continued)
Single-Factor Tests for Conditional Variance Risk Premia

	Maturity (months)					
	1	2	3	6	9	12
Panel C. Predicting Next-Day S&P 500 Straddle Returns						
$LEVEL_t$	−0.51	1.48	2.92	3.74	2.28	2.37
	(9.22)	(5.72)	(4.61)	(3.31)	(2.75)	(2.60)
$SLOPE_t$	−37.19**	−29.56**	−28.45**	−21.63**	−18.09**	−14.88**
	(9.97)	(6.26)	(5.08)	(3.51)	(2.96)	(2.77)
$CURVE_t$	−18.74	−20.08**	−17.82**	−9.81*	−6.54	−3.08
	(10.67)	(7.05)	(5.87)	(4.09)	(3.40)	(3.08)
$PC4_t$	14.34	10.52	7.95	4.52	3.42	3.75
	(10.18)	(6.31)	(5.05)	(3.42)	(2.90)	(2.58)
$PC5_t$	3.31	0.20	−0.89	−2.73	−4.03*	−3.55*
	(6.63)	(4.02)	(2.97)	(2.28)	(1.84)	(1.76)
$PC6_t$	2.67	−2.35	−0.46	−4.16	−1.73	−2.23
	(8.43)	(4.98)	(4.23)	(3.09)	(2.55)	(2.37)
Adj. R^2	0.44%	1.07%	1.60%	1.91%	1.83%	1.43%
SLOPE adj. R^2	0.38%	0.73%	1.14%	1.51%	1.56%	1.27%
OOS R^2	0.22%	0.80%	1.28%	1.55%	1.45%	1.05%
SLOPE OOS R^2	0.30%	0.64%	1.05%	1.41%	1.45%	1.15%
Single-factor χ^2	5.84	12.17*	11.77*	10.65	10.03	7.97
p-value	(32.2%)	(3.2%)	(3.8%)	(5.9%)	(7.4%)	(15.8%)
Panel D. Predicting Next-Month S&P 500 Variance Swap Returns						
$LEVEL_t$	−0.53	−0.84	−0.04	−0.16	0.20	−0.55
	(3.50)	(2.73)	(2.05)	(1.24)	(1.11)	(0.84)
$SLOPE_t$	−12.88**	−12.47**	−11.82**	−8.08**	−6.71**	−4.32**
	(4.47)	(3.35)	(2.71)	(1.54)	(1.31)	(0.95)
$CURVE_t$	−5.77	−3.19	−1.90	−0.86	−0.40	0.19
	(4.83)	(3.30)	(2.30)	(1.28)	(0.97)	(0.96)
$PC4_t$	−4.52	−2.82	−0.54	0.75	0.84	0.42
	(5.57)	(3.91)	(2.69)	(1.14)	(0.95)	(0.82)
$PC5_t$	10.26	5.96*	3.63*	1.79	0.38	0.62
	(5.59)	(2.96)	(1.71)	(0.94)	(0.70)	(0.68)
$PC6_t$	−5.05	−3.23	−1.78	−0.35	0.15	0.60
	(3.79)	(2.46)	(1.59)	(0.85)	(0.75)	(0.78)
Adj. R^2	2.07%	3.96%	6.42%	8.86%	9.61%	4.30%
SLOPE adj. R^2	1.02%	2.86%	5.70%	8.36%	9.49%	4.13%
OOS R^2	−1.96%	−0.99%	1.11%	3.59%	3.45%	0.70%
SLOPE OOS R^2	−0.53%	1.09%	3.56%	6.27%	7.23%	2.78%
Single-factor χ^2	5.39	5.26	5.42	4.51	1.74	2.92
(*p*-value)	(37.0%)	(38.5%)	(36.7%)	(47.9%)	(88.4%)	(71.2%)
Panel E. Predicting Next-Month VIX Futures Returns						
$LEVEL_t$	1.17	1.73	2.15	2.14	1.99	2.39**
	(1.23)	(1.24)	(1.10)	(1.11)	(0.94)	(0.76)
$SLOPE_t$	−6.92**	−7.02**	−5.32**	−6.22**	−5.04**	−3.37**
	(1.52)	(1.34)	(1.19)	(1.10)	(0.94)	(0.88)
$CURVE_t$	−1.27	−0.87	−1.07	−1.32	−0.63	1.53*
	(1.14)	(0.98)	(0.80)	(0.73)	(0.71)	(0.69)
$PC4_t$	−0.56	−0.01	0.43	0.89	−0.45	0.16
	(1.35)	(1.14)	(0.97)	(1.02)	(0.94)	(0.61)
$PC5_t$	−1.47	−1.18	−0.55	−0.54	−0.51	0.17
	(1.23)	(1.13)	(0.87)	(0.91)	(0.98)	(0.79)
$PC6_t$	1.24	1.17	0.63	0.66	0.71	−0.31
	(0.98)	(0.92)	(0.77)	(0.73)	(0.79)	(0.63)
Adj. R^2	12.93%	15.32%	13.40%	19.61%	14.11%	10.49%
SLOPE adj. R^2	11.44%	13.68%	11.00%	16.45%	11.81%	6.20%
OOS R^2	1.19%	3.70%	1.06%	6.25%	4.05%	0.72%
SLOPE OOS R^2	6.52%	8.72%	4.18%	8.82%	6.76%	3.34%
Single-factor χ^2	3.78	5.92	9.68	15.67**	7.07	30.06**
p-value	(58.1%)	(31.4%)	(8.5%)	(0.8%)	(21.5%)	(0.0%)

(continued on next page)

TABLE 4 (continued)
Single-Factor Tests for Conditional Variance Risk Premia

	Maturity (months)					
	1	2	3	6	9	12
Panel F. Predicting Next-Month S&P 500 Straddle Returns						
$LEVEL_t$	−0.90	−0.45	−0.31	−0.08	−0.38	−0.23
	(1.26)	(0.90)	(0.76)	(0.64)	(0.47)	(0.44)
$SLOPE_t$	−6.97**	−4.96**	−4.68**	−3.73**	−3.21**	−2.81**
	(1.47)	(0.89)	(0.73)	(0.56)	(0.46)	(0.42)
$CURVE_t$	−0.70	−0.90	−0.88	−0.88	−0.57	−0.45
	(1.69)	(1.00)	(0.75)	(0.54)	(0.45)	(0.40)
$PC4_t$	0.89	0.65	0.57	0.58	0.59	0.69
	(1.33)	(0.81)	(0.70)	(0.54)	(0.44)	(0.38)
$PC5_t$	3.81**	2.38**	1.84**	1.04**	0.54	0.40
	(1.15)	(0.68)	(0.53)	(0.37)	(0.31)	(0.29)
$PC6_t$	−0.98	−0.70	−0.44	−0.38	−0.04	−0.23
	(1.01)	(0.62)	(0.50)	(0.38)	(0.30)	(0.31)
Adj. R^2	5.07%	7.18%	9.92%	12.34%	13.26%	13.34%
SLOPE adj. R^2	3.80%	5.56%	8.23%	10.63%	12.06%	11.99%
OOS R^2	1.62%	3.45%	5.97%	7.76%	8.97%	8.74%
SLOPE OOS R^2	2.29%	4.01%	6.67%	8.91%	10.39%	10.23%
Single-factor χ^2	16.08**	17.25**	16.79**	13.20*	9.21	9.77
p-value	(0.7%)	(0.4%)	(0.5%)	(2.2%)	(10.1%)	(8.2%)

FIGURE 2
Conditional and Unconditional Straddle Strategy Returns

Figure 2 presents cumulative returns for two trading strategies using constant-maturity Standard & Poor's (S&P) 500 straddle portfolios. The unconditional strategy sells the 30-day straddle portfolio every day, holds for 1 day, and then rebalances to the next 30-day portfolio. The daily returns for selling straddles account for the margin requirement, 20% of the S&P 500 index level. The conditional strategy estimates $SLOPE_t$, the second principal component (PC) of the VIX term structure, using PC analysis on past data. The strategy buys straddles if $SLOPE_t$ is in the bottom quintile of its historical distribution and sells straddles (with the appropriate margin) otherwise. The gray vertical stripes indicate periods during which the conditional strategy buys straddles. Because the conditional strategy is purely out of sample, it requires a training period, and so each strategy's cumulative returns start at $1 in 2000, meaning the sample contains 3,437 daily observations from 2000 through 2013.

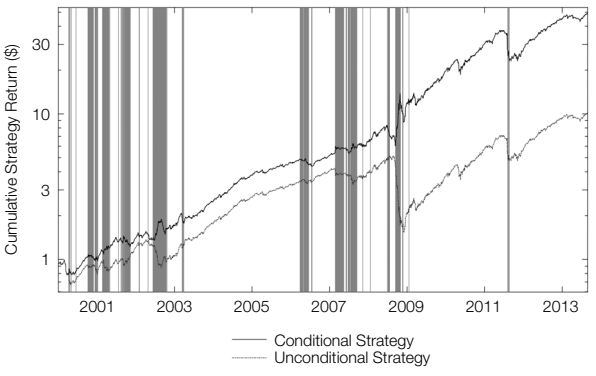

C. Economic Significance of SLOPE as a Predictor

Given the results in Table 4, the remainder of my tests treat SLOPE as a summary of all information in the VIX term structure about variance risk premia and examine the economic significance, robustness, and incremental power of SLOPE as a predictor of variance-asset returns. For brevity, I focus on next-day

variance asset returns, though my results are qualitatively identical for next-month returns.

I assess the economic significance of the predictability afforded by SLOPE using two approaches, the first of which is to estimate the performance of an OOS SLOPE-based trading strategy, as discussed in Section IV.F. The second is to compute average next-day excess variance asset returns across SLOPE quintiles. In addition to measuring economic significance, this approach allows me to assess any nonlinearities in the relation between SLOPE and variance asset returns. Table 5 shows that the difference between high- and low-SLOPE quintiles is statistically significant for all but the 6-month VIX futures returns, and economically enormous: between 29 bps and 181 bps per day.[5] Furthermore, the relation does appear to be nonlinear. In quintiles 2–5, SLOPE and variance asset returns are moderately negatively related. However, the relation becomes much stronger in quintile 1, for which average returns increase dramatically. This pattern suggests there are huge intertemporal changes in the loadings of these assets on variance risk, the price of variance risk, or both.

D. Robustness of SLOPE as a Predictor

A limitation of my analysis is the relatively short 1996–2013 sample period featuring a historic financial crisis. To make my results as convincing as possible given this limitation, I show the predictive relation between SLOPE and variance asset returns is robust to many alternate specifications. The first alternative specification is predicting monthly rather than daily returns. Table 4 shows SLOPE negatively predicts next-month variance asset returns to the exclusion of other factors in the term structure. Table 6 provides an additional robustness check, as well as a measure of economic significance, by examining differences in next-month returns across extreme SLOPE quintiles for 18 variance assets, along with the next-day return differences for comparison. The economic magnitudes are large, ranging from −8.25% to −36.67%, and statistically significant at the 1% level for all 18 test assets.

Another potential concern is that the predictive relation between SLOPE and variance asset returns is driven by extreme observations of both SLOPE and variance asset returns during the financial crisis. Figure 1 shows SLOPE occasionally spikes downward, for example reaching a low almost 10 standard deviations below its mean in 2008. Furthermore, Table 5 shows that variance asset returns following days in the lowest SLOPE quintile are substantially different from those following days in the other 4 quintiles. It is therefore possible that my main results are driven entirely by a few days with extremely low SLOPE and abnormally positive future variance asset returns.

Table 6 alleviates this concern by examining the differences in next-day variance asset returns across SLOPE quintiles in 2 subsamples: one without the financial crisis (defined as all of 2008 and 2009) and one without the bottom 5% of days by SLOPE. In both cases, the magnitude of the predictability afforded by SLOPE

[5] As a benchmark, Table 2 shows average daily returns for these assets are between 0 bps and 136 bps.

TABLE 5

Next-Day Variance Asset Returns across SLOPE Quintiles

Table 5 presents average next-day returns for 18 variance-sensitive investments across 5 subsamples of equal size sorted by SLOPE, the second principal component (PC) of the VIX term structure. Within each subsample, I compute average next-day excess returns for each variance asset and present the difference between the average in the highest quintile and the lowest quintile of SLOPE. Panel A presents average next-day returns for synthetic Standard & Poor's (S&P) 500 variance swaps, Panel B for VIX futures, and Panel C for at-the-money S&P 500 straddles, all net of the risk-free rate. Returns are in percentages, and standard errors are in parentheses. * and ** indicate significance at the 5% and 1% levels, respectively. The sample contains 4,445 daily observations from 1996 through 2013 for variance swaps and straddles, and 2,375 daily observations from 2004 through 2013 for VIX futures.

	Maturity (months)					
SLOPE Quintile	1	2	3	6	9	12
Panel A. Average Next-Day Variance Swap Returns by SLOPE Quintile						
1 (Low)	−0.25	0.46	0.60*	0.44*	0.38**	0.31*
	(0.51)	(0.33)	(0.26)	(0.18)	(0.15)	(0.14)
2	−1.63	−0.83	−0.50	−0.27	−0.22	−0.17
	(0.50)	(0.33)	(0.26)	(0.18)	(0.15)	(0.14)
3	−1.67**	−0.77*	−0.46	−0.25	−0.12	−0.02
	(0.50)	(0.33)	(0.26)	(0.18)	(0.15)	(0.14)
4	−1.30*	−0.70	−0.36	−0.35	−0.20	−0.09
	(0.51)	(0.33)	(0.26)	(0.18)	(0.15)	(0.14)
5 (High)	−1.96**	−1.35**	−1.01**	−0.58**	−0.45**	−0.41**
	(0.51)	(0.33)	(0.26)	(0.18)	(0.15)	(0.14)
High – Low	−1.70*	−1.81**	−1.61**	−1.01**	−0.83**	−0.72**
	(0.72)	(0.47)	(0.37)	(0.25)	(0.21)	(0.20)
Panel B. Average Next-Day VIX Futures Returns by SLOPE Quintile						
1 (Low)	0.21	0.20	0.14	0.15	0.18	0.05
	(0.18)	(0.17)	(0.15)	(0.15)	(0.15)	(0.17)
2	−0.32	−0.34*	−0.31*	−0.15	−0.10	−0.18
	(0.17)	(0.17)	(0.15)	(0.15)	(0.15)	(0.17)
3	−0.13	−0.11	−0.09	0.01	−0.08	0.07
	(0.17)	(0.16)	(0.15)	(0.15)	(0.15)	(0.17)
4	−0.11	−0.14	−0.05	−0.05	−0.03	−0.11
	(0.17)	(0.17)	(0.15)	(0.15)	(0.15)	(0.17)
5 (High)	−0.62**	−0.62**	−0.38*	−0.37*	−0.32*	−0.25
	(0.18)	(0.17)	(0.15)	(0.15)	(0.15)	(0.17)
High – Low	−0.83**	−0.82**	−0.53*	−0.52*	−0.50*	−0.29
	(0.25)	(0.23)	(0.22)	(0.21)	(0.22)	(0.24)
Panel C. Average Next-Day S&P 500 Straddle Returns by SLOPE Quintile						
1 (Low)	0.44*	0.42**	0.47**	0.39**	0.34**	0.29**
	(0.19)	(0.11)	(0.09)	(0.06)	(0.05)	(0.04)
2	−0.38*	−0.23*	−0.18*	−0.11	−0.07	−0.03
	(0.19)	(0.12)	(0.09)	(0.06)	(0.05)	(0.04)
3	−0.61**	−0.35**	−0.21**	−0.14**	−0.10*	−0.09*
	(0.19)	(0.12)	(0.09)	(0.06)	(0.05)	(0.04)
4	−0.43*	−0.30**	−0.24**	−0.13*	−0.11*	−0.06
	(0.19)	(0.11)	(0.09)	(0.06)	(0.05)	(0.04)
5 (High)	−0.65**	−0.44**	−0.35**	−0.21**	−0.16**	−0.11*
	(0.19)	(0.11)	(0.09)	(0.06)	(0.05)	(0.04)
High – Low	−1.09**	−0.86**	−0.81**	−0.59**	−0.50**	−0.40**
	(0.27)	(0.16)	(0.13)	(0.08)	(0.07)	(0.06)

is smaller than in the baseline results but remains statistically and economically significant.

A final concern addressed in Table 6 is that the definition of SLOPE I use for my main tests is based on principal components analysis of the VIX term structure over the entire sample, introducing a possible look-ahead bias. Because principal components analysis is a simple linear rotation of the VIX term structure that does not use return data, any look-ahead bias is likely to be small. Regardless, if my interpretation of the second PC as "slope" is correct, exogenous measures of the term structure's SLOPE should also predict future variance asset returns.

TABLE 6

Robustness of Variance Asset Returns across SLOPE Quintiles

Table 6 presents average future returns for 18 variance-sensitive investments across quintiles of SLOPE, the second principal component (PC) of the VIX term structure. I create 5 subsamples of equal size sorted by SLOPE. Within each subsample, I compute average future excess returns for each variance asset and present the difference between the average in the highest quintile and the lowest quintile of SLOPE. Panel A presents differences in average returns for synthetic Standard & Poor's (S&P) 500 variance swaps across SLOPE quintiles, Panel B for VIX futures, and Panel C for at-the-money S&P 500 straddles, all net of the risk-free rate. For each asset, I present the baseline next-day return differences, next-month return differences, and next-day return differences with the crisis (Jan. 1, 2008–Dec. 31, 2009) removed, the smallest 5% of SLOPE days removed, and SLOPE defined as $VIX^2_{12,t} - VIX^2_{1,t}$. Returns are in percentages, and standard errors are in parentheses. * and ** indicate significance at the 5% and 1% levels, respectively.

	Maturity (months)					
Variation	1	2	3	6	9	12
Panel A. Difference in Variance Swap Returns across Extreme SLOPE Quintiles						
Baseline	−1.70*	−1.81**	−1.61**	−1.01**	−0.83**	−0.72**
	(0.72)	(0.47)	(0.37)	(0.25)	(0.21)	(0.20)
Next-month returns	−36.67*	−36.02**	−32.51**	−21.61**	−17.29**	−12.60**
	(14.36)	(10.55)	(8.16)	(4.50)	(3.81)	(3.55)
Crisis removed	−1.25	−1.44**	−1.33**	−0.87**	−0.82**	−0.77**
	(0.76)	(0.48)	(0.39)	(0.26)	(0.21)	(0.21)
Bottom 5% SLOPE removed	−1.56*	−1.60**	−1.30**	−0.88**	−0.72**	−0.60**
	(0.72)	(0.47)	(0.37)	(0.25)	(0.20)	(0.20)
$SLOPE \equiv VIX^2_{12} - VIX^2_1$	−0.53	−1.08*	−1.18**	−0.71**	−0.60**	−0.58**
	(0.71)	(0.47)	(0.37)	(0.25)	(0.20)	(0.20)
Panel B. Difference in VIX Futures Returns across Extreme SLOPE Quintiles						
Baseline	−0.83**	−0.82**	−0.53*	−0.52*	−0.50*	−0.29
	(0.25)	(0.23)	(0.22)	(0.21)	(0.22)	(0.24)
Next-month returns	−21.10**	−20.90**	−15.25**	−16.63**	−13.74**	−12.07**
	(5.44)	(5.11)	(4.44)	(4.28)	(3.95)	(3.36)
Crisis removed	−0.63*	−0.67*	−0.48*	−0.41	−0.32	−0.36
	(0.28)	(0.27)	(0.24)	(0.23)	(0.22)	(0.23)
Bottom 5% SLOPE removed	−0.61*	−0.60*	−0.36	−0.37	−0.23	−0.15
	(0.24)	(0.23)	(0.22)	(0.21)	(0.21)	(0.23)
$SLOPE \equiv VIX^2_{12} - VIX^2_1$	−0.66**	−0.75**	−0.51*	−0.58**	−0.57**	−0.71**
	(0.24)	(0.23)	(0.21)	(0.21)	(0.21)	(0.24)
Panel C. Difference in VIX Futures Returns across Extreme SLOPE Quintiles						
Baseline	−1.09**	−0.86**	−0.81**	−0.59**	−0.50**	−0.40**
	(0.27)	(0.16)	(0.13)	(0.08)	(0.07)	(0.06)
Next-month returns	−21.13**	−15.40**	−14.10**	−11.43**	−9.67**	−8.25**
	(4.83)	(3.19)	(2.56)	(1.93)	(1.49)	(1.34)
Crisis removed	−0.94**	−0.73**	−0.73**	−0.57**	−0.47**	−0.35**
	(0.29)	(0.17)	(0.13)	(0.09)	(0.07)	(0.07)
Bottom 5% SLOPE removed	−0.92**	−0.70**	−0.66**	−0.48**	−0.40**	−0.32**
	(0.28)	(0.16)	(0.12)	(0.08)	(0.07)	(0.06)
$SLOPE \equiv VIX^2_{12} - VIX^2_1$	−0.57*	−0.51**	−0.52**	−0.33**	−0.29**	−0.25**
	(0.27)	(0.16)	(0.12)	(0.08)	(0.07)	(0.06)

Table 6 shows that sorting the sample by one such alternate measure, $VIX^2_{12} - VIX^2_1$, results in statistically significant but muted predictability.[6]

E. SLOPE as an Incremental Predictor

As discussed in Section II, the related literature documents other timely indicators for variance risk premia. Table 7 shows the ability of SLOPE to predict future variance asset returns is incremental to these other indicators and cannot be explained by mismeasurement or liquidity. As in Table 4, I scale each independent variable to have standard deviation of 1 to make economic magnitudes easier to interpret. I detail each result in turn, but the important conclusion is that SLOPE remains a statistically and economically significant predictor of next-day returns for all but 2 of the assets it predicts in Table 4. Even for the few assets for which SLOPE does not incrementally predict future returns, the coefficient remains negative and economically significant.

The first control variable is the day t variance asset return $r_{i,t}$, which addresses the concern that my results are driven by measurement errors. This concern arises because both SLOPE and time t variance swap and straddle prices are computed from the same potentially noisy option prices, making it possible that measurement error in the option prices drives the predictability I document. This possibility is mitigated by the fact that my results hold when predicting returns on day $t+2$ using SLOPE on day t,[7] and for VIX futures that are traded separately and therefore not subject to the same measurement error as SLOPE. Nevertheless, I address this possibility by controlling for $r_{i,t}$ in Table 7. To the extent measurement errors drive my results, there should be negative autocorrelation in variance asset returns and the predictive power of SLOPE should disappear with after controlling for $r_{i,t}$. Controlling for $r_{i,t}$ also accounts for any compensation liquidity providers receive in the form of short-term reversals in variance asset returns. Table 7 shows that past returns are negative and significant incremental predictors for only 3 of 18 test assets, indicating that measurement error and short-term reversal are not significant for these assets.

The second control variable in Table 7 is $CRASH_{t-20,t}$, an indicator for whether there was a "crash" in the S&P 500 over the 21 trading days ending with t. I define a "crash" as a day with excess market returns in the bottom 1% of my sample, −3.36% or worse. Variance risk premia could be larger after market crashes because the risk aversion of traders increases, because the probability of subsequent crashes increases, or a combination of the two (see Todorov (2010), Aıt-Sahalia et al. (2015)). The results in Table 7 indicate there is no significant relation between $CRASH_{t-20,t}$ and future variance asset returns incremental to the other controls. If anything, future variance asset returns are higher following market crashes, indicating the magnitude of variance risk premia actually decreases.

[6]This exogenous SLOPE definition is likely less effective because it is −68% correlated with the LEVEL of the term structure, which is unrelated to variance asset returns. The negative correlation is due to mean reversion in volatility, which makes the term structure downward sloping when its level is high. A more effective alternative is to sort by $2VIX^2_{12,t} - VIX^2_{1,t}$, which is much less correlated with LEVEL.

[7]Results are available from the author.

The third control variable in Table 7 is $VIX^2_{1,t}$. Although the single-factor tests in Table 4 support the notion that SLOPE summarizes all relevant information in the VIX term structure, they assess only whether the term structure adds predictability to SLOPE without controlling for any other factors. Given thestrong

TABLE 7
Slope as an Incremental Predictor of Variance Asset Returns

Table 7 tests whether SLOPE predicts variance asset returns incrementally to a variety of controls. For each of 18 variance assets, I regress next-day excess returns $r_{i,t+1}$ on 8 potential predictors: $SLOPE_t$, the second principal component of the VIX term structure; $r_{i,t}$, the day t excess return of the variance asset; $CRASH_{t-20,t}$, an indicator for whether excess market returns were in the bottom 1% of the entire sample on 1 of the prior 21 days; $VIX^2_{1,t}$; $VIX^2_{1,t} - \mathbb{E}_t(RV^2_{t+1})$, implied minus expected variance; SP_SKEW_t, the option-implied 30-day skewness of the Standard & Poor's (S&P) 500 index; $NOISE_t$, the illiquidity measure from Hu, Pan, and Wang (2013); and $DEALER\_LEVERAGE_t$ ($\times 10^{-2}$), the most recent ratio of assets to equity for the aggregate broker-dealer sector, as reported in Federal Reserve Flow of Funds, Table L.128. All predictors are scaled to have a standard deviation of 1 except for $CRASH_{t-20,t}$. Panel A tests the incremental predictability of SLOPE for synthetic S&P 500 variance swap returns, Panel B for VIX futures returns, and Panel C for at-the-money S&P 500 straddle returns, all net of the risk-free rate. Returns are in percentages, and standard errors are in parentheses. * and ** indicate significance at the 5% and 1% levels, respectively.

	Maturity (months)					
Predictor	1	2	3	6	9	12
Panel A. Predicting Next-Day S&P 500 Variance Swap Returns						
$SLOPE_t$	−42.87	−58.21*	−55.23**	−28.48*	−22.44*	−24.27*
	(32.82)	(22.91)	(17.39)	(11.80)	(10.37)	(10.41)
$r_{i,t}$	−72.91**	−14.65	−5.44	13.43	6.54	−2.06
	(25.82)	(19.90)	(16.21)	(11.81)	(9.02)	(9.27)
$CRASH_{t-20,t}$	227.58	139.97	79.32	43.77	30.54	24.97
	(141.96)	(94.93)	(71.53)	(47.44)	(37.33)	(35.91)
$VIX^2_{1,t}$	−43.40	−34.68	−4.56	10.56	7.96	−8.87
	(62.82)	(44.48)	(32.66)	(20.21)	(16.03)	(14.49)
$VIX^2_{1,t} - \mathbb{E}_t(RV^2_{t+1})$	−93.25**	−73.08**	−60.06**	−53.70**	−37.22**	−10.71
	(34.76)	(23.01)	(18.92)	(15.05)	(13.22)	(10.52)
SP_SKEW_t	27.36	2.69	−2.09	−9.93	−8.71	−10.12
	(27.50)	(18.04)	(13.80)	(9.55)	(7.78)	(7.26)
$NOISE_t$	24.78	44.05	30.00	15.13	9.55	13.57
	(39.59)	(26.82)	(19.98)	(15.81)	(13.28)	(13.18)
$DEALER\_LEVERAGE_t$	60.59	24.13	8.69	3.18	4.67	−1.78
	(69.14)	(49.64)	(34.46)	(19.31)	(15.88)	(14.85)
Adj. R^2	0.99%	1.10%	1.02%	1.00%	0.76%	0.22%
Panel B. Predicting Next-Day VIX Futures Returns						
$SLOPE_t$	−28.35*	−24.55*	−12.90	−17.00	−20.34	−16.45
	(12.98)	(12.31)	(12.51)	(12.74)	(11.99)	(10.68)
$r_{i,t}$	−3.14	−3.57	−4.74	−10.19	−40.52**	−48.60**
	(9.82)	(9.93)	(9.60)	(9.39)	(10.93)	(15.22)
$CRASH_{t-20,t}$	65.90	51.15	33.26	36.53	30.72	47.41
	(51.10)	(45.99)	(40.05)	(38.41)	(46.00)	(35.55)
$VIX^2_{1,t}$	−0.89	4.65	10.42	15.30	9.83	−8.51
	(24.97)	(22.59)	(17.76)	(17.31)	(18.43)	(22.76)
$VIX^2_{1,t} - \mathbb{E}_t(RV^2_{t+1})$	−35.98*	−38.39*	−38.29**	−38.25**	−15.92	−33.52**
	(17.54)	(16.09)	(11.89)	(8.80)	(9.44)	(10.93)
SP_SKEW_t	−13.76	−10.94	−15.61*	−13.05	−7.97	−9.86
	(8.74)	(8.52)	(7.85)	(7.81)	(7.54)	(7.27)
$NOISE_t$	22.19	23.30	24.71	19.58	15.58	52.85**
	(20.55)	(18.96)	(16.88)	(17.83)	(17.18)	(17.36)
$DEALER\_LEVERAGE_t$	1.62	0.91	2.20	−1.98	0.31	−12.07
	(28.49)	(26.12)	(24.97)	(26.11)	(24.55)	(17.10)
Adj. R^2	1.50%	1.49%	1.43%	1.58%	2.20%	2.87%

(continued on next page)

TABLE 7 (continued)
Slope as an Incremental Predictor of Variance Asset Returns

Predictor	Maturity (months) 1	2	3	6	9	12
Panel C. Predicting Next-Day S&P 500 Straddle Returns						
$SLOPE_t$	−38.88**	−27.78**	−26.21**	−19.05**	−14.81**	−11.26**
	(11.20)	(7.07)	(5.44)	(3.90)	(3.08)	(2.79)
$r_{l,t}$	−11.10	7.04	7.74	9.69*	11.76**	10.05**
	(17.02)	(10.03)	(6.91)	(4.15)	(3.64)	(3.08)
$CRASH_{t-20,t}$	69.74	34.64	24.44	15.74	4.95	0.12
	(48.22)	(30.02)	(23.94)	(16.13)	(13.16)	(11.36)
$VIX^2_{1,t}$	1.20	10.82	12.52	13.53	13.00*	14.68**
	(25.09)	(15.55)	(11.83)	(7.99)	(6.18)	(5.54)
$VIX^2_{1,t} - \mathbb{E}_t(RV^2_{t+1})$	−33.73*	−20.56	−14.30	−9.81	−7.30	−6.85
	(16.43)	(10.66)	(8.34)	(6.02)	(4.93)	(4.64)
SP_SKEW_t	0.56	−1.73	−1.42	−2.87	−2.61	−2.35
	(10.39)	(5.99)	(4.67)	(3.22)	(2.61)	(2.30)
$NOISE_t$	13.45	1.25	−1.37	−4.31	−6.93	−8.86*
	(13.45)	(8.24)	(6.59)	(5.00)	(4.12)	(3.85)
$DEALER\_LEVERAGE_t$	−0.34	−2.72	−4.56	−4.29	−2.16	−0.77
	(19.46)	(11.81)	(8.53)	(5.34)	(4.09)	(3.67)
Adj. R^2	0.68%	0.98%	1.37%	2.07%	2.40%	2.12%

role VIX_1 plays in pricing across many markets, the importance of the level of volatility in most asset pricing models, and the time-varying correlation between LEVEL and SLOPE in Figure 1, it is possible that $VIX^2_{1,t}$ has incremental information after controlling for indicators outside the VIX term structure. Table 7 shows this is not the case, as $VIX^2_{1,t}$ is an insignificant or incrementally positive predictor of variance asset returns.

I also control for the difference between option-implied and expected realized variance, $VIX^2_{1,t} - \mathbb{E}_t(RV^2_{t+1})$, in Table 7. As discussed above, Bollerslev et al. (2009) and Drechsler and Yaron (2011) use implied minus expected variance as an ex ante measure for conditional variance risk premia. I follow Drechsler and Yaron and use the fitted value from a full-sample time-series regression of RV^2_{t+1} on RV^2_t and $VIX^2_{1,t}$ as a proxy for $\mathbb{E}_t\left(RV^2_{t+1}\right)$.[8] This measure uses information outside the VIX term structure, past realized S&P 500 volatility, and therefore could provide information about the variance risk premium that is incremental to SLOPE. Consistent with this hypothesis, I find that $VIX^2_{1,t} - \mathbb{E}_t(RV^2_{t+1})$ negatively predicts future variance asset returns incremental to SLOPE for 11 of 18 test assets. This implies that when option-implied variance is abnormally higher than expected variance, variance risk premia are larger and therefore variance assets have particularly negative abnormal returns.

Kozhan, Neuberger, and Schneider (2013) show that a large part of unconditional variance risk premia is a reflection of a skew risk premia. To rule out the possibility that SLOPE's relation to variance risk premia is due to its being correlated with the conditional skewness of the S&P 500, I add model-free implied skewness as a control variable, estimated as in Kozhan et al. The results in Table 7 indicate conditional S&P 500 skewness does not predict future variance

[8] Also following Drechsler and Yaron (2011), I estimate RV from realized 5-minute S&P 500 futures returns and annualize it to be comparable to VIX^2.

asset returns incremental to the other indicators for most test assets. This alone does not contradict the results in Kozhan et al. because I study conditional, rather than unconditional, variance risk premia.

To rule out the possibility that SLOPE's relation to variance risk premia is due to its being correlated with aggregate liquidity levels, I include the Hu et al. (2013) NOISE measure for illiquidity as a control variable. NOISE, based on the noise in the Treasury yield curve, is more likely to be related to SLOPE than other aggregate liquidity proxies because it is also a daily measure and constructed from a term structure of asset prices. Table 7 shows that NOISE plays a small roll in predicting variance risk premia incremental to the other controls, being slightly positive but statistically insignificant for most assets.

Barras and Malkhozov (2016) argue variance risk premia are larger when intermediaries' risk-bearing capacity is abnormally low. The primary proxy in Barras and Malkhozov for intermediaries' risk-bearing capacity is the aggregate leverage (assets divided by equity) of broker-dealers available quarterly in data from the Federal Reserve Flow of Funds, Table L.128. To assess whether the relation between SLOPE and variance risk premia is driven by intermediaries' risk-bearing capacity, I add a final control to Table 7: DEALER_LEVERAGE$_t$, the most recent leverage ratio of broker-dealers as of time t. I find no significant relation between DEALER_LEVERAGE$_t$ and future variance asset returns incremental to the other controls.

Several untabulated robustness checks are worth mentioning. The single-factor tests in Table 4 yield identical, in some cases stronger, results if tested by regressing next-day variance asset returns on the 6 components of the VIX term structure without first rotating into PCs. The predictability documented in Table 7 is robust to winsorizing SLOPE at the 1% or 5% levels, using quintiles of SLOPE, and using alternate geometric definitions of SLOPE. The results of these analyses are available from the author.

F. OOS Predictive Performance of SLOPE

A natural question is whether SLOPE performs well as an OOS predictor of variance asset returns. Goyal and Welch (2008) argue that OOS performance provides an additional falsifiable prediction of the no-predictability hypothesis and an indicator of the economic value of the predictor to real investors. I assess the OOS performance of SLOPE using both regression and trading strategy approaches. As described previously, Table 4 presents OOS R^2s for both the restricted model that predicts variance asset returns using only SLOPE and the unrestricted model that uses the whole VIX term structure. The OOS R^2s afforded by SLOPE in Table 4 are positive in 33 of 36 cases for the univariate SLOPE regressions. Furthermore, although mechanically smaller than the in-sample R^2, in most cases the two are close, indicating SLOPE's predictability is stable over time and not driven by in-sample overfitting. The OOS R^2 for the unrestricted multivariate model, however, is substantially lower than for both the OOS R^2 offered by SLOPE and the in-sample R^2 in the unrestricted model. The unrestricted model performs poorly OOS because the 6 highly colinear predictors are more prone to overfitting than SLOPE alone.

Although the OOS R^2 in Table 4 provides another rejection of the no-predictability null, a trading strategy approach is better suited to measure the economic value to investors of SLOPE as a predictor. Figure 2 presents cumulative returns for two trading strategies using S&P 500 straddles. The first unconditionally sells S&P 500 straddles with a constant maturity of 30 days by shorting the straddle portfolio I use in my main tests. When shorting straddles, Regulation T requires the proceeds, along with 20% of the index value, be posted as margin. I therefore compute short-straddle returns as:

$$r_{t+1}^{\text{SHORT\_STRADDLE}} = \frac{\text{STRADDLE}_t - \text{STRADDLE}_{t+1}}{0.2 \times \text{SP500}_t}. \tag{16}$$

As documented in Coval and Shumway (2001), average straddle returns are negative, and so the unconditional strategy in Figure 2 results in a substantial positive cumulative return.

The second strategy in Figure 2 is an OOS conditional strategy using SLOPE as an indicator for whether to buy or sell straddles. Specifically, for each day following the 1996–1999 training period, I use the following procedure:

i) Compute SLOPE_t using principal components analysis on VIX term structure data from the beginning of the sample through t.

ii) If SLOPE_t is in the bottom quintile of its historical distribution, buy straddles at t and sell them at $t+1$. Otherwise, short straddles at t and buy them back at $t+1$.

Figure 2 presents the cumulative returns of this conditional strategy, which outperforms the unconditional strategy by a factor of 4.8 over the 14-year window, 11.9% per year.[9] This additional performance comes from days, highlighted in gray in Figure 2, on which the conditional strategy deviates from the unconditional strategy and buys (rather than sells) straddles. A substantial portfolio of the OOS improvement comes from late 2008, when the conditional strategy was long straddles and markets crashed. However, as discussed above, the predictability afforded by SLOPE is robust to removing 2008 and 2009 from the sample.

G. Empirical Patterns for Future Work to Explain

The contribution of this article is to show that a single factor, SLOPE, summarizes all information about variance risk premia in the VIX term structure and is a significant and robust predictor of variance asset returns across all maturities. In doing so, my results provide three puzzling empirical patterns for future work on variance risk premia to explain.

The first puzzling empirical pattern is the insignificant relation between LEVEL and variance risk premia for all maturities. Most models, by contrast, predict that variance risk premia are larger (i.e., more negative) when variance is higher. The second is the direction of the relation between SLOPE and variance risk premia for all maturities. The fact that unconditional variance risk premia are higher for shorter maturities, as documented in Dew-Becker et al. (2017) and

[9] I compute the returns of both strategies using midpoints of the bid–ask spread, which can be substantial in the options market, meaning they overstate returns available to real-time investors.

in Table 2, suggests that variance risk premia disproportionately affect short-term VIX, which would mean that when variance risk premia are large, short-term VIX should be higher than long-term VIX, making the term structure downward sloping. Instead, my results indicate that variance risk premia are large when short-term VIX is low relative to long-term VIX.

A final puzzling empirical pattern is the magnitude of the predictability afforded by SLOPE, as documented in Section IV.C. The magnitude of the predictability is so strong that in the lowest SLOPE quintile, Table 5 shows 17 of the 18 assets have positive abnormal returns, 10 of which are statistically significant. Explaining these results requires intertemporal changes in either the assets' exposure to variance risk or the price of variance risk so substantial they occasionally change the sign of variance risk premia.

V. Conclusion

Changes in the shape of the VIX term structure convey information about time-varying variance risk premia rather than expected changes in VIX, a rejection of the expectations hypothesis. Using daily returns of synthetic S&P variance swaps, VIX futures, and S&P 500 straddles for different maturities, I show that a single factor, SLOPE, summarizes all information in the VIX term structure about variance risk premia. SLOPE predicts returns for all maturities and to the exclusion of the rest of the term structure. The predictability is economically significant, robust, and incremental to other predictors from the literature.

Appendix. Construction of Variance Asset Returns

Carr and Madan (1998) show the VIX^2 index approximates the price of a variance swap traded at time t and maturing at time $t+T$. Because options expiring exactly T months from t are not always traded, the VIX is calculated using a linear interpolation between variance swap rates for the two nearest expiration dates to T. VIX^2 without annualization, an estimate of the variance swap price at time t, is therefore:

$$\text{(A-1)} \qquad \hat{p}_{t,T} = \frac{T-S_1}{S_2-S_1}\sum_K \frac{\Delta K_i}{K^2}\text{OPTION}_t(K;t+S_1) + \frac{S_2-T}{S_2-S_1}\sum_K \frac{\Delta K}{K^2}\text{OPTION}_t(K;t+S_2),$$

where $\text{OPTION}_t(K;t+S)$ is the price at t of the option with strike K and expiration date $t+S$ that is out of the money at time t, ΔK is the difference between K and the nearest strike price, and $S_1 \le T \le S_2$ are the two nearest expiration dates to T.[10]

Note that equation (A-1) is the price of a specific, tradable, portfolio of out-of-the-money options with times to expiration equal to S_1 and S_2. Therefore, the return of a variance swap from day t to day $t+1$ can be approximated by the return of the day t replicating

[10] See https://www.cboe.com/micro/vix/vixwhite.pdf for details on how the set of strikes K is selected.

portfolio:

$$
\begin{aligned}
&r_{T,t+1}^{\text{VAR\_SWAP}} \\
&= \frac{\sum_K \left(\frac{T-S_1}{S_2-S_1} \frac{\Delta K_i}{K^2} \text{OPTION}_{t+1}(K; t+S_1) + \frac{S_2-T}{S_2-S_1} \frac{\Delta K}{K^2} \text{OPTION}_{t+1}(K; t+S_2) \right)}{\sum_K \left(\frac{T-S_1}{S_2-S_1} \frac{\Delta K_i}{K^2} \text{OPTION}_t(K; t+S_1) + \frac{S_2-T}{S_2-S_1} \frac{\Delta K}{K^2} \text{OPTION}_t(K; t+S_2) \right)} - 1 \\
&= \sum_K w_{1,t}(K) \cdot r_{t+1}^{\text{OPTION}}(K; t+S_1) + w_{2,t}(K) \cdot r_{t+1}^{\text{OPTION}}(K; t+S_2), \\
&w_{1,t}(K) \equiv \frac{\frac{T-S_1}{S_2-S_1} \frac{\Delta K}{K^2} \text{OPTION}_t(K; t+S_1)}{\hat{p}_{t,T}}, w_{2,t}(K) \equiv \frac{\frac{S_2-T}{S_2-S_1} \frac{\Delta K}{K^2} \text{OPTION}_t(K; t+S_2)}{\hat{p}_{t,T}},
\end{aligned}
$$

where $\text{OPTION}_t(K; t+S)$ is the price at t of the option with strike K and expiration date $t+S$ that is out of the money at time t, and

$$
r_{t+1}^{\text{OPTION}}(K; t+S) = \frac{\text{OPTION}_{t+1}(K; t+S)}{\text{OPTION}_t(K; t+S)} - 1.
$$

I compute the returns for each constant maturity VIX futures strategy using:

$$
r_{T,t+1}^{\text{VIX\_FUT}} = \frac{\frac{T-S_1}{S_2-S_1} \text{VIX\_FUT}_{t+1}(t+S_1) + \frac{S_2-T}{S_2-S_1} \text{VIX\_FUT}_{t+1}(t+S_2)}{\frac{T-S_1}{S_2-S_1} \text{VIX\_FUT}_t(t+S_1) + \frac{S_2-T}{S_2-S_1} \text{VIX\_FUT}_t(t+S_2)} - 1,
$$

where $\text{VIX\_FUT}_t(t+S)$ is the day t price of a VIX futures contract with maturity date $t+S$, and S_1 and S_2 are the two closest times to maturity to the target time to maturity T.

I compute the returns for each constant maturity S&P straddle strategy using:

$$
r_{T,t+1}^{\text{STRADDLE}} = \frac{\frac{T-S_1}{S_2-S_1} \text{STRADDLE}_{t+1}(t+S_1) + \frac{S_2-T}{S_2-S_1} \text{STRADDLE}_{t+1}(t+S_2)}{\frac{T-S_1}{S_2-S_1} \text{STRADDLE}_t(t+S_1) + \frac{S_2-T}{S_2-S_1} \text{STRADDLE}_t(t+S_2)} - 1, \tag{A-2}
$$

where $\text{STRADDLE}_t(t+S)$ is the day t price of an at-the-money straddle with expiration date $t+S$, and S_1 and S_2 are the two closest times to expiration to the target T.

References

Adrian, T., and J. Rosenberg. "Stock Returns and Volatility: Pricing the Short-Run and Long-Run Components of Market Risk." *Journal of Finance*, 63 (2008), 2997–3030.

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

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

Bakshi, G., and N. Kapadia. "Delta-Hedged Hains and the Negative Market Volatility Risk Premium." *Review of Financial Studies*, 16 (2003), 527–566.

Bakshi, G., and D. Madan. "A Theory of Volatility Spreads." *Management Science*, 52 (2006), 1945–1956.

Bakshi, G.; G. Panayotov; and G. Skoulakis. "Improving the Predictability of Real Economic Activity and Asset Returns with Forward Variances Inferred from Option Portfolios." *Journal of Financial Economics*, 100 (2011), 475–495.

Barras, L., and A. Malkhozov. "Does Variance Risk Have Two Prices? Evidence from the Equity and Option Markets." *Journal of Financial Economics*, 121 (2016), 79–92.

Bekaert, G., and M. Hoerova. "The VIX, the Variance Premium and Stock Market Volatility." *Journal of Econometrics*, 183 (2014), 181–192.

Black, F., and M. Scholes. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy*, 81 (1973), 637–654.
Bollerslev, T.; G. Tauchen; and H. Zhou. "Expected Stock Returns and Variance Risk Premia." *Review of Financial Studies*, 22 (2009), 4463–4492.
Breeden, D. T., and R. H. Litzenberger. "Prices of State-Contingent Claims Implicit in Option Prices." *Journal of Business*, 51 (1978), 621–651.
Broadie, M.; M. Chernov; and M. Johannes. "Understanding Index Option Returns." *Review of Financial Studies*, 22 (2009), 4493–4529.
Campbell, J. Y.; S. Giglio; C. Polk; and R. Turley. "An Intertemporal CAPM with Stochastic Volatility." *Journal of Financial Economics*, forthcoming (2017).
Carr, P., and D. Madan. "Towards a Theory of Volatility Trading." In *Volatility*, Vol. I, R. Jarrow, ed. Ann Arbor, MI: Risk Books (1998), 417–427.
Carr, P., and L. Wu. "Variance Risk Premiums." *Review of Financial Studies*, 22 (2009), 1311–1341.
Christoffersen, P.; S. Heston; and K. Jacobs. "The Shape and Term Structure of the Index Option Smirk: Why Multifactor Stochastic Volatility Models Work So Well." *Management Science*, 55 (2009), 1914–1932.
Christoffersen, P.; K. Jacobs; C. Ornthanalai; and Y. Wang. "Option Valuation with Long-Run and Short-Run Volatility Components." *Journal of Financial Economics*, 90 (2008), 272–297.
Cochrane, J. H., and M. Piazzesi. "Bond Risk Premia." *American Economic Review*, 95 (2005), 138–160.
Corradi, V.; W. Distaso; and A. Mele. "Macroeconomic Determinants of Stock Volatility and Volatility Premiums." *Journal of Monetary Economics*, 60 (2013), 203–220.
Coval, J. D., and T. Shumway. "Expected Option Returns." *Journal of Finance*, 56 (2001), 983–1009.
Dew-Becker, I.; S. Giglio; A. Le; and M. Rodriguez. "The Price of Variance Risk." *Journal of Financial Economics*, 123 (2017), 225–250.
Drechsler, I., and A. Yaron. "What's Vol Got to Do with It." *Review of Financial Studies*, 24 (2011), 1–45.
Egloff, D.; M. Leippold; and L. Wu. "The Term Structure of Variance Swap Rates and Optimal Variance Swap Investments." *Journal of Financial and Quantitative Analysis*, 45 (2010), 1279–1310.
Eraker, B., and Y. Wu. "Explaining the Negative Returns to VIX Futures and ETNs: An Equilibrium Approach." Working Paper, University of Wisconsin (2014).
Feunou, B.; J.-S. Fontaine; A. Taamouti; and R. Tédongap. "Risk Premium, Variance Premium, and the Maturity Structure of Uncertainty." *Review of Finance*, 18 (2014), 219–269.
Filipović, D.; E. Gourier; and L. Mancini. "Quadratic Variance Swap Models." *Journal of Financial Economics*, 119 (2016), 44–68.
Garleanu, N.; L. H. Pedersen; and A. M. Poteshman. "Demand-Based Option Pricing." *Review of Financial Studies*, 22 (2009), 4259–4299.
Goyal, A., and I. Welch. "A Comprehensive Look at the Empirical Performance of Equity Premium Prediction." *Review of Financial Studies*, 21 (2008), 1455–1508.
Heston, S. L. "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options." *Review of Financial Studies*, 6 (1993), 327–343.
Hu, G. X.; J. Pan; and J. Wang. "Noise as Information for Illiquidity." *Journal of Finance*, 68 (2013), 2341–2382.
Kozhan, R.; A. Neuberger; and P. Schneider. "The Skew Risk Premium in the Equity Index Market." *Review of Financial Studies*, 26 (2013), 2174–2203.
Martin, I. "Simple Variance Swaps." Working Paper, London School of Economics (2013).
Mencía, J., and E. Sentana. "Valuation of VIX Derivatives." *Journal of Financial Economics*, 108 (2013), 367–391.
Merton, R. C. "An Intertemporal Capital Asset Pricing Model." *Econometrica*, 41 (1973), 867–887.
Mixon, S. "The Implied Volatility Term Structure of Stock Index Options." *Journal of Empirical Finance*, 14 (2007), 333–354.
Neuberger, A. "The Log Contract." *Journal of Portfolio Management*, 20 (1994), 74–80.
Newey, W. K., and K. D. West. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica*, 55 (1987), 703–708.
Todorov, V. "Variance Risk-Premium Dynamics: The Role of Jumps." *Review of Financial Studies*, 23 (2010), 345–383.