Federal Reserve Bank of New York
Staff Reports

# Global Variance Term Premia and Intermediary Risk Appetite

Peter Van Tassel
Erik Vogt

Staff Report No. 789
August 2016

**Global Variance Term Premia and Intermediary Risk Appetite**
Peter Van Tassel and Erik Vogt
*Federal Reserve Bank of New York Staff Reports*, no. 789
August 2016
JEL classification: G12, G13, C58

**Abstract**

Sellers of variance swaps earn time-varying risk premia for their exposure to realized variance, the level of variance swap rates, and the slope of the variance swap curve. To measure risk premia, we estimate a dynamic term structure model that decomposes variance swap rates into expected variances and term premia. Empirically, we document a strong global factor structure in variance term premia across the U.S., U.K., Europe, and Japan. We further show that variance term premia are negatively correlated with the risk appetite of hedge funds, broker-dealers, and mutual funds. Our results support the hypothesis that financial intermediaries are marginal investors in the variance swap market.

Key words: variance swap, variance risk premium, term structure, empirical asset pricing, volatility, financial intermediaries

---

Van Tassel, Vogt: Federal Reserve Bank of New York (e-mail: peter.vantassel@ny.frb.org, erik.vogt@ny.frb.org). The authors thank Tobias Adrian, Tim Bollerslev, Nina Boyarchenko, and Richard Crump for helpful comments as well as seminar participants at Duke University and the Federal Reserve Bank of New York. 

# 1 Introduction

The financial crisis of 2008 saw some of the highest levels of aggregate stock market volatility in recorded history. During that time, the cost of insuring additional shocks to volatility, as measured by the price of a one-month variance swap, rose to an all-time high. By comparison, the cost of insuring shocks to future volatility, as indicated by long-dated variance swap forwards, was much lower. Modern asset pricing theory predicts that the difference in the valuation of these claims can be explained by declining investor expectations of future volatility or a downward sloping term structure of variance risk premia. As a result, variance swaps contain valuable information regarding investor expectations about the resolution of uncertainty as well as insights into investor preferences for risk across different horizons. The ability to extract and disentangle these components is therefore of vital importance for understanding the expected quantity and pricing of risk.

This paper proposes a new affine term structure model that decomposes the variance swap curve into expected variances and risk premia. In contrast to recent approaches in the variance swap literature, our model involves no latent stochastic volatility factors, which translates into substantial gains in terms of transparency and tractability. These gains arise in large part from allowing realized variance to play the central role as the primitive state variable in our model's economy. Our approach therefore shifts attention from modeling unobserved stochastic volatility to modeling realized variance directly. This is an important distinction as realized variance is observable in our setting from its definition as the payoff of the floating leg of a variance swap. As a result, our model is remarkably fast to implement, requires only linear regression for estimation, and does not necessitate the filtering of latent volatility factors. Moreover, by adhering to theoretical no-arbitrage restrictions, our model accurately prices variance swaps and significantly predicts variance swap returns. In what follows, we exploit this tractability and accuracy to make three empirical contributions.

The first empirical contribution of the paper is to show that a three-factor model explains variance swap prices and returns. The first factor is realized variance, which arises naturally in our setting due to no-arbitrage restrictions. The other factors are the level and slope of the variance swap curve, which correspond to the first two principal components of the variance swap term structure. Similar to their roles in fixed income, the level factor captures shocks that affect all swap rates equally while the slope factor captures shocks that asymmetrically affect the short-end versus the long-end of the curve. We find that not only are each of these factors critical for accurate variance swap pricing, but also that they govern the conditional properties of variance risk premia across the term structure. We refer to these risk premia as variance term premia in analogy to the fixed income literature. Interestingly, we find that variance term premia are volatile and that each of our factors contributes significantly to their time variation. In particular, realized variance and the level of variance swap rates play offsetting roles: increases in realized variance have a negative effect on variance term premia that is strongest at the short end of the curve. In contrast, increases in the level of variance swap rates have a positive effect on term premia that is roughly equal across the curve. The slope factor contributes additionally by governing relative term premia at the short and long end of the curve. Taken together, these factors combine to imply a rich set of dynamics for variance term premia that accurately forecast variance swap returns.

The second empirical contribution of the paper is to document a strong global factor structure in variance swap rates and term premia. We find that most of the variation in international variance swaps is driven by a single factor. In particular, a principal component analysis reveals that a global level factor, which is extracted from all international variance swap curves simultaneously, accounts for more than 87 percent of the variation in international variance swap rates. Further, we find that the level of variance swap rates in each of the international markets loads heavily on the global level factor. This leads us to investigate whether the price of risk for the level of variance swap rates is driven purely by the global component. To that end, we decompose the level factors for each region into a global component and an orthogonal local component. We then estimate a four-factor model that includes both components and find that it is predominantly global risk, and not idiosyncratic local risk, that is priced into variance swap returns. In addition to these results, we also document a strong correlation between our estimates of international variance term premia.

The global term structure results suggest that variance swap markets are integrated in the sense that there is an economic agent who is marginal in setting the level of global variance swap rates. To investigate this hypothesis further, we explore empirically whether financial intermediaries are the marginal investors in the variance swap market. Given the observation that banks and hedge funds created the variance swap market to hedge and speculate on volatility (Carr and Lee (2009)), it seems natural that banks and other types of financial intermediaries would be the important participants in the variance swap market. Moreover, intermediaries with global trading operations are likely to be active in multiple volatility markets, which is consistent with our findings of a strong factor structure in global variance swap markets and the previous findings of the literature (Garleanu et al. (2009)).

The third empirical contribution of the paper is therefore an investigation into whether empirical proxies for financial intermediary risk appetite can account for the time-variation in the level and slope of variance swap rates, as well as variance term premia. Our investigation is motivated by intermediary asset pricing theories which build on the hypothesis that the stochastic discount factor that prices assets is driven by the marginal value of wealth for a representative financial intermediary, rather than a household (He and Krishnamurthy (2013), Adrian et al. (2014), Brunnermeier and Sannikov (2014)). These models predict that risk premia are driven by the time varying effective risk aversion of financial intermediaries. To test this prediction, we regress variance term premia from our term structure model on three distinct measures that proxy for financial intermediary risk appetite including hedge fund drawdowns, broker-dealer leverage, and mutual fund redemptions. In support of our hypothesis, we show that all three measures of risk appetite comove significantly with variance term premia. In addition, we find that broker-dealer leverage varies with short-term premia, whereas hedge fund drawdown and mutual fund redemptions are strongly related to long-term premia.

These results shed new light on the economics of variance swaps. While Bollerslev et al. (2009), Drechsler and Yaron (2011), and Drechsler (2013) evaluate consumption-based asset pricing models for their ability to match the size of the variance risk premium and its predictability of equity returns, Andries et al. (2015) and Dew-Becker et al. (2016) argue that these approaches are inconsistent with the stylized facts pertaining to the term structure of variance swap returns. Intuitively, an investor with Epstein-Zin preferences wants to hedge shocks to future consumption growth that will have a negative impact on lifetime utility.

This can be accomplished by trading long dated variance swaps and results in a Sharpe ratio for variance swaps that is relatively constant across maturities in benchmark long-run risk and disaster models. Notably, this contrasts the data where the Sharpe ratio for receiving fixed in variance swaps is largest at a one-month horizon and is then strongly downward sloping. As an alternative approach, our results point to intermediary asset pricing theories as complementary avenues for further exploration. In particular, our results support the hypothesis that financial intermediaries are marginal investors in the variance swap market: as measures of funding constraints tighten, intermediary risk appetite declines, and the required return to selling variance increases. Furthermore, the relationship between variance term premia at the short end of the curve with broker-dealer leverage and variance term premia at the long end of the curve with hedge fund drawdowns and mutual fund redemptions is indicative of market segmentation or a preferred habitat arrangement in the variance swap market.

The remainder of the paper proceeds as follows. Section 2 provides a literature review. Section 3 builds the variance swap term structure model. Section 4 describes the data and estimation approach. Section 5 presents the main empirical results including our estimates of global variance term premia. Section 6 extends these results by relating variance term premia to measures of financial intermediary risk appetite. Section 7 concludes and the Appendix reports our robustness checks.

# 2 Literature Review

Our paper joins a small but rapidly growing literature on variance swap term structures. Many of the existing papers feature two factor stochastic volatility models that govern variance swap dynamics at the short end and long end of the curve. This setup can be motivated by Adrian and Rosenberg (2008) who find that stock market volatility has a short-run and long-run component that capture market skewness and business cycle risk respectively. In particular, Egloff et al. (2010) and Aït-Sahalia et al. (2015) estimate continuous time affine term structure models with two stochastic volatility factors. While Egloff et al. (2010) solve for the optimal portfolio choice in their estimated model, Aït-Sahalia et al. (2015) estimate the joint term structures of variance risk premia and equity risk premia in a model that allows for jumps using a likelihood-based approach that relies on polynomial approximations of the transition density (Aït-Sahalia (2002)). In addition, Amengual (2008) is an early paper that advocates for a term structure model with two stochastic volatility factors from a Bayesian perspective, while Johnson (2016) provides recent evidence that the second principal component of variance swap rates is significant in return predictability regressions. Filipović et al. (2016) present a different perspective by deriving a quadratic variance swap model. Similar to previous studies, they advocate for two stochastic volatility factors and provide an analysis of optimal variance swap trading strategies. In discrete time, Dew-Becker et al. (2016) consider a three factor model that augments realized variance with short term and long term volatility factors, as opposed to using the level and slope of the variance swap curve. Their approach is closest to the maximum likelihood estimation that we perform in the Appendix as a robustness check for our regression based approach. Beyond these term structure models, Amengual and Xiu (2015) use variance swaps to study whether downward jumps in volatility are related to the resolution of policy uncertainty around FOMC meetings. Also, Barras and Malkhozov (2016) document a significant difference in the estimates of the one-month variance risk premium that are obtained from equity data and option data, which they relate to broker-dealer leverage and interpret as evidence of market segmentation. In addition, Chen et al. (2014) relate broker-dealer leverage to the demand for out-of-the money put options and the equity risk premium in a disaster model.

Compared to the literature, our term structure model represents a distinct, yet complementary approach to modeling variance swaps. A unifying characteristic of the existing models is that they take the underlying stock return as primitive and the variance factors as latent. This approach originates from the option pricing literature, where the underlying return is central to the pricing of options and spot volatility is latent. In contrast, we take realized variance as the primitive state variable that is directly observable from the definition of the floating leg of a variance swap. This more closely resembles the recent bond pricing literature, which assumes that the state vector driving the economy may be either directly observed or inferred from a subset of yields that are priced without error (e.g. Adrian et al. (2013), Joslin et al. (2011), Joslin et al. (2014)). A benefit of this approach is that it drastically simplifies pricing, estimation, and inference. For example, our model is the first model of variance swap term structures that adheres to theoretical no-arbitrage restrictions, requires only linear regressions for estimation, and uses observable pricing factors. We exploit this flexibility to make inferences about the global price of variance risk across international variance swap markets, a topic which has not been considered previously in the literature. At the same time, our approach is not without cost. By focusing on realized variance, we do not model the joint term structures of equity risk premia and variance risk premia as in Aït-Sahalia et al. (2015).

Our work relating the pricing of variance risk to financial intermediaries also complements and extends the recent papers by Chen et al. (2014) and Barras and Malkhozov (2016) in several ways. In particular, Chen et al. (2014) investigates how broker-dealer leverage relates to the equity risk premium and implied volatility skewness. In contrast, we study how leverage relates to variance term premia. In addition, while Barras and Malkhozov (2016) focus on the one-month variance risk premium, we show that financial intermediary risk appetite is related to variance swap pricing across the entire term structure. In particular, the link between broker-dealer leverage and the variance risk premia occurs mainly through the slope of the term structure in our analysis. Second, we show that beyond broker-dealer leverage, hedge fund drawdowns and mutual fund redemptions are also related to the level of variance swap rates. We interpret this as evidence of a rich structure of financial intermediaries that are active in the variance swap market. Finally, our finding that tightened funding constraints are associated with higher prices of risk for global level and slope factors is new to the literature, supporting the hypothesis that it is globally active financial intermediaries that are the marginal investors in the variance market.

Our paper also fits into the broader literature on studying term structures outside of the bond market. Recent important contributions include van Binsbergen et al. (2012) and van Binsbergen et al. (2013) who study term structures of equity yields using option prices and dividend swaps. Their findings suggest that the term structure of equity risk premia is downward sloping, with the short-end of the yield curve dominating the total equity risk premium both in magnitude and time-variation. Along similar lines, Dew-Becker et al. (2016) and Andries et al. (2015) study the unconditional term structure of variance risk premia and find that the Sharpe ratio from shorting variance swaps is downward sloping. These results are notable because they pose challenges for the leading consumption based asset pricing models that predict either flat or upward sloping term structures of equity risk premia and relatively flat term structures of variance risk premia. Moreover, they demonstrate how term structures are useful in testing asset pricing theories and they provide motivation for the investigation of alternative hypotheses, such as intermediary based asset pricing theories, for explaining the unconditional and conditional term structure of variance risk premia.

In addition to the consumption based asset pricing literature, the variance risk premium has also been studied extensively in the option pricing literature. For example, see Bakshi et al. (1997), Bates (2000), Chernov and Ghysels (2000), Pan (2002), Eraker et al. (2003), Broadie et al. (2007), and Andersen et al. (2015). Building on the Black and Scholes (1973) model, these papers derive option prices from stock price dynamics. Often, risk premia are then modeled as the difference in parameters between the risk-neutral and physical measures, which are typically assumed to be in the same affine parametric class to allow for closed-form option pricing (Duffie et al. (2000)). To summarize the results on the variance risk premium, although studies obtain the same sign using different time periods and estimation strategies, finding precise estimates has been challenged by the limited use of long histories and long-maturity option prices. As Singleton (2009) suggests in a review of the option pricing literature, analyzing the variance term-structure, in analogy with the fixed income literature, is an alternative way to obtain precise estimates for the variance risk premium. We pursue that approach in this paper.

# 3 Modeling Variance Swaps

## 3.1 Dynamic Asset Pricing Model

In this section, we build a dynamic term structure model to decompose the variance swap curve into expected variances and risk premia. We begin by assuming that the systematic risk in the economy can be summarized by a $K \times 1$ vector of state variables $X_t$ that follow a stationary vector autoregression as motivated by the intertemporal capital asset pricing model of Merton (1973) or the arbitrage pricing theory of Ross (1976),

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

The innovations $v_{t+1}$ are conditionally mean-zero with covariance matrix $\Sigma_v$. In addition, we assume the existence of a stochastic discount factor,

$$\frac{M_{t+1} - E_t[M_{t+1}]}{E_t[M_{t+1}]} = -\lambda_t' \Sigma_v^{-1/2} v_{t+1}, \tag{2}$$

with affine prices of risk,

$$\lambda_t = \Sigma_v^{-1/2} \left(\Lambda_0 + \Lambda_1 X_t\right). \tag{3}$$

This setup results in the following risk-neutral dynamics,

$$X_{t+1} = \mu^{\mathbb{Q}} + \Phi^{\mathbb{Q}} X_t + v_{t+1}^{\mathbb{Q}}, \tag{4}$$

which are linked to the physical dynamics by the prices of risk $\Lambda_0$ and $\Lambda_1$ through $\mu^{\mathbb{Q}} = \mu - \Lambda_0$ and $\Phi^{\mathbb{Q}} = \Phi - \Lambda_1$. It then follows by standard arguments and the absence of arbitrage, $E_t[M_{t+1} R_{n,t+1}] = 0$, that we can express the excess return for an $n$-month variance swap in the conditional beta representation,

$$R_{n,t+1} = \beta_{n,t}' \left(\Lambda_0 + \Lambda_1 X_t\right) + \beta_{n,t}' v_{t+1} + e_{n,t+1}. \tag{5}$$

The return pricing error $e_{n,t+1}$ is conditionally orthogonal to the innovations $v_{t+1}$ in the state variables. Last, we assume that betas are constant $\beta_{n,t} = \beta_n$ for our baseline approach.

## 3.2 Variance Swap Rates

With this setup, we guess and verify that variance swap rates are affine in the economy's state variables,

$$VS_{n,t} = A_n + B_n' X_t. \tag{6}$$

To derive expressions for $A_n$ and $B_n$, note that the seller of an $n$-month variance swap at time $t$ is entitled to the time $(t+n)$ payoff,

$$VS_{n,t} - \sum_{i=1}^{n} RV_{t+i}, \tag{7}$$

multiplied by the notional amount of the swap where $RV_{t+i}$ is realized variance. Following the market convention, we define realized variance $RV_{t+i}$ as the sum of squared daily log returns in month $t+i$. We also assume that interest rates are either deterministic or independent from realized variance over the $n$-month horizon for which we compute variance swap rates. This facilitates the computation of variance swap rates $VS_{n,t}$ which are set at inception to have a price of zero. In particular, the $n$-month variance swap rate is equal to,

$$VS_{n,t} = E_t^{\mathbb{Q}} \left[ \sum_{i=1}^{n} RV_{t+i} \right]. \tag{8}$$

In deriving variance swap rates it is convenient to work with forward variance swaps. To that end, define the forward variance swap curve as,

$$F_{n,t} = E_t^{\mathbb{Q}} \left[ RV_{t+n} \right]. \tag{9}$$

Forward variance swaps decompose the variance swap curve into the monthly contribution of expected realized variance under the risk-neutral measure. Put differently, variance swap rates are the sum of forward variance swap rates,

$$VS_{n,t} = \sum_{i=1}^{n} F_{n,t}. \tag{10}$$

Now, consider the strategy of selling an $n$-month forward variance swap at time $t$ and buying it back at time $(t+1)$, where $n > 1$. The floating legs $RV_{t+n}$ exactly offset in this trade. Moreover, the trade costs zero dollars at initiation. As a result, it defines the excess return,

$$Rx_{n,t+1} = F_{n,t} - F_{n-1,t+1}. \tag{11}$$

With the returns from trading forward variance swaps now defined, we conjecture that forward variance swap rates are affine in the economy's state variables,

$$F_{n,t} = \tilde{A}_n + \tilde{B}_n' X_t. \tag{12}$$

Plugging this guess into the risk-neutral pricing equation reveals that,

$$\begin{aligned} E_t^{\mathbb{Q}}\left[Rx_{n,t+1}\right] &= E_t^{\mathbb{Q}}\left[F_{n,t} - F_{n-1,t+1}\right] \\ &= \tilde{A}_n + \tilde{B}_n' X_t - \tilde{A}_{n-1} - \tilde{B}_{n-1}' \left(\mu^{\mathbb{Q}} + \Phi^{\mathbb{Q}} X_t\right) \\ &= 0. \end{aligned} \tag{13}$$

Since this equation must hold state by state, matching coefficients results in the following system of recursive equations that determine variance swap rates,

$$\begin{aligned} \tilde{A}_n &= \tilde{A}_{n-1} + \tilde{B}_{n-1}' \mu^{\mathbb{Q}} \\ \tilde{B}_n' &= \tilde{B}_{n-1}' \Phi^{\mathbb{Q}}. \end{aligned} \tag{14}$$

In comparison to zero-coupon bond pricing, note that the variance of the innovations does not enter the recursions. This occurs because our stochastic discount factor is affine as opposed to exponentially affine. As a result, conditional heteroskedasticity, which refers to time varying kurtosis of the underlying, does not affect variance swap pricing.

To close the model, we set the initial condition so that $F_{0,t} = RV_t$. This allows the model to price realized variance exactly and ensures that forward variance swap returns coincide with variance swap returns at a one-month horizon as is required by the absence of arbitrage,

$$\begin{aligned} Rx_{1,t+1} &= F_{1,t} - F_{0,t+1} \\ &= VS_{1,t} - RV_{t+1}. \end{aligned} \tag{15}$$

Furthermore, the initial condition makes clear how realized variance plays a central role in variance swap pricing. We must include realized variance as one of the model's state variables to ensure that $F_{0,t} = RV_t$. To accomplish this, we set the first element of the state vector equal to the standardized value of realized variance,

$$
\begin{aligned}
RV_t &= \mu_{RV} + \left[\sigma_{RV}\ \vec{0}\right] \cdot X_t \\
&= \tilde{A}_0 + \tilde{B}_0' X_t \\
&= F_{0,t}.
\end{aligned} \tag{16}
$$

This defines the initial condition $\tilde{A}_0 = \mu_{RV}$ and $\tilde{B}_0 = [\sigma_{RV}\,\vec{0}]$. The linear relationship between variance swaps and forward variance swaps then allows us to derive the variance swap pricing coefficients in (6) which are given by[1]

$$
A_n = \sum_{i=1}^{n} \tilde{A}_i, \qquad B_n = \sum_{i=1}^{n} \tilde{B}_i. \tag{17}
$$

Last, we can relate the excess returns from trading variance swaps to forward variance swap returns and the conditional beta relationship from the previous section. In particular, variance swap excess returns are the sum of forward variance swap excess returns,

$$
R_{n,t+1} = \sum_{i=1}^{n} Rx_{i,t+1}. \tag{18}
$$

This is similar to the relationship between variance swap rates and forward variance swap rates. The implication for the factor loadings is,

$$
\beta_n = \sum_{i=1}^{n} \tilde{\beta}_i, \tag{19}
$$

where $\tilde{\beta}_i$ is the factor loading for forward variance swap returns. Finally, to relate the pricing coefficients to the beta representation, note that $\tilde{\beta}_n = -\tilde{B}_{n-1}$ as these coefficients multiply the state variable innovations in excess returns.[2]

[1] As an alternative derivation, one can also proceed directly by iterating on the vector autoregression to compute the expectation of future realized variance under the risk-neutral measure. This results in explicit formulas for the coefficients $\tilde{A}_n = \delta_1' \sum_{i=1}^{n} (\Phi^{\mathbb{Q}})^{i-1} \mu^{\mathbb{Q}} + \delta_0$ and $\tilde{B}_n = \delta_1' (\Phi^{\mathbb{Q}})^n$ which match the recursions above. Further, this highlights the connection between the solution to the pricing difference equation and the probabilistic evaluation of the risk-neutral expectation, analogous to the Feynman-Kac formula in continuous time settings.

[2] This follows from the observation that excess returns of variance swaps and forward variance swaps are equivalent at a one-month maturity. For maturities greater than one-month, the excess return from selling an $n$-month variance swap rate at time $t$ and buying it back at time $t+1$ is $R_{n,t+1} = VS_{n,t} - VS_{n-1,t+1} - RV_{t+1}$. It follows that $Rx_{n,t+1} = R_{n,t+1} - R_{n-1,t+1}$.

# 4 Data and Estimation

## 4.1 Variance Swap Data

We estimate our model with synthetic variance swap rates that are constructed from S&P 500 index option prices using the OptionMetrics database. In addition, we also obtain synthetic variance swap rates from Bloomberg and over-the-counter variance swap rates from Markit Totem, which are available starting in November 2008 and September 2006 respectively. In comparison to the Bloomberg and Markit data, our synthetic rates are advantageous as they allow for a longer sample period that runs from January 1996 to August 2015.[3]

As a robustness check, we confirm in the appendix that our synthetic rates closely match the other datasets during the latter years in the sample when all of the datasets are available. In addition, we show that the estimation results and qualitative conclusions are similar across the datasets. Finally, for our international application, we use Bloomberg data for variance swaps written on the S&P 500, Nikkei 225, FTSE 100, and STOXX 50, which are broad indexes that track equity markets in the US, Japan, England, and Europe.

Before discussing our estimation approach and results, we provide a brief discussion of our synthetic variance swap construction. In particular, we compute synthetic variance swap rates from the price of a replicating portfolio that takes a static position in a continuum of out-of-the money European option prices (e.g. Demeterfi et al. (1999), Bakshi and Madan (2000), and Britten-Jones and Neuberger (2000)). Carr and Lee (2009) provide a review of this well-known nonparametric approach.[4] We apply their derivation of the variance swap rate as a weighted average of implied variances. In particular, for each day in our sample, we use OptionMetrics implied volatilities from closing mid-quotes of out-of-the money European option prices to estimate a piecewise quadratic implied volatility function following Broadie et al. (2007) with a knot point that is set at the value of the index. We do this for

[3]Synthetic variance swap rates are also of interest as they can be compared to over-the-counter variance swap rates which include additional jump, credit, and liquidity risks. For example, see Martin (2013) and Aït-Sahalia et al. (2015). To avoid the strike truncation and discretization error implicit in the VIX methodology for computing synthetic variance swap rates, we follow the Carr and Wu (2009) approach.

[4]The replicating portfolio is exact when there are no jumps and interest rates are constant. In recent work, Martin (2013) has shown that the assumption of a continuous underlying can be relaxed by using simple returns as opposed to log returns to compute realized variance for the floating leg payoff.

each maturity that expires on the third Friday of the month with at least three out-of-the money call and put options and at least one-hundred-thousand dollars of outstanding vega.[5] We then compute synthetic variance swap rates at the observed maturities by integrating over the implied volatility function, where we extrapolate beyond the observed strikes by appending log-normal tails. Last, we interpolate between the observed maturities which are fixed in calendar time onto a monthly grid from one-month to two-years for our empirical application.[6]

Figure 1 provides an example of this procedure for the S&P 500 Index on July 20th, 2015. As the top plot illustrates, implied volatility fitting errors are small while the range of observed moneyness is large. This indicates that the estimated implied volatility function provides an accurate description of the risk-neutral distribution. The bottom plot reports the resulting term-structure of synthetic variance swap rates quoted in annualized volatility. The term-structure is upward sloping in this example indicating that expected forward variance is higher than expected short-dated variance under the risk-neutral measure. To illustrate that these results are representative of the broader sample, we report summary statistics for the option prices we use and the implied volatility fitting errors in the Appendix. Throughout the early and latter part of the sample period, the observed range of moneyness is large and the implied volatility fitting errors are small. In addition, our synthetic rates closely match the Bloomberg and Markit data during the latter part of the sample when all of the datasets are available.

To estimate the model, we use monthly variance swap returns. In particular, we compute the excess return from receiving fixed in an $n$-month swap as

$$R_{n,t+1} = VS_{n,t} - VS_{n-1,t+1} - RV_{t+1}. \tag{20}$$

This stems from the following trade. We receive fixed in an $n$-month variance swap at time

[5] In recent years the CBOE has introduced SPX Weeklys, End-of-Month, and PM options. As the liquidity for these contracts may differ from the traditional contracts, particularly when they are initially introduced, we focus our attention on options with the traditional AM settlement on the third Friday of the month.

[6] We use a shape-preserving cubic spline to interpolate the variance swap curve onto a monthly grid. To avoid arbitrage opportunities in our interpolated rates, we check that the associated forward variance swap curve is non-negative. In the rare cases where we observe a negative forward rate (which occurs on less than .50% of days in our sample), we remove the negative forward rate and replace it with a small positive value.

$t$ and then pay fixed in an $n-1$ month swap at time $t+1$. This defines an excess return as no cash exchanges hands at inception and all payoffs are known at time $t+1$. With this definition in hand, we can compute variance swap returns using our synthetic variance swap curves and a measure of realized variance. Following the market convention, we define the realized variance in month $t+i$ as the sum of squared daily log returns.[7]

Table 1 reports summary statistics for our monthly synthetic variance swap rates. We find the variance swap curve is unconditionally upward sloping and positively skewed. In addition, we observe that the curve is more skewed at the front-end, which reflects the fact that the curve tends to invert during periods of market stress. Figure 2 illustrates this by plotting the time-series of the variance swap curve and different examples of the shapes the curve can take on, including upward sloping, inverted, and hump-shaped patterns. During our sample period, the financial crisis coincides with the largest increase in short-dated and long-dated variance swap rates. Last, we can note that the long-end of the curve tends to be less volatile and more autocorrelated than the front-end of the curve.

In addition, Table 1 also reports summary statistics for our monthly excess returns (20) from receiving fixed in variance swaps at various maturities. We can observe that the returns are positive on average, negatively skewed, autocorrelated, and statistically significant at the front-end of the curve. The volatility and CAPM beta of the returns also increase with maturity. This follows mechanically from our definition of excess returns which are not annualized. In particular, the $n$-month variance swap $VS_{n,t}$ exchanges a fixed rate for the amount of realized variance from $t$ to $t+n$. By way of contrast, a one-month swap only exchanges realized variance for one-month. In addition to the increasing volatility, we also see that CAPM alphas decline with maturity both in magnitude and statistical significance. This result is similar to the finding in Dew-Becker et al. (2016) and Andries et al. (2015) who document larger Sharpe ratios for short-dated as opposed to long-dated variance swaps and straddles. Last, we can also note that the returns at the front-end are predominantly positive, with only 16% of one-month returns being negative during our sample period. This result, combined with the negative skewness and positive average returns, motivates

[7] In particular, we assume that each month has 21 trading days which corresponds to 252 trading days per year. To compute returns, we multiply the annualized variance swap rate by $n/12$ to obtain $VS_{n,t}$ and define $RV_{t+1}$ as 21 times the average squared daily log return in month $t+1$.

interpreting the returns from receiving fixed in variance swaps as a form of selling market insurance, which is distinct from but similar to put-writing strategies.

## 4.2 Factor Selection: Initial Evidence

The state variables $X_t$ in our model must explain both the time-variation in expected variance swap returns and the cross-section of variance swap prices. We therefore select factors that meet both of these criteria.[8]

To start, we note that the no-arbitrage recursions in Section 3 require realized variance to appear as a state variable, which determines one of the components of $X_t$. Beyond realized variance, we also use the first two principal components of standardized variance swap rates as state variables. Our motivation for using two principal components is multifold. First, the first two principal components can explain over 99 percent of the cross-sectional variation in variance swap rates. This contrasts fixed income settings where three principal components, which are typically interpreted as level, slope, and curvature factors, are required to explain the cross-sectional variation in the yield curve (e.g. Litterman and Scheinkman (1991)). To illustrate the factor structure in the variance swap market, Figure 3 plots the factor loadings, percentage of variance explained, and the time-series of the first two principal components. As the plot indicates, the first principal component is a level factor that explains over 95 percent of the variation in the curve. The second principal component is a slope factor that explains over 4 percent of the variation in the curve. We also find similar results across subsamples and internationally, which leads us to conclude that two factors are sufficient for pricing the cross-section of variance swap rates.

Of course, the model also requires the state variables to explain the time-series properties of variance swap returns in addition to the cross-section of variance swap rates. In particular,

[8] We can also extend our baseline model and decompose the state vector $X'_{t+1} = [X'_{C,t+1}\ X'_{F,t+1}]$ into cross-sectional variables $X_{C,t+1}$ for pricing and unspanned variables $X_{F,t+1}$ for forecasting. For related applications in the fixed income literature see Collin-Dufresne and Goldstein (2002), Adrian et al. (2013), and Joslin et al. (2014). Unspanned factors are potentially advantageous as they do not increase the number of columns in the beta matrix. In particular, our estimation and inference assumes that $B$ is full column rank. For cases when $B$ is weakly identified see Kleibergen (2009).

recall that expected variance swap returns are equal to,

$$E_t\left[R_{n,t+1}\right] = \beta_n' \left(\Lambda_0 + \Lambda_1 X_t\right). \tag{21}$$

This implies that the state vector should include variables that forecast variance swap returns. To that end, we examine variance swap return predictability for our long US sample in Table 2. As the table indicates, the first two principal components are the strongest predictors of variance swap returns at various horizons. In particular, we run forecasting regressions of the form,

$$\sum_{i=1}^{h} R_{t+i}^{p} = \beta_{0h} + \beta_h' f_t + \epsilon_{t,h}, \tag{22}$$

where $R_{t+i}^{p}$ is an excess return from a variance swap portfolio that is equal-weighted by volatility across maturities and $f_t$ is a forecasting variable.[9] The results indicate that high levels of variance swap rates forecast high returns from receiving fixed in variance swaps. In addition, a more-inverted variance swap curve, which corresponds to a higher value for the second principal component, predicts lower returns. In contrast to these results, the higher order principal components do not provide significant forecasts of variance swap returns across the various horizons considered.

As a robustness check, we also add realized variance and other well-known forecasting variables to the regressions. Panel B includes the US Treasury 10-year minus 3-month yield (Term Slope), Moody's Baa minus Aaa credit spread (Credit), the dividend-to-price ratio (DP), and a three-month moving average of log differences in industrial production (IP Growth). While adding these variables does not change our qualitative conclusions about the level factor, we see that realized variance subsumes the predictive content of the slope factor and that none of the other variables significantly forecast variance swap returns. In addition, we also find that the level factor remains a strong predictor of returns when we use Ibragimov and Müller (2010) as opposed to Newey and West (1987) t-statistics, which are reported in the Appendix, and that the level factor predicts returns during the first half of

[9] We equal-weight by volatility due to our observation in Table 1 that volatility is increasing in maturity for the variance swap excess returns $R_{n,t+1}$ which are not annualized. Similar results hold for equal-weighted returns.

the sample which excludes the financial crisis. In summary, the analysis indicates that the level and slope factors help predict variance swap returns and almost entirely summarize the factor structure in the covariance matrix of variance swap rates. This provides the initial support for our use of realized variance and the level and slope factors as the state variables in our baseline model. In Section 5, we present further evidence in favor of this specification when examining pricing performance in the US and internationally.

### 4.3 Model Estimation

To estimate our model, we apply recently developed methods from the term structure literature. Following Adrian et al. (2013), we assume that return pricing errors, as opposed to variance swap pricing errors, are independent over time. This approach is appealing from an economic and empirical perspective. While there is significant evidence that variance swap pricing errors are autocorrelated at a monthly frequency, there is comparatively less evidence that return pricing errors are predictable.[10] In what follows, we briefly sketch our regression-based approach which closely follows Adrian et al. (2015). The estimated model and empirical results are then discussed in the following sections.

In particular, we assume that the $K \times 1$ state vector $X_t$ is observable and stack our model from the previous section as,

$$\begin{aligned} R &= B\Lambda_0\iota_T' + B\Lambda_1 X_- + BV + E \\ X &= \mu\iota_T' + \Phi X_- + V, \end{aligned} \tag{23}$$

where $R$ is a $N \times T$ matrix of excess returns, $B$ is a $N \times K$ matrix whose rows $\beta_i'$ equal the factor loadings for each maturity, $\iota_T$ is a $T \times 1$ vector of ones, $X = [X_1 \cdots X_T]$ is a $K \times T$ matrix of state variables, $X_- = [X_0 \cdots X_{T-1}]$ is a $K \times T$ matrix of lagged state variables, $V = [v_1 \cdots v_T]$ is a $K \times T$ matrix of state variable innovations, and $E = [e_1 \cdots e_T]$ is a $N \times T$ matrix of return pricing errors.

[10] Similar observations have been made in the option pricing and term structure literatures. For example, see Bliss (1997) and Bates (2000). As Adrian et al. (2013) discuss, assuming independent measurement errors for variance swaps implies return predictability that is not found in the data. In our estimated model, the average autocorrelation across maturities is .40 [5.27] and .23 [2.97] for variance swap prices and -.16 [-2.47] and .06 [.81]. for variance swap returns.

With this setup, we can estimate the model from the following regressions. First, we estimate the physical parameters from a vector autoregression of the observed state variables,

$$X = \hat{\mu}\iota_T' + \hat{\Phi} X_- + \hat{V}. \tag{24}$$

Second, we rewrite the return generating process in the seemingly-unrelated regression $R = AZ + E$ where $Z = [\iota_T\, X_-'\, V']'$ and $A = [A_0\, A_1\, B]$. Replacing the unobserved innovations $V$ with our first-step estimate $\hat{V}$, we can estimate $A$ from the regression,

$$\hat{A} = R\hat{Z}'\left(\hat{Z}\hat{Z}'\right)^{-1} = [\hat{A}_0\, \hat{A}_1\, \hat{B}]. \tag{25}$$

Third, observing that $A_0 = B\Lambda_0$ and $A_1 = B\Lambda_1$, we can estimate the market prices of risk from the regressions,

$$\hat{\Lambda}_0 = \left(\hat{B}'\hat{B}\right)^{-1}\hat{B}'\hat{A}_0, \quad \hat{\Lambda}_1 = \left(\hat{B}'\hat{B}\right)^{-1}\hat{B}'\hat{A}_1. \tag{26}$$

Adrian et al. (2015) provide conditions under which these estimators are consistent and asymptotically normal, and they derive the asymptotic covariance matrix for the parameters taking into account the estimation uncertainty that stems from the multi-step approach. In addition, Adrian et al. (2015) show how to conduct inference when imposing linear restrictions on $B$ and $\Lambda$, and extend this approach to allow for time-varying betas $\beta_{n,t}$.

For our empirical application, a regression-based approach has several advantages. First, it accommodates a potentially large number of factors, which can be financial or real variables, while remaining computationally fast and asymptotically consistent. In addition, we avoid the typical numerical challenges associated with trying to find the global maximum of a likelihood function that has a large number of parameters. However, while (23) does impose significant rank restrictions across the intercept and slope coefficients as well as the state vector evolution, it does not impose all of the restrictions implied by the model. In particular, linear regressions do not impose the nonlinear no-arbitrage recursions on the factor loadings or the internal-consistency constraints on the observed state variables.[11] Naturally,

[11]In particular, our baseline model uses the first two principal components of standardized variance swap rates as observed state variables. Under the null hypothesis that these variables are observed without error,

one might wonder whether these constraints contain useful information or if imposing them would result in efficiency gains.

Empirically, we show that the constraints have little impact on the results. First, one can observe that the constraints are nearly satisfied despite the fact that they are not imposed in estimation. In particular, the factor loadings implied by the no-arbitrage recursions are almost entirely within the 95 percent confidence intervals for the estimated betas. In addition, the model is found to price the principal components used in estimation very accurately. This suggests that imposing the constraints will have little impact on the results.

To confirm this intuition and as a further robustness check, we compute the maximum likelihood estimates which impose the no-arbitrage and internal consistency constraints. We report these results in the appendix. In summary, we obtain similar results using a regression or likelihood-based approach, which leaves our substantive conclusions unchanged. To be specific, we compute the maximum likelihood estimates by applying the insight from Joslin et al. (2011) that the likelihood function factors into the return-pricing errors and a vector-autoregression for the state variables that separately identify the risk-neutral and physical parameters. An implication of Joslin et al. (2011) is that the physical parameters can be estimated from a vector-autoregression of the state variables. As a result, the maximum likelihood estimates for the physical parameters are identical to our first-step estimates. Moreover, this result implies that the pricing restrictions from the no-arbitrage model are not useful in forecasting the evolution of the state vector. For the risk-neutral parameters, the separation of the likelihood function also helps to significantly reduce the computational burden associated with maximizing the likelihood function. Using our regression-based estimates as an initial condition, we find the likelihood function converges rapidly to an estimate of the risk-neutral parameters that is close to our regression-based estimate.

---

the model needs to price them exactly. This allows us to infer the state variables from the observed principal components. As an alternative, one could leave the state variables as latent, which would require us to solve a filtering problem as the state variables would be unobserved.

# 5 Variance Term Premia

This section presents our estimation results and examines the model's main economic content. In particular, while the model is estimated on variance swap excess returns, equation (6) shows that it also prices variance swaps. We therefore compare model-based prices to observed prices as a preliminary gauge of the model and estimator's performance. We then use the model's estimated prices of risk to extract variance term premia. We first focus on the US, and then extend the analysis internationally.

## 5.1 Pricing Performance in the U.S.

The estimation procedure in equations (23)-(26) delivers estimates for the state vector evolution $\mu, \Phi$ and price of risk parameters $\Lambda_0, \Lambda_1$ which can be used to price variance swaps via equation (6). Specifically, by setting $\hat{\mu}^{\mathbb{Q}} = \hat{\mu} - \hat{\Lambda}_0$ and $\hat{\Phi}^{\mathbb{Q}} = \hat{\Phi} - \hat{\Lambda}_1$, we can iterate on the recursions in (14)-(16) to obtain the sequences $\tilde{A}_n$ and $\tilde{B}_n$, which determine forward variance prices. We then convert $\tilde{A}_n$ and $\tilde{B}_n$ to the variance swap coefficients $A_n$ and $B_n$ by equation (17), which are the coefficients needed to form model-based variance swap prices for any maturity $n$.

We show that the three-factor model, when fit to returns, produces variance swap prices with low pricing errors. This result is novel as it indicates that three-step linear regression delivers accurate pricing for variance swaps. To be concrete, Table 3 reports summary statistics for the model's pricing errors. Across maturities, the average pricing errors are close to zero for both variance swap rates and variance swap returns. Moreover, the standard deviation of the variance swap pricing errors are less than 1 percent for the 1-month maturity and .50 percent for longer dated maturities in annualized volatility units. From an economic perspective, we can note that these pricing errors are well within the typical bid-offer spreads in the variance swap market.[12] Figure 4 further illustrates this result by plotting the fitting errors for 1-month, 6-month, and 12-month swaps. As the plot demonstrates, the estimated model captures the dynamics of the variance swap curve throughout the sample.

[12] For example, Dew-Becker et al. (2016) reports that typical bid-offer spreads are 1-2% for maturities up to 1 year and 2-3% for maturities between 1 and 2 years.

## 5.2 Variance Term Premia in the U.S.

Table 4 reports the model's estimated parameters for the full 1996 to 2015 sample using 1, 3, 6, 9, 12, 15, 18, 21 and 24 month variance swap returns. We first note that $\hat{\mu}$ is close to zero. This follows from the fact that we have standardized the state variables under the physical measure. As a result, we can interpret the point estimate $\hat{\Lambda}_0$ as the expected value of the state variables under the risk-neutral measure multiplied by negative one, which is the unconditional risk premium. Thus, we can see from the table that realized variance is unconditionally priced whereas the other state variables are not. This is consistent with the well-documented fact that the unconditional variance risk premium is statistically and economically significant. To interpret the point estimate for realized variance, we define the annualized variance term premium as,

$$VTP_{n,t} = \frac{12}{n}\left(E_t^{\mathbb{Q}}\left[\sum_{i=1}^{n} RV_{t+i}\right] - E_t^{\mathbb{P}}\left[\sum_{i=1}^{n} RV_{t+i}\right]\right). \tag{27}$$

We refer to this object as a variance term premium due to its direct analogy with the term premium in the fixed income literature. To see this, note that for the zero-coupon yield curve, the term premium measures the difference between the yield on a long-maturity bond and the physical expectation of the returns to a sequence of investments that roll over monthly at the short-rate. It measures the compensation that investors earn for tying funds up in a long-maturity bond relative to a position that rolls over a series of short-term investments. By direct analogy, the variance term premium (27) measures the expected holding period return for receiving fixed in a long-maturity variance swap. Note that the one-month term premium is just the familiar variance risk premium studied in the literature (e.g. Carr and Wu (2009)), which is also equivalent to one-month expected variance swap returns.

With this definition, the estimated unconditional one-month variance risk premium in the model is,

$$\hat{E}\left[VTP_{1,t}\right] = -12 \cdot \hat{\Lambda}_{0,1} \cdot \hat{\sigma}_{RV} = .0129. \tag{28}$$

Decomposing the variance risk premium into the variance swap rate and the forecast of realized variance, this amount is equivalent to 21.2% - 18.1% = 3.1% in annualized volatility units. In contrast, the sample average is

$$12\left(\frac{1}{T}\sum_{t=1}^{T}VS_{1,t}-\frac{1}{T}\sum_{t=1}^{T}RV_t\right)=.0145, \tag{29}$$

which is equivalent to 21.5% - 17% = 4.5% in annualized volatility units. The model therefore captures a large and statistically significant unconditional variance risk premium, which is slightly lower than the estimate obtained from the sample averages of the one-month variance swap rate and realized variance.

Of course, the advantage of a dynamic model is that it provides a coherent framework to study the time series properties of the variance term premium and to investigate the variance term premium at different maturities. To that end, we use the model's expression for (27), which is

$$VTP_{n,t}=\frac{12}{n}\left(\left(A_n-A_n^{\mathbb{P}}\right)+\left(B_n'-B_n^{\mathbb{P}\prime}\right)X_t\right). \tag{30}$$

This expression shows that the variance term premium is affine in the state variables $X_t$, with coefficients $A_n$ and $B_n$ coming from the usual recursions on $\mu^{\mathbb{Q}}=\mu-\Lambda_0$ and $\Phi^{\mathbb{Q}}=\Phi-\Lambda_1$ in (14)-(17). By contrast, the physical measure coefficients $A_n^{\mathbb{P}}$ and $B_n^{\mathbb{P}}$ result from using the same recursions, but with prices of risk $\Lambda_0$ and $\Lambda_1$ set to zero. This suggests that $\Lambda_0$ and $\Lambda_1$ critically drive the wedge between objective and risk neutral expectations of realized variance, with $\Lambda_0$ governing the unconditional variance term premium, and $\Lambda_1$ governing the conditional variance term premium.

Turning again to Table 4, we see that the dynamic prices of risk $\hat{\Lambda}_1$ for each of the state variables are strongly statistically significant, which contrasts with the unconditional results that were just discussed. Interestingly, our results suggest that while only realized variance is priced unconditionally, all of the state variables are priced conditionally and therefore represent distinct sources of time variation in variance term premia. Finally, the table also shows that the state variables in the vector autoregression are persistent and significant in forecasting each other from the estimate of $\hat{\Phi}$. This finding has implications for the no-arbitrage recursions (14) and therefore the shape of VTP term structure.

To further interpret the dynamic prices of risk, Figure 5 plots the estimated 1-month variance risk premium and 12-month variance term premium against the realized holding period returns from receiving fixed in variance swaps. The results indicate that there is substantial time variation in the variance term premium, which increases during periods of economic distress. For example, we can see that the LTCM crisis in 1998, the financial crisis in 2008-09, and the European sovereign debt crises in 2010-11 all coincide with increases in the 12-month variance term premium. In addition, the plot indicates that the estimated variance term premium explains realized holding period returns for 1-month and 12-month variance swaps with an $R^2_{adj}$ of 19 and 28 percent respectively. This is particularly notable for the long horizon returns. Despite the fact that we estimate the model with 1-month returns and observe a large forecast error during the financial crisis, the estimated variance term premium still delivers a large $R^2_{adj}$ for 12-month holding period returns.[13]

To analyze the variance term premium further, the plots on the right-hand side of Figure 5 provide a decomposition of the variance term premium into the annualized variance swap rate and the realized variance forecast from the vector-autoregression. The plot reveals that the 1-month variance swap rate and realized variance forecast are more volatile than the corresponding 12-month variance swap rate and forecast. This follows from the mean reversion of realized variance. Moreover, the plot indicates that realized variance mean reverts faster under the physical measure than it does under the risk-neutral measure. To see this, we can observe that shocks to the level of realized variance are followed by realized variance forecasts that quickly mean revert. By way of contrast, the variance swap rate tends to remain elevated after periods of economic distress, particularly at the 12-month horizon. This generates time variation in the variance term premium that has forecast variance swap returns, particularly in the period that has followed the financial crisis.

Beyond the time series dynamics, we can also interpret the estimated prices of risk $\hat{\Lambda}_1$ by plotting the model's comparative statics. In particular, Figure 6 plots the impact of a one-standard deviation increase in the state variables on the annualized variance term premium at different horizons. The results illustrate how the estimated prices of risk translate into

[13]In addition, we also estimated the model using 5-minute intraday returns to compute realized variance. For 5-minute returns, the explanatory power of the variance term premium increases to an $R^2_{adj}$ of 23% and 41% for 1-month and 12-month holding period returns respectively. This is consistent with the high frequency data providing a more precise estimate of realized variance than daily returns. Beyond this improvement, most of our qualitative conclusions are unchanged when using 5-minute or daily log returns, which result in realized variance estimates that are 98% correlated at a monthly frequency. As a result, we report the more conservative baseline results that use daily log returns which follows the variance swap market convention.

time variation in the variance term premium. As the plot indicates, an increase in the level of variance swap rates results in an increase in the variance term premium that is fairly constant across the curve. In contrast, an increase in realized variance results in a decrease in the variance term premium that is strongest at the short-end of the curve. This analysis highlights two important aspects of the estimated model. First, the level of variance swaps plays an offsetting role to realized variance. Second, shocks to realized variance mean revert faster than shocks to the level of variance swap rates, which explains why realized variance has a more pronounced impact at the short-end of the curve. In addition, we can also see that the slope factor differs from realized variance and the level of variance swaps. When slope increases, the variance term premium increases at the short-end of the curve and decreases at the long-end of the curve, albeit with a magnitude that is relatively small compared to the other factors.

To connect these results with our previous work, we can also observe that the comparative statics in the estimated model are consistent with the intuition that emerged from the unconstrained forecasting regressions in Table 2. Compared to the reduced form regressions, the model can price variance swaps and forecast variance swap returns despite the additional structure and reduction in parameters that results from the no-arbitrage restrictions.

## 5.3 Pricing Performance Internationally

Our international analysis relies on Bloomberg synthetic variance swap data that is available from 2009 to 2015. While this is a relatively short sample in comparison to the longer US sample that runs from 1996 to 2015, the Bloomberg data allows us to study variance swap pricing internationally and is available at a daily frequency for analysis at higher frequencies. In addition, the Bloomberg data provides an important robustness check for our previous results. Despite using Bloomberg data during a different sample period that excludes the financial crisis, we continue to find similar results to our full sample analysis. In particular, a three-factor model that is estimated by linear regression using realized volatility, the level of variance swap rates, and the slope of the variance swap curve continues to provide accurate pricing even for international variance swaps.

Table 5 reports summary statistics for the international variance swap data. Similar to the US, we see the international variance swap curves are upward sloping on average and positively skewed. The long-end of the curve is less volatile and more highly autocorrelated than the short-end of the curve. In addition, the returns from receiving fixed in variance swaps are negatively skewed and predominantly positive, in line with the view that earning the variance term premium is akin to selling market insurance. The largest Sharpe ratio is obtained at the 1-month maturity across the different indexes, which is also the maturity where we observe statistically significant and positive CAPM alphas. In terms of differences, we can see that the level of variance swap rates has been higher and that variance swap returns have been more volatile in Europe and Japan relative to the US and UK during the recent recovery. This seems consistent with the European sovereign debt crisis and the Abenomics reforms in Japan having both occurred during the 2009 to 2015 sample period. Also, we can note that receiving fixed in US variance swaps has produced the largest Sharpe ratio. Of course, one must be careful in interpreting the Sharpe ratios over this recent sample which has been a benign period for global financial markets and an attractive time for receiving fixed in variance swaps.

To see that the three-factor model accurately prices international variance swaps, Table 6 reports summary statistics for the model fitting errors for each of the different indexes. Similar to our previous results, the average root mean squared pricing errors are less than 1% in annualized volatility units across the indexes and across maturities. On average, the pricing errors appear smallest for Japan and Europe and are slightly larger for the US and UK. In addition, the pricing errors are also positively autocorrelated at both 1-month and 6-month horizons and exhibit excess kurtosis. On the whole, however, we note that there do not appear to be systematic pricing biases across countries or maturities.

## 5.4 International Variance Term Premia

Table 7 reports estimates for the model parameters that pin down international variance term premia. As with our full US sample, the intercepts $\hat{\mu}$ to the factor evolution equation are relatively small, never exceeding one or two tenths of a standard deviation. This allows us to interpret $\hat{\Lambda}_0$ as approximately equal to the unconditional price of variance risk. We see from Table 7 that unconditional price of realized variance risk is negative and strongly statistically significant across all four regions. Economically, this says that sellers of one-month variance swaps are predominantly compensated for their exposure to realized variance by our definition of the term premium. The size of the $t$-statistics are also noteworthy given the relatively shorter sample. In contrast to our findings on the full US sample, however, we also find that level risk is unconditionally priced, although the effect is statistically slightly weaker than the realized variance price of risk.

Table 7 also reports estimates for the conditional price of risk $\hat{\Lambda}_1$ and shows that all three factors are priced in the UK, Europe, and Japan. We note that among the components of $\hat{\Lambda}_1$ that are precisely estimated, the signs and magnitudes of the estimated prices of risk are remarkably similar across countries, suggesting strong similarities in the way variance risk is priced across markets. The sole exception appears to be US, where slope risk, captured by PC2, no longer appears statistically significant in the shortened sample. We conjecture that this is due to the exclusion of the financial crisis, which represented a significant tail event that resulted in the variance term structure inverting, with short-term swaps priced highest. The uncoupling of time series variation in short-term variance swap rates from long-term variance swap rates is essentially what the slope factor captures. We note, however, that in the $\hat{\Phi}$ estimates in the bottom panel of Table 7, the slope factor nonetheless is helpful in forecasting the level factor and realized variance. This implies that slope impacts the multi-step forecasts of realized variance $E_t^{\mathbb{P}}\left[\sum_{i=1}^{n} RV_{t+i}\right]$, and is therefore helpful for constructing the physical measure leg of the term premium. Finally, we note that also among the components of $\Phi$ that are precisely estimated, magnitudes and signs are similar across countries.

Because of the similarities across countries in estimated prices of risk $\Lambda_0$, $\Lambda_1$ and physical parameters $\mu, \Phi$, equations (6) and (30) along with our no-arbitrage recursions suggest that international swap rates and variance term premia should be similar as well. To gain further insight, Figure 7 plots the variance swap rates and variance term premium at a 1-month and 12-month maturity for the different indexes. From the plot, it is clear that the international variance swap rates and risk premia are indeed highly correlated. Pointwise correlations with the US are reported in the legends. For example, in the UK and Europe, we can see that both short- and long-dated variance swap rates are between 89 to 97 percent correlated with the US. Moreover, the correlation in term premia is similarly striking. The UK one-month variance risk premium is 96 percent correlated with the US, and the 12-month variance term premium is 92 percent correlated with the US. For the Euro STOXX swaps, these correlations are 92 and 95 percent. These results suggest a strong common factor structure in variance term premia across the US, UK, and Europe.

For Japan, there is a high correlation with the other markets, but also evidence of important idiosyncratic events. In particular, we can see that the spike in Japanese variance swap rates in 2011 coincides with the Tōhoku earthquake and tsunami while the increases in 2013 coincide with the Bank of Japan's announcements regarding their extensive quantitative easing and reform program. Across the markets, we can also note that the largest increases in variance swap rates and risk premia are contemporaneous with concerns about the European sovereign debt crisis and with negative stock market returns, which is consistent with the well-known leverage effect.

Taken together, these results indicate that there is a strong factor structure in international variance swap rates and term premia. To investigate this within the context of our model, we compute a global level factor, which is the first principal component across all international variance swap term structures. Panel A of Table 8 reports regressions of the region-specific level factors onto the global factor. The results strongly indicate that the global level factor can explain most of the variation in local level factors. For example, in the US, UK, and Europe, the global level factor explains 97, 98, and 93 percent of the variation in the level of variance swap rates, respectively. For Japan, the global level factor explains 77 percent of the variation. The regression coefficients, furthermore, are strongly statistically significant. In Panel B, we reestimate our international models, but with country-specific level factors split into two components: The first component represents the projection of the local level factor onto the global level factor, using the estimated OLS equations from Panel A. The second component is the residual from each regression and represents variation in the local level factor that is orthogonal to the global level factor. We then reestimate our term structure model in each country including both the global level component and the orthogonalized local level component as separate factors.

The results for $\hat{\Lambda}_1$ indicate that only global level risk is priced, in that it is significant for explaining the time variation in the variance term premium. The idiosyncratic level risk is not priced. In summary, our results suggest that the price of volatility risk is tightly linked to a global level factor that summarizes common risks across the different variance swap markets.

# 6 Intermediary Asset Pricing Interpretation

In this section, we investigate the economic drivers of variance term premia from an intermediary asset pricing perspective. Because of the strong factor structure across international variance swap rates and term premia, it seems natural to hypothesize that there is an economic agent who is marginal in setting the price of global variance swap rates. Given their extensive trading operations and activity in volatility markets, we conjecture that financial intermediaries are likely to be the marginal investors. To that end, this section investigates whether measures of intermediary risk appetite are connected to our variance swap pricing factors and estimates of variance term premia.

## 6.1 Motivating Intermediary Asset Pricing Theories

In their historical account of the volatility derivatives market, Carr and Lee (2009) note that volatility trading has been largely dominated by financial intermediaries. In the nascent days of the variance swap market, hedge funds received fixed in variance swaps to gain exposure to the high levels of implied volatility that followed the 1998 LTCM crisis. Banks were readily available counterparties in these transactions as they could exploit the emerging research on variance swap pricing to hedge their volatility exposure in the options market, thereby earning a cross market bid-ask spread. Since then, institutional demand to hedge and speculate on volatility has been a constant source of trading among financial intermediaries.

To extend these observations to more recent periods, we obtain disaggregated information about volatility exposure by investor type for the VIX futures market from the Commitment of Traders Report that is provided by the Commodity Futures Trading Commission (CFTC). Figure 8 summarizes the results. The top panel shows that open interest in VIX futures has grown rapidly in recent years while the bottom panel shows that trading has been dominated by financial intermediaries. Throughout the sample, broker-dealers and leveraged funds have had the largest net open interest positions. Consistent with Carr and Lee (2009)'s account, we see that broker-dealers have been net long volatility against leveraged funds, such as hedge funds and commodity trading advisors, that have been net short. In addition, we can also see the positions held by institutional asset managers and non-financial traders. The institutional category includes buy-side market participants such as pension funds, endowments, insurance companies, and mutual funds. As the plot indicates, institutional investors have been net long volatility, albeit with smaller positions than broker-dealers. Last, we can see that non-financial traders in the other category have had relatively small positions throughout the sample. Of course, an important caveat for these results is that net open interest does not necessarily reflect overall volatility positions, as it does not capture offsetting positions in related markets. Nonetheless, the figure highlights the important role that hedge funds, broker-dealers, and asset managers play in the market for volatility.

The dominant participation of financial intermediaries in VIX futures suggests that they are likely to be active in the variance swap market as well. This implies that asset pricing theories in which intermediaries are marginal investors should be particularly relevant for explaining variance swap prices and risk premia. In these theories, risk premia are driven by the effective risk aversion of financial intermediaries as opposed to the risk aversion of a representative household. This provides a testable hypothesis: if financial intermediaries are indeed the marginal investors in the variance swap market, then variance swap rates and term premia should be related to measures of intermediary risk appetite. We investigate this hypothesis empirically by relating our pricing factors to hedge fund drawdowns, broker-dealer leverage, and mutual fund redemptions. These measures are drawn from the intermediary asset pricing literature and share the common feature that in periods when intermediaries become financially constrained, whether through the tightening of funding conditions or binding capital constraints, their effective risk aversion increases.

## 6.2 Measuring Intermediary Risk Appetite

**Hedge Funds** In the intermediary asset pricing model of He and Krishnamurthy (2013), households gain exposure to the economy's risky asset by purchasing equity issued by the intermediary, but the amount that they invest is bounded above by a “skin in the game” constraint. If the intermediary is a hedge fund, this constraint means that households may invest only in proportion to the hedge fund manager’s personal wealth in the fund. Thus, if the hedge fund suffers losses, the fund manager’s wealth declines, and the household’s ability to invest in the economy’s risky asset is diminished, which raises risk premia. He and Krishnamurthy (2013) interpret the “skin in the game” constraint as the 20% of returns that are typically paid to the hedge fund manager. In practice, since hedge fund fees are a function of the fund exceeding its previous high-water mark, the equity funding mechanism of He and Krishnamurthy (2013) is likely related to high-water marks in the context of hedge funds.

Panageas and Westerfield (2009) and Drechsler (2014) explicitly study how high-water mark contracts affect the optimal portfolio choice decision of hedge fund managers. In both models, high-water marks induce risk aversion in otherwise risk-neutral hedge fund managers. The key state variables that emerge from this analysis are hedge fund wealth $W_t$ and the associated high-water mark $H_t = \max_{s \leq t} W_s$. In the Drechsler (2014) model in particular, the manager’s value function is homogenous in the high-water mark and the risky asset allocation is a function of drawdown,

$$DD_t = \frac{W_t}{H_t}, \tag{31}$$

which measures the gap between fund manager wealth and the high-water mark that must be obtained to begin earning the performance fee. Thus, the fund manager ends up behaving like a myopic risk-averse investor with time-varying relative risk aversion that is a decreasing function of the drawdown ratio defined above. Motivated by this result, we use drawdown as a measure of hedge fund risk appetite and construct an empirical drawdown measure by assuming that the HFRI Fund-Weighted Composite Index represents the wealth $W_t$ of the hedge fund sector.

**Broker-Dealers** Adrian et al. (2014) provide evidence that broker-dealer leverage is a proxy for financial intermediaries’ marginal value of wealth. As funding conditions tighten, intermediaries are forced to delever. Moreover, because the asset side of broker-dealers' balance sheets consist primarily of risky assets whereas the liability side consists substantially of short-term secured funding, leverage is a measure of risk appetite, with low leverage corresponding to low risk appetite. Similar implications arise in models where intermediaries face risk-based capital constraints (e.g. Danielsson et al. (2011) and Adrian and Boyarchenko (2015)). Intuitively, risk-neutral intermediaries choose leverage in a way that leaves no slack in their risk-based capital constraints. Thus, when the volatility of their risky asset holdings is low, intermediaries choose high leverage, which links leverage to the effective risk aversion of financial intermediaries.

We therefore use broker-dealer leverage as a measure of the risk appetite of broker-dealers. Following Adrian et al. (2014) and Adrian and Boyarchenko (2015), we obtain aggregate leverage information from Table L.129 of the Federal Reserve Flow of Funds.[14] Specifically, we define leverage as the ratio of broker-dealers' total financial assets divided by total equity, where the latter is given by the difference between total assets and total liabilities.

**Asset-Managers** In Vayanos (2004), investors are fund managers who are subject to redemptions when performance falls below a threshold. Because the probability of the threshold being breached increases when aggregate volatility is high, redemption risk is a function of volatility. In equilibrium, asset returns are governed by a two-factor CAPM as a result of fund manager's volatility hedging demand. In our context, because long positions in variance swaps are volatility hedges, a fund manager could mitigate redemption risks by paying the fixed leg on a variance swap to receive protection from increases in volatility. This is also consistent with the net long open interest positions for asset managers in VIX futures plotted in Figure 8.

Thus, because redemption risks motivate variance swap trading, we measure fund manager risk appetite using redemption data from the Investment Company Institute's (ICI) Trends in Mutual Fund Activity. ICI provides monthly statistics on US domiciled mutual funds including their redemptions, which are defined as the dollar value of money returned

[14] Because the rest of our analysis is performed at a monthly frequency, we linearly interpolate the quarterly Flow of Funds figures to obtain a monthly series. In the Appendix, we repeat our analysis with the quarterly data and find virtually identical results.

to investors who have sold shares in the fund. To capture the behavior of funds with risky holdings, we use the redemption for equity funds relative to the total size of equity funds in the database.

### 6.3 Empirical Results

We begin by illustrating the results. Figure 9 plots the level factor (PC1) in blue on the left axis against drawdown $1 - DD_t$ in orange on the right axis for the HFRI Fund-Weighted Composite Index. This relationship is especially strong. As the plot indicates, the two series exhibit significant comovement, which translates into an unconditional correlation of 77% during our full sample period. By comparison, drawdown only has a 55% correlation with realized variance and a -3% correlation with the slope factor. Beyond correlations, the figure also shows that drawdown took its highest value during the financial crisis when it increased by more than 20% relative to the previous high-water mark. During this period, PC1 shows that variance swap rates also reached their highest levels. From an intermediary asset pricing perspective, one interpretation is that the severe drawdown during the crisis increased hedge fund manager risk aversion, which resulted in higher required returns for receiving fixed in variance swaps to induce fund managers to maintain their short positions.

To break out this result for different financial intermediaries, Table 9 shows the results of regressing the level (PC1) and slope (PC2) factors onto the three measures of intermediary risk appetite. Our focus on level and slope stems from the hypothesis that variance swap rates are related to the pricing kernel of financial intermediaries. To interpret the results, we standardize the left and right hand side variables in these regressions to be mean zero and standard deviation one. The top panel of the table uses PC1 and PC2 from our full sample from 1996 to 2015. The bottom panel uses international variance swap rates from 2006 to 2015 that are available from Markit to extract the global PC1 and PC2 factors. We also report $t$-statistics in brackets using Newey-West standard errors with 12 lags to correct for serial correlation and heteroskedasticity, keeping in mind that both the principal components and measures of intermediary risk appetite are persistent when we interpret the statistical significance of our results.

The top panel shows a negative relationship between hedge fund (HF) drawdown and the US level factor, suggesting that as hedge fund wealth declines from its high-water mark, variance swap rates across the curve shift upward. The top right panel shows that drawdown is not related to the slope factor. In contrast, broker-dealer (BD) leverage appears unrelated to level, but strongly related to slope. The sign suggests that as BD leverage increases, so does the steepness of the variance swap curve. Next, the mutual fund (MF) redemption measure is also related to level, but not slope, with increases in redemptions corresponding to upward shifts in variance swap rates. These relationships hold even when all intermediary asset pricing factors are included in the regressions, with a total $R^2$ of 69% for PC1 and 22% for PC2. Further, the bottom panel shows the same patterns between the global principal components and the intermediary asset pricing variables, with hedge fund drawdown and mutual fund redemptions being related to the global level factor, and broker-dealer leverage's relationship to the global slope strengthening relative to the US-based regression. The $R^2$ for the regression including all intermediary asset pricing variables is 81% for PC1 and 24% for PC2.

Beyond the pricing factors, we also conjecture that intermediary risk appetite is related to the price of variance risk as indicated by our estimate of variance term premia. To explore this connection, Table 10 shows regressions of US variance term premia at various horizons onto our intermediary asset pricing factors. The regression uses our full monthly sample from 1996 to 2015 and again reports Newey-West $t$-statistics with 12 lags. The table shows a rich and statistically strong relationship between variance term premia at different horizons and the various intermediary asset pricing factors. Whereas the ability of broker-dealer leverage to explain variance term premia diminishes with the horizon, the ability of hedge fund drawdown and mutual fund redemptions to explain term premia increases by horizon, in terms of both statistical significance as well as coefficient magnitudes. In particular, while a one standard deviation decline in HF drawdown translates into a 0.32 standard deviation increase in the one-month term premium, the same decline in drawdown translates to a larger 0.72 standard deviation increase in the 24-month term premium. As the fourth column indicates, the effects continue to hold when all of the risk appetite variables are included simultaneously.

Broker-dealer leverage has the opposite pattern as hedge fund drawdown, with economic and statistical magnitudes largest for the 1 and 6 month term premium. At the one-month horizon (top left panel), a one-standard deviation increase in broker-dealer leverage translates into roughly the same decline in the one-month risk premium as hedge fund drawdown. Moreover, the sign is intuitive: As leverage increases, the term premium declines, as broker-dealers require less compensation for selling variance swaps when their risk aversion is low. Finally, mutual fund redemptions, like hedge fund drawdowns, are more strongly related to long-horizon term premia, although the magnitudes are smaller both economically and statistically. Here, too, the sign of the relationship is intuitive: when redemptions rise as a fraction of fund size, the fund manager's risk aversion increases, creating a motive to buy long positions in variance swaps. This results in a rise in risk premia for strategies that are net short volatility.

As a robustness check, we repeated the analysis in Tables 9 and 10 on quarterly data to avoid interpolating the broker-dealer leverage measure. The results are virtually identical to the monthly case and are therefore included in the Appendix.

## 6.4 Discussion and Interpretation

The preceding results have important implications for asset pricing theories that seek to explain risk premia in the variance swap market. While the consumption-based asset pricing literature has relied on the variance risk premium to distinguish between equilibrium asset pricing theories, the preceding results point to intermediary asset pricing theories as complementary avenues for further exploration. In particular, our results provide strong indications that hedge fund drawdowns, broker-dealer leverage, and mutual fund redemptions are related to the pricing of risk in the variance swap market. The regressions support the hypothesis that certain financial intermediaries are marginal investors in the variance swap market: as measures of their risk appetite decline, the required return to selling variance (as measured by the term premium) increases. We found this pattern to hold for three different types of intermediaries.

Furthermore, embedded in many intermediary asset pricing models (e.g. He and Krishnamurthy (2013) and Brunnermeier and Sannikov (2014)) is a market segmentation assumption in which households cannot directly participate in the risky asset market in question.

This assumption amounts to a friction that can be interpreted in the context of the hedge fund replication literature (Fung and Hsieh (2001), Mitchell and Pulvino (2001), Jurek and Stafford (2015)). The hedge fund replication literature has shown that hedge funds pursue trading strategies with "option-like" characteristics, in that their returns are often nonlinearly related to benchmark risk factors like the equity market. Thus, this literature has argued that a key reason for the existence of hedge funds is to help investors earn risk premia for bearing exposure to these complicated nonlinear risk factors which they may not be able to earn on their own. The "risky asset" in the intermediary pricing models would therefore represent the collection of dynamic trading strategies that the hedge fund puts on on behalf of households. This participation constraint would explain why hedge funds are marginal investors in the market for variance, and potentially why standard household pricing kernels are unable to generate the Sharpe ratios observed in the variance term structure.

Finally, the broker-dealer leverage connection to the slope factor sheds some light on the debate about the term structure of variance risk premia. In particular, Dew-Becker et al. (2016) have shown that unconditional Sharpe ratios for selling 1-month variance are substantially higher than Sharpe ratios for selling forward variance. Furthermore, our Tables 9 and 10 suggest that broker-dealer leverage's relation to slope helps explain short-run variance risk premia (1-6 months) as opposed to longer-run variance term premia (12-24 months). In contrast, hedge funds and mutual funds appear primarily connected to the pricing of variance risk at long maturities. This finding is suggestive of a preferred habitat arrangement among financial intermediaries in the variance swap market. Taken together, these results imply scope for heterogeneous intermediary asset pricing theories whose investors display horizon-dependent risk aversion.

# 7 Conclusion

We develop a new and tractable affine term structure model of variance swaps. The model explains the cross-section and predictability of variance swap returns. It is also fast to estimate and requires no filtering of latent volatility factors. We apply the model to decompose variance swap curves into the expected realized variance and variance term premium over different horizons, thereby measuring the compensation that investors earn for selling long-dated variance swaps. In addition, we investigate the pricing of variance swaps not just in the US, but also in international markets including the UK, Europe, and Japan.

We report three central empirical findings. First, we show that variance swap returns and prices are well explained by only three factors: realized variance, the level of variance swap rates, and the slope of the variance swap curve. Second, we document a strong factor structure in international variance swap rates and term premia. A global level factor explains over 93 percent of the variation in the level of variance swaps in the US, the UK, and Europe, as well as 77 percent of the variation in the level of Japanese variance swaps. Moreover, our results suggest that it is predominantly global risk, and not idiosyncratic local risk, that is priced into variance swap returns. Last, we show that variance swap rates and term premia are highly correlated with measures of risk appetite for financial intermediaries including hedge funds, broker-dealers, and mutual funds. Moving forward, our results motivate further investigation into whether intermediary asset pricing theories can explain risk premia in global variance swap markets.

# References


Adrian, T. and N. Boyarchenko (2015). Intermediary leverage cycles and financial stability. *Federal Reserve Bank of New York Staff Reports* (567).

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

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

Adrian, T., E. Etula, and T. Muir (2014). Financial intermediaries and the cross-section of asset returns. *The Journal of Finance 69*(6), 2557–2596.

Adrian, T. and J. Rosenberg (2008). Stock returns and volatility: Pricing the short-run and long-run components of market risk. *The Journal of Finance 63*(6), 2997–3030.

Aït-Sahalia, Y. (2002). Maximum likelihood estimation of discretely sampled diffusions: A closed-form approximation approach. *Econometrica 70*(1), 223–262.

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

Amengual, D. (2008). The term structure of variance risk premia. *Working Paper*.

Amengual, D. and D. Xiu (2015). Resolution of policy uncertainty and sudden declines in volatility. *Chicago Booth Working Paper*.

Andersen, T. G., N. Fusari, and V. Todorov (2015). The risk premia embedded in index options. *Journal of Financial Economics 117*(3), 558–584.

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

Bakshi, G., C. Cao, and Z. Chen (1997). Empirical performance of alternative option pricing models. *The Journal of Finance 52*(5), 2003–2049.

Bakshi, G. and D. Madan (2000). Spanning and derivative-security valuation. *Journal of Financial Economics 55*(2), 205–238.

Barras, L. and A. Malkhozov (2016). Does variance risk have two prices? evidence from the equity and option markets. *Journal of Financial Economics 121*(1), 79–92.

Bates, D. (2000). Post-'87 crash fears in the s&p 500 futures option market. *Journal of Econometrics 94*, 181–238.

van Binsbergen, J., M. Brandt, and R. Koijen (2012). On the timing and pricing of dividends. *American Economic Review 102*(4), 1596–1618.

van Binsbergen, J., W. Hueskes, R. Koijen, and E. Vrugt (2013). Equity yields. *Journal of Financial Economics 110*(3), 503–519.

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

Bliss, R. R. (1997). Testing term structure estimation methods. *Advances in Futures and Options Research 9*, 97–231.

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

Britten-Jones, M. and A. Neuberger (2000). Option prices, implied price processes, and stochastic volatility. *The Journal of Finance 55*(2), 839–866.

Broadie, M., M. Chernov, and M. Johannes (2007). Model specification and risk premia: Evidence from futures options. *The Journal of Finance 62*(3), 1453–1490.

Brunnermeier, M. K. and Y. Sannikov (2014). A macroeconomic model with a financial sector. *American Economic Review 104*(2), 379–421.

Carr, P. and R. Lee (2009). Volatility derivatives. *Annu. Rev. Financ. Econ. 1*(1), 319–339.

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

Chen, H., S. Joslin, and S. X. Ni (2014). Demand for crash insurance, intermediary constraints, and stock return predictability. *Working Paper*.

Chernov, M. and E. Ghysels (2000). A study towards a unified approach to the joint estimation of objective and risk neutral measures for the purpose of options valuation. *Journal of Financial Economics 56*(3), 407–458.

Collin-Dufresne, P. and R. S. Goldstein (2002). Do bonds span the fixed income markets? theory and evidence for unspanned stochastic volatility. *The Journal of Finance 57*(4), 1685–1730.

Danielsson, J., H. S. Shin, and J.-P. Zigrand (2011). Balance sheet capacity and endogenous risk. *Financial Markets Group, London School of Economics and Political Science*.

Demeterfi, K., E. Derman, M. Kamal, and J. Zou (1999). A guide to volatility and variance swaps. *The Journal of Derivatives 6*(4), 9–32.

Dew-Becker, I., S. Giglio, A. Le, and M. Rodriguez (2016). The price of variance risk. Technical report, National Bureau of Economic Research.

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

Drechsler, I. (2014). Risk choice under high-water marks. *Review of Financial Studies 27*(7), 2052–2096.

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

Duffie, D., J. Pan, and K. Singleton (2000). Transform analysis and asset pricing for affine jump-diffusions. *Econometrica 68*(6), 1343–1376.

Egloff, D., M. Leippold, and L. Wu (2010). The term structure of variance swap rates and optimal variance swap investments. *Journal of Financial and Quantitative Analysis 45*(5), 1279–1310.

Eraker, B., M. Johannes, and N. Polson (2003). The impact of jumps in volatility and returns. *The Journal of Finance 58*(3), 1269–1300.

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

Fung, W. and D. A. Hsieh (2001). The risk in hedge fund strategies: Theory and evidence from trend followers. *Review of Financial Studies 14*(2), 313–341.

Garleanu, N., L. H. Pedersen, and A. M. Poteshman (2009). Demand-based option pricing. *Review of Financial Studies 22*(10), 4259–4299.

He, Z. and A. Krishnamurthy (2013). Intermediary Asset Pricing. *American Economic Review 103*(2), 732–770.

Ibragimov, R. and U. K. Müller (2010). t-Statistic based correlation and heterogeneity robust inference. *Journal of Business & Economic Statistics 28*(4), 453–468.

Johnson, T. L. (2016). Risk premia and the vix term structure. *Working Paper*.

Joslin, S., M. Priebsch, and K. J. Singleton (2014). Risk premiums in dynamic term structure models with unspanned macro risks. *The Journal of Finance 69*(3), 1197–1233.

Joslin, S., K. J. Singleton, and H. Zhu (2011). A new perspective on gaussian dynamic term structure models. *Review of Financial Studies 24*(3), 926–970.

Jurek, J. W. and E. Stafford (2015). The cost of capital for alternative investments. *The Journal of Finance 70*(5), 2185–2226.

Kleibergen, F. (2009). Tests of risk premia in linear factor models. *Journal of Econometrics 149*(2), 149–173.

Litterman, R. B. and J. Scheinkman (1991). Common factors affecting bond returns. *The Journal of Fixed Income 1*(1), 54–61.

Martin, I. (2013). Simple variance swaps. *NBER Working Paper 16884*.

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

Mitchell, M. and T. Pulvino (2001). Characteristics of risk and return in risk arbitrage. *the Journal of Finance 56*(6), 2135–2175.

Newey, W. and K. West (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica: Journal of the Econometric Society*, 703–708.

Pan, J. (2002). The jump-risk premia implicit in options: Evidence from an integrated time-series study. *Journal of Financial Economics 63*(1), 3–50.

Panageas, S. and M. M. Westerfield (2009). High-water marks: High risk appetites? convex compensation, long horizons, and portfolio choice. *The Journal of Finance 64*(1), 1–36.

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

Singleton, K. J. (2009). *Empirical dynamic asset pricing: model specification and econometric assessment*. Princeton University Press.

Vayanos, D. (2004). Flight to quality, flight to liquidity, and the pricing of risk. Technical report, National Bureau of Economic Research.

**Table 1: Variance Swap Summary Statistics**

This table reports summary statistics for synthetic variance swap rates and returns computed at a monthly frequency from 1996 to 2015. Panel A indicates that the variance swap curve is upward sloping on average and positively skewed. Larger skewness at the front-end reflects the tendency of the variance swap curve to invert during periods of market distress. In addition, the table indicates that the long-end is less volatile and more highly autocorrelated at longer lags. Panel B reports the returns from receiving fixed and paying floating in a synthetic variance swap. The table indicates these returns are positive on average, negatively skewed, autocorrelated, and statistically significant for short-dated maturities. To account for the observed autocorrelation in returns, we report Newey and West (1987) t-statistics using 5 lags.

Panel A: Variance Swap Rates (annualized volatility %)

| Maturity | 1 | 3 | 6 | 9 | 12 | 18 | 24 |
|---|---|---|---|---|---|---|---|
| Mean | 21.54 | 21.89 | 22.46 | 22.67 | 22.89 | 23.32 | 23.57 |
| Std. Deviation | 8.14 | 7.25 | 6.64 | 6.25 | 6.03 | 5.81 | 5.71 |
| Minimum | 10.46 | 11.53 | 11.80 | 12.17 | 12.30 | 12.43 | 12.50 |
| Median | 19.96 | 20.80 | 21.56 | 22.01 | 22.15 | 22.63 | 23.08 |
| Maximum | 61.88 | 53.98 | 50.29 | 47.77 | 46.41 | 43.89 | 42.96 |
| Skewness | 1.69 | 1.47 | 1.27 | 1.16 | 1.04 | 0.85 | 0.77 |
| Kurtosis | 7.37 | 6.33 | 5.47 | 4.99 | 4.62 | 3.98 | 3.73 |
| $\rho_1$ | 0.82 | 0.87 | 0.90 | 0.91 | 0.92 | 0.94 | 0.94 |
| $\rho_3$ | 0.66 | 0.75 | 0.80 | 0.82 | 0.84 | 0.86 | 0.88 |
| $\rho_6$ | 0.46 | 0.53 | 0.59 | 0.61 | 0.64 | 0.68 | 0.70 |

Panel B: Variance Swap Returns (monthly)

| Maturity | 1 | 3 | 6 | 9 | 12 | 18 | 24 |
|---|---|---|---|---|---|---|---|
| Mean | 0.0012 | 0.0012 | 0.0014 | 0.0013 | 0.0016 | 0.0016 | 0.0017 |
| Std. Deviation | 0.0039 | 0.0080 | 0.0118 | 0.0142 | 0.0165 | 0.0211 | 0.0258 |
| Sharpe Ratio | 0.3085 | 0.1533 | 0.1187 | 0.0937 | 0.0985 | 0.0756 | 0.0655 |
| t-Statistic | 3.6581 | 1.9903 | 1.5780 | 1.2481 | 1.3258 | 1.0351 | 0.9260 |
| Minimum | -0.0385 | -0.0801 | -0.1154 | -0.1321 | -0.1495 | -0.1810 | -0.2093 |
| Maximum | 0.0104 | 0.0243 | 0.0354 | 0.0439 | 0.0500 | 0.0610 | 0.0791 |
| Skewness | -5.4856 | -4.9982 | -4.4795 | -3.7836 | -3.4992 | -2.9007 | -2.5368 |
| Kurtosis | 53.6750 | 49.3152 | 45.0421 | 36.6971 | 33.4251 | 26.8390 | 22.4819 |
| Autocorrelation | 0.3957 | 0.3301 | 0.2798 | 0.2589 | 0.2165 | 0.1977 | 0.1373 |
| Percent Negative | 0.1617 | 0.2681 | 0.3319 | 0.3574 | 0.3617 | 0.4213 | 0.4426 |
| $\beta_{CAPM}$ | 0.0447 | 0.1140 | 0.1710 | 0.2077 | 0.2430 | 0.3001 | 0.3382 |
| $\alpha_{CAPM}$ | 0.0009 | 0.0006 | 0.0004 | 0.0002 | 0.0003 | -0.0001 | -0.0002 |
| $\alpha_{\text{t-Statistic}}$ | 2.9647 | 1.1191 | 0.5983 | 0.1843 | 0.2570 | -0.0765 | -0.1455 |
| $R^2_{adj}$ | 0.2761 | 0.4193 | 0.4399 | 0.4455 | 0.4529 | 0.4185 | 0.3552 |

**Table 2: Variance Swap Return Forecasting Regressions**

This table reports variance swap return forecasting regressions for different monthly horizons. The return for each horizon is annualized and equal-weighted across maturities by volatility given the observation in Table 1 that volatility is increasing in maturity. Panel A shows that the level and slope factors are the strongest predictors of variance swap returns among the principal components. Panel B confirms that high values of variance swap rates forecast high returns from selling variance swaps after controlling for well-known forecasting variables. Panel C reports the correlation matrix of the forecasting variables which are standardized to be mean zero and standard deviation one in each of the regressions.

Panel A: Principal Components

| | $h=1$ | $h=3$ | $h=6$ | $h=9$ | $h=12$ |
|---|---|---|---|---|---|
| $PC1$ | 0.020** | 0.026*** | 0.027*** | 0.022*** | 0.020*** |
| | [2.44] | [4.43] | [7.66] | [6.89] | [7.70] |
| $PC2$ | -0.036*** | -0.013* | -0.008** | -0.006*** | -0.005*** |
| | [-3.07] | [-1.95] | [-2.43] | [-2.67] | [-2.70] |
| $PC3$ | -0.039* | -0.014 | -0.009 | -0.001 | -0.000 |
| | [-1.71] | [-1.50] | [-1.49] | [-0.38] | [-0.12] |
| $PC4$ | -0.024** | -0.007 | 0.001 | -0.000 | 0.000 |
| | [-2.24] | [-1.60] | [0.25] | [-0.06] | [0.27] |
| $PC5$ | -0.003 | -0.000 | 0.001 | -0.002 | 0.002 |
| | [-0.38] | [-0.12] | [0.52] | [-1.05] | [1.01] |
| $R^2_{adj}$ | 0.20 | 0.12 | 0.20 | 0.21 | 0.24 |

Newey-West t-statistics $\max(3, 2h)$ lags: *p<0.10, **p<0.05, ***p<0.01

Panel B: Principal Components and Forecasting Variables

| | $h=1$ | $h=3$ | $h=6$ | $h=9$ | $h=12$ |
|---|---|---|---|---|---|
| $PC1$ | 0.083*** | 0.065*** | 0.047*** | 0.030*** | 0.024*** |
| | [3.75] | [5.16] | [4.09] | [4.09] | [4.09] |
| $PC2$ | 0.008 | 0.009 | 0.000 | -0.005 | -0.002 |
| | [0.58] | [1.20] | [0.05] | [-1.31] | [-0.64] |
| $RV$ | -0.090*** | -0.044*** | -0.016*** | -0.001 | -0.002 |
| | [-2.90] | [-5.43] | [-2.70] | [-0.34] | [-0.64] |
| Term Slope | -0.011 | -0.007 | -0.005 | -0.000 | 0.004 |
| | [-1.34] | [-0.98] | [-0.84] | [-0.03] | [1.04] |
| Credit | 0.023 | 0.001 | -0.007 | -0.008 | -0.001 |
| | [0.85] | [0.04] | [-0.60] | [-1.02] | [-0.15] |
| DP | 0.009 | 0.010 | 0.011* | 0.007* | 0.003 |
| | [1.00] | [1.16] | [1.85] | [1.70] | [0.86] |
| IP Growth | 0.029 | 0.019 | 0.014 | 0.008 | 0.010** |
| | [1.06] | [1.02] | [0.95] | [1.08] | [2.10] |
| $R^2_{adj}$ | 0.22 | 0.20 | 0.25 | 0.24 | 0.29 |

Newey-West t-statistics $\max(3, 2h)$ lags: *p<0.10, **p<0.05, ***p<0.01

Panel C: Correlation Matrix

| $PC1$ | $PC2$ | $RV$ | Term | Credit | DP | IP |
|---|---|---|---|---|---|---|
| 1.00 | 0.00 | 0.74 | 0.23 | 0.71 | 0.41 | -0.47 |
| | 1.00 | 0.46 | -0.15 | -0.01 | 0.02 | -0.17 |
| | | 1.00 | 0.13 | 0.61 | 0.39 | -0.53 |
| | | | 1.00 | 0.35 | 0.39 | -0.14 |
| | | | | 1.00 | 0.70 | -0.73 |
| | | | | | 1.00 | -0.46 |
| | | | | | | 1.00 |

**Table 3: Model Fit**

This table summarizes the model fitting errors during our 1996-2015 sample period. Panel A reports variance swap pricing errors in annualized volatility units. Panel B reports monthly return pricing errors. To illustrate these results, Figure 4 plots the fitting errors.

| Maturity | Avg. | 1 | 3 | 6 | 9 | 12 | 15 | 18 | 21 | 24 |
|---|---|---|---|---|---|---|---|---|---|---|
| Panel A: Variance Swap Pricing Errors | | | | | | | | | | |
| Mean (%) | 0.12 | 0.27 | -0.21 | 0.10 | 0.17 | 0.22 | 0.28 | 0.20 | 0.08 | -0.06 |
| Std. Dev (%) | 0.42 | 0.97 | 0.44 | 0.45 | 0.42 | 0.43 | 0.26 | 0.16 | 0.24 | 0.38 |
| Skewness | -0.01 | 0.77 | -0.10 | -0.10 | -0.17 | -1.74 | 1.09 | -0.01 | 0.09 | 0.10 |
| Kurtosis | 5.99 | 6.72 | 5.56 | 3.12 | 3.51 | 13.81 | 9.07 | 5.02 | 3.56 | 3.57 |
| $\rho_1$ | 0.40 | 0.08 | 0.25 | 0.45 | 0.60 | 0.37 | 0.43 | 0.35 | 0.50 | 0.54 |
| $\rho_6$ | 0.13 | 0.07 | 0.15 | 0.33 | 0.24 | 0.28 | -0.17 | 0.02 | 0.13 | 0.15 |
| Panel B: Return Pricing Errors | | | | | | | | | | |
| Mean (%) | 0.00 | 0.02 | -0.01 | -0.01 | -0.01 | 0.01 | 0.01 | 0.00 | 0.00 | -0.01 |
| Std. Dev (%) | 0.17 | 0.05 | 0.07 | 0.11 | 0.13 | 0.25 | 0.18 | 0.13 | 0.21 | 0.37 |
| Skewness | -0.11 | 1.96 | -1.59 | 0.56 | -0.77 | -0.53 | -2.89 | -0.28 | 1.02 | 1.49 |
| Kurtosis | 14.84 | 16.24 | 13.43 | 8.23 | 7.62 | 18.90 | 33.31 | 8.32 | 10.74 | 16.78 |
| $\rho_1$ | -0.16 | 0.24 | 0.18 | -0.18 | -0.22 | -0.43 | -0.24 | -0.07 | -0.34 | -0.35 |
| $\rho_6$ | -0.01 | 0.02 | 0.03 | -0.09 | -0.01 | 0.37 | -0.21 | -0.16 | 0.02 | -0.02 |

**Table 4: Model Estimate**

This table reports the prices of risk and physical parameters that we estimate by three-step linear regression. The state variables include standardized realized variance and the first two principal components of standardized variance swap rates. The results indicate that realized variance is priced unconditionally. In addition, each of the state variables contributes significantly to the time-variation in expected variance swap returns.

Panel A: Prices of Risk

| | $\Lambda_0$ | $\Lambda_{1,1}$ | $\Lambda_{1,2}$ | $\Lambda_{1,3}$ |
|---|---|---|---|---|
| $RV$ | -0.21*** | 0.45*** | -0.16*** | -0.15 |
| | [-4.29] | [4.50] | [-5.18] | [-1.40] |
| $PC1$ | -0.05 | 0.99*** | -0.31*** | -0.07 |
| | [-0.62] | [5.78] | [-5.88] | [-0.37] |
| $PC2$ | -0.02 | 0.37*** | -0.11*** | -0.29*** |
| | [-0.57] | [4.34] | [-4.19] | [-3.10] |

Panel B: Physical Parameters

| | $\mu$ | $\Phi_{1,1}$ | $\Phi_{1,2}$ | $\Phi_{1,3}$ |
|---|---|---|---|---|
| $RV$ | 0.00 | 0.61** | 0.03 | 0.18* |
| | [0.08] | [2.38] | [0.55] | [1.89] |
| $PC1$ | 0.02 | 0.81* | 0.69*** | -0.37 |
| | [0.22] | [1.92] | [7.10] | [-1.40] |
| $PC2$ | -0.00 | 0.13 | -0.08*** | 0.64*** |
| | [-0.02] | [1.14] | [-2.69] | [7.16] |

robust t-statistics: *p<0.10, **p<0.05, ***p<0.01

**Table 5: Global Variance Swap Summary Statistics: 2009 to 2015**

This table reports summary statistics for international variance swap rates and returns computed at a monthly frequency from 2009 to 2015 using Bloomberg data. Similar to the synthetic swap rate data during the full sample, we see that the international variance swap curves are upward sloping on average and positively skewed. In addition, the long-end of the curve is less volatile than the short-end and more highly autocorrelated at longer lags. Beyond these similarities to the full sample, we can also note that variance swap rates have been higher in Europe and Japan relative to the US and the UK during the recent recovery. At the same time, Panel B shows that average returns have been similar across the different countries and most volatile in Europe and Japan. As a result, the largest Sharpe ratios are observed for short-dated variance swaps in the US and UK.

Panel A: Variance Swap Rates (annualized volatility %)

| Index | SPX | | FTSE | | STOXX | | NKY | |
|---|---|---|---|---|---|---|---|---|
| Maturity | 1 | 12 | 1 | 12 | 1 | 12 | 1 | 12 |
| Mean | 20.21 | 24.89 | 19.33 | 23.34 | 25.13 | 27.52 | 24.98 | 27.83 |
| Std. Deviation | 7.61 | 6.27 | 6.73 | 6.38 | 7.15 | 6.40 | 6.73 | 5.73 |
| Minimum | 11.31 | 17.11 | 9.84 | 15.30 | 14.49 | 18.46 | 15.38 | 17.57 |
| Median | 17.73 | 23.34 | 18.05 | 22.61 | 23.45 | 26.08 | 23.44 | 26.57 |
| Maximum | 44.37 | 42.96 | 42.29 | 45.71 | 46.99 | 44.19 | 52.07 | 54.20 |
| Skewness | 1.47 | 0.97 | 1.23 | 0.99 | 0.97 | 0.64 | 1.64 | 1.98 |
| Kurtosis | 4.86 | 3.41 | 4.60 | 3.96 | 3.64 | 2.67 | 6.71 | 8.98 |
| $\rho_1$ | 0.79 | 0.91 | 0.79 | 0.91 | 0.76 | 0.90 | 0.70 | 0.66 |
| $\rho_3$ | 0.66 | 0.83 | 0.68 | 0.85 | 0.66 | 0.83 | 0.53 | 0.55 |
| $\rho_6$ | 0.31 | 0.57 | 0.40 | 0.64 | 0.28 | 0.58 | 0.10 | 0.23 |

Panel B: Variance Swap Returns (monthly)

| Index | SPX | | FTSE | | STOXX | | NKY | |
|---|---|---|---|---|---|---|---|---|
| Maturity | 1 | 12 | 1 | 12 | 1 | 12 | 1 | 12 |
| Mean | 0.0012 | 0.0049 | 0.0012 | 0.0048 | 0.0014 | 0.0041 | 0.0015 | 0.0047 |
| Std. Deviation | 0.0025 | 0.0162 | 0.0023 | 0.0159 | 0.0035 | 0.0195 | 0.0038 | 0.0270 |
| Sharpe Ratio | 0.4969 | 0.3022 | 0.5302 | 0.2984 | 0.4080 | 0.2115 | 0.3850 | 0.1754 |
| t-Statistic | 4.2510 | 2.9282 | 3.7437 | 2.8565 | 3.3646 | 2.0544 | 3.1748 | 2.1490 |
| Minimum | -0.0122 | -0.0495 | -0.0066 | -0.0496 | -0.0151 | -0.0644 | -0.0197 | -0.0739 |
| Maximum | 0.0076 | 0.0505 | 0.0100 | 0.0508 | 0.0077 | 0.0483 | 0.0176 | 0.0734 |
| Skewness | -1.8169 | -0.2659 | -0.0419 | -0.2855 | -1.7674 | -0.5804 | -1.4112 | -0.1557 |
| Kurtosis | 12.9421 | 5.2771 | 6.9355 | 5.4282 | 9.0920 | 4.6650 | 17.0824 | 3.8688 |
| Autocorrelation | -0.0271 | -0.0063 | 0.1607 | -0.1100 | 0.0915 | -0.0710 | 0.0055 | -0.2985 |
| Percent Negative | 0.1500 | 0.2875 | 0.1875 | 0.3000 | 0.2375 | 0.3500 | 0.1500 | 0.3375 |
| $\beta_{CAPM}$ | 0.0280 | 0.2735 | 0.0177 | 0.2225 | 0.0369 | 0.3231 | 0.0108 | 0.2783 |
| $\alpha_{CAPM}$ | 0.0009 | 0.0014 | 0.0010 | 0.0019 | 0.0009 | -0.0000 | 0.0013 | 0.0012 |
| $\alpha_{\text{t-Statistic}}$ | 2.9693 | 0.9854 | 2.2494 | 1.0073 | 1.7926 | -0.0118 | 2.3233 | 0.4686 |
| $R^2_{adj}$ | 0.2232 | 0.5244 | 0.1011 | 0.3550 | 0.2033 | 0.5037 | 0.0022 | 0.1880 |

**Table 6: Global Model Fit: 2009 to 2015**

This table summarizes the fitting errors in the estimated model for the international variance swap data from Bloomberg from 2009 to 2015. The results are similar to the full sample. As before, the model prices variance swaps accurately. The pricing errors are within the typical bid-ask spread in the variance swap market. In addition, we can see the pricing errors exhibit positive autocorrelation at a monthly frequency and excess kurtosis. In unreported results, we confirm the model has similar pricing errors using the synthetic variance swap data over this subsample.

| Maturity | Avg. | 1 | 3 | 6 | 9 | 12 | 15 | 18 | 21 | 24 |
|---|---|---|---|---|---|---|---|---|---|---|
| Panel A: SPX Variance Swap Pricing Errors | | | | | | | | | | |
| Mean (%) | 0.30 | -0.23 | 0.40 | 0.70 | 0.61 | 0.47 | 0.44 | 0.25 | 0.06 | -0.05 |
| Std. Dev (%) | 0.42 | 0.76 | 0.46 | 0.56 | 0.48 | 0.39 | 0.29 | 0.20 | 0.24 | 0.37 |
| Skewness | 0.38 | -0.48 | 1.36 | 1.00 | 0.38 | -0.03 | 1.04 | 0.53 | -0.46 | 0.07 |
| Kurtosis | 3.89 | 2.86 | 7.34 | 4.56 | 3.26 | 3.60 | 4.46 | 3.76 | 2.78 | 2.40 |
| $\rho_1$ | 0.47 | 0.54 | 0.45 | 0.59 | 0.57 | 0.39 | 0.43 | 0.27 | 0.38 | 0.63 |
| $\rho_6$ | 0.11 | 0.26 | 0.08 | 0.19 | 0.10 | 0.25 | -0.32 | 0.21 | -0.03 | 0.25 |
| Panel B: FTSE Variance Swap Pricing Errors | | | | | | | | | | |
| Mean (%) | 0.58 | 0.03 | -0.11 | 0.50 | 0.70 | 0.81 | 0.95 | 0.96 | 0.82 | 0.54 |
| Std. Dev (%) | 0.61 | 1.05 | 0.44 | 0.69 | 0.70 | 0.62 | 0.54 | 0.51 | 0.45 | 0.51 |
| Skewness | 0.41 | -0.27 | -0.10 | 0.32 | 0.56 | 0.57 | 0.28 | 0.84 | 1.05 | 0.45 |
| Kurtosis | 3.79 | 3.56 | 5.07 | 3.67 | 4.05 | 3.03 | 2.06 | 3.39 | 4.24 | 5.00 |
| $\rho_1$ | 0.34 | -0.14 | 0.07 | 0.27 | 0.36 | 0.49 | 0.67 | 0.74 | 0.51 | 0.07 |
| $\rho_6$ | 0.18 | -0.18 | 0.03 | 0.02 | 0.24 | 0.32 | 0.48 | 0.37 | 0.21 | 0.09 |
| Panel C: STOXX Variance Swap Pricing Errors | | | | | | | | | | |
| Mean (%) | 0.25 | -0.09 | -0.13 | 0.51 | 0.56 | 0.49 | 0.51 | 0.28 | 0.09 | 0.03 |
| Std. Dev (%) | 0.44 | 0.61 | 0.45 | 0.53 | 0.42 | 0.46 | 0.39 | 0.34 | 0.35 | 0.42 |
| Skewness | -0.16 | 0.38 | 0.03 | 0.38 | 1.05 | -1.09 | -0.03 | -1.37 | -0.42 | -0.35 |
| Kurtosis | 4.59 | 3.45 | 3.13 | 3.33 | 4.78 | 5.73 | 5.29 | 7.74 | 3.71 | 4.14 |
| $\rho_1$ | 0.33 | 0.06 | 0.06 | 0.38 | 0.46 | 0.21 | 0.42 | 0.52 | 0.51 | 0.36 |
| $\rho_6$ | -0.10 | -0.14 | 0.10 | -0.06 | -0.16 | 0.28 | -0.42 | -0.08 | -0.29 | -0.16 |
| Panel D: NKY Variance Swap Pricing Errors | | | | | | | | | | |
| Mean (%) | -0.03 | -0.37 | -0.53 | -0.24 | 0.06 | 0.21 | 0.22 | 0.18 | 0.12 | 0.07 |
| Std. Dev (%) | 1.03 | 1.54 | 0.80 | 0.77 | 1.05 | 1.25 | 1.17 | 1.05 | 0.93 | 0.75 |
| Skewness | 0.33 | -0.01 | -0.95 | -0.91 | -0.35 | 0.28 | 0.51 | 0.96 | 1.47 | 1.92 |
| Kurtosis | 5.47 | 2.65 | 4.81 | 4.98 | 3.68 | 3.67 | 3.95 | 5.23 | 8.28 | 11.97 |
| $\rho_1$ | 0.21 | 0.16 | 0.19 | 0.02 | 0.28 | 0.31 | 0.29 | 0.26 | 0.21 | 0.17 |
| $\rho_6$ | 0.04 | 0.10 | -0.03 | -0.13 | 0.10 | 0.12 | 0.08 | 0.04 | 0.03 | 0.06 |

**Table 7: Global Model Estimate: 2009 to 2015**

This table reports the estimated prices of risk and physical parameters for the international variance swap data from Bloomberg from 2009 to 2015. The state variables include realized variance and the first two principal components of variance swaps, which are standardized. Similar to the full sample, we see that realized variance is priced unconditionally across all of the countries. In addition, the table highlights how realized variance, level, and slope contribute differentially to the time-variation in expected variance swap returns across the panel.

Panel A: Price of Risk Estimates

Panel A.I: SPX 01/09-08/15 ($T = 80$)

| | $\Lambda_0$ | $\Lambda_{1,1}$ | $\Lambda_{1,2}$ | $\Lambda_{1,3}$ |
|---|---|---|---|---|
| $RV$ | -0.39*** | -0.08 | -0.10* | 0.16 |
| | [-4.66] | [-0.54] | [-1.94] | [0.77] |
| $PC1$ | -0.36*** | 0.87*** | -0.38*** | -0.21 |
| | [-2.67] | [3.42] | [-4.63] | [-0.61] |
| $PC2$ | -0.05 | 0.11 | -0.04 | -0.06 |
| | [-0.71] | [1.01] | [-1.03] | [-0.46] |

Panel A.II: FTSE 01/09-08/15 ($T = 80$)

| | $\Lambda_0$ | $\Lambda_{1,1}$ | $\Lambda_{1,2}$ | $\Lambda_{1,3}$ |
|---|---|---|---|---|
| $RV$ | -0.48*** | -0.21 | -0.14** | 0.30 |
| | [-3.80] | [-1.19] | [-2.34] | [0.97] |
| $PC1$ | -0.40*** | 0.81*** | -0.44*** | -0.38 |
| | [-2.91] | [3.65] | [-6.30] | [-1.04] |
| $PC2$ | -0.11 | 0.28*** | -0.12*** | -0.61*** |
| | [-1.45] | [2.80] | [-3.44] | [-3.47] |

Panel A.III: STOXX 01/09-08/15 ($T = 80$)

| | $\Lambda_0$ | $\Lambda_{1,1}$ | $\Lambda_{1,2}$ | $\Lambda_{1,3}$ |
|---|---|---|---|---|
| $RV$ | -0.34*** | 0.07 | -0.13*** | 0.15 |
| | [-3.49] | [0.44] | [-2.69] | [0.62] |
| $PC1$ | -0.24 | 0.68*** | -0.33*** | -0.14 |
| | [-1.55] | [2.71] | [-4.23] | [-0.37] |
| $PC2$ | -0.05 | 0.22* | -0.10*** | -0.23 |
| | [-0.70] | [1.96] | [-2.68] | [-1.25] |

Panel A.IV: NKY 01/09-08/15 ($T = 80$)

| | $\Lambda_0$ | $\Lambda_{1,1}$ | $\Lambda_{1,2}$ | $\Lambda_{1,3}$ |
|---|---|---|---|---|
| $RV$ | -0.45*** | -0.32 | -0.19*** | 0.15 |
| | [-3.60] | [-1.53] | [-3.79] | [0.47] |
| $PC1$ | -0.34** | -0.31 | -0.34*** | 0.98*** |
| | [-2.16] | [-1.25] | [-5.06] | [3.00] |
| $PC2$ | -0.10 | 0.20 | 0.03 | -0.67** |
| | [-0.86] | [1.14] | [0.71] | [-2.21] |

Panel B: Physical Parameters

Panel B.I: SPX 01/09-08/15 ($T = 80$)

| | $\mu$ | $\Phi_{1,1}$ | $\Phi_{1,2}$ | $\Phi_{1,3}$ |
|---|---|---|---|---|
| $RV$ | -0.03 | -0.09 | 0.23*** | 0.78** |
| | [-0.34] | [-0.51] | [3.36] | [2.57] |
| $PC1$ | -0.13 | 0.86** | 0.59*** | -0.63 |
| | [-1.01] | [2.09] | [5.16] | [-1.42] |
| $PC2$ | 0.00 | 0.09 | -0.05 | 0.67*** |
| | [0.12] | [0.90] | [-1.47] | [5.17] |

Panel B.II: FTSE 01/09-08/15 ($T = 80$)

| | $\mu$ | $\Phi_{1,1}$ | $\Phi_{1,2}$ | $\Phi_{1,3}$ |
|---|---|---|---|---|
| $RV$ | -0.02 | 0.07 | 0.15* | 0.72** |
| | [-0.18] | [0.31] | [1.72] | [2.39] |
| $PC1$ | -0.17 | 0.63* | 0.62*** | -0.52 |
| | [-1.29] | [1.67] | [6.69] | [-1.18] |
| $PC2$ | 0.01 | 0.02 | -0.03 | 0.54*** |
| | [0.14] | [0.24] | [-0.87] | [4.15] |

Panel B.III: STOXX 01/09-08/15 ($T = 80$)

| | $\mu$ | $\Phi_{1,1}$ | $\Phi_{1,2}$ | $\Phi_{1,3}$ |
|---|---|---|---|---|
| $RV$ | -0.00 | 0.12 | 0.13*** | 0.66*** |
| | [-0.04] | [0.80] | [2.79] | [3.41] |
| $PC1$ | -0.12 | 0.61* | 0.67*** | -0.47 |
| | [-0.79] | [1.94] | [7.81] | [-1.16] |
| $PC2$ | 0.01 | 0.10 | -0.07*** | 0.64*** |
| | [0.20] | [1.53] | [-3.00] | [5.35] |
| robust t-stats: *p<0.10, **p<0.05, ***p<0.01 | | | | |

Panel B.IV: NKY 01/09-08/15 ($T = 80$)

| | $\mu$ | $\Phi_{1,1}$ | $\Phi_{1,2}$ | $\Phi_{1,3}$ |
|---|---|---|---|---|
| $RV$ | -0.03 | -0.17 | 0.11** | 0.54*** |
| | [-0.36] | [-1.31] | [2.11] | [3.36] |
| $PC1$ | -0.20 | -0.44** | 0.67*** | 0.78*** |
| | [-1.38] | [-2.57] | [10.20] | [3.54] |
| $PC2$ | 0.00 | -0.01 | 0.05* | 0.37*** |
| | [0.05] | [-0.07] | [1.72] | [3.06] |
| robust t-stats: *p<0.10, **p<0.05, ***p<0.01 | | | | |

**Table 8: The Price of Global Level Risk: 2009 to 2015**

This table reports estimated prices of risk for global $PC1G$ and local $PC1L$ factors for the level of variance swap rates. The global factor is the first principal component from the stacked international data. The local factor is the residual from a regression of the index specific level factor onto the global factor. Panel A reports a regression of the index specific level factors onto the global factor. It shows that the global factor explains the vast majority of the variation in variance swap rates across the different markets. Panel B reports the estimated prices of risk when all the state variables have been standardized. The global factor is the only common variable across the subpanels. As the results indicate, the global factor is priced in each of the countries and plays a similar role to the index specific level factor in the previous results. By way of contrast, the local factor is only priced in Japan where it is marginally statistically significant.

Panel A: Explanatory Power of Global Level Risk

| | $PC1_{SPX}$ | $PC1_{FTSE}$ | $PC1_{STOXX}$ | $PC1_{NKY}$ |
|---|---|---|---|---|
| $PC1G$ | 0.98*** | 0.99*** | 0.96*** | 0.88*** |
| | [21.20] | [58.39] | [11.64] | [5.85] |
| $R^2_{adj}$ | 0.97 | 0.98 | 0.93 | 0.77 |

Panel B: Price of Risk Estimates

Panel B.I: SPX 01/09-08/15 ($T = 80$)

| | $\Lambda_0$ | $\Lambda_{1,1}$ | $\Lambda_{1,2}$ | $\Lambda_{1,3}$ | $\Lambda_{1,4}$ |
|---|---|---|---|---|---|
| $RV$ | -0.40*** | -0.08 | -0.29** | -0.05 | 0.07 |
| | [-4.45] | [-0.53] | [-2.02] | [-0.58] | [0.75] |
| $PC1G$ | -0.01 | 0.38** | -0.36** | -0.13 | -0.04 |
| | [-0.12] | [2.40] | [-2.41] | [-1.25] | [-0.37] |
| $PC1L$ | -0.62 | -0.41 | -0.13 | 0.19 | -0.01 |
| | [-1.61] | [-0.55] | [-0.19] | [0.37] | [-0.03] |
| $PC2$ | -0.15 | 0.27 | -0.28 | -0.21* | -0.12 |
| | [-0.90] | [1.10] | [-1.21] | [-1.91] | [-0.97] |

Panel B.II: FTSE 01/09-08/15 ($T = 80$)

| | $\Lambda_0$ | $\Lambda_{1,1}$ | $\Lambda_{1,2}$ | $\Lambda_{1,3}$ | $\Lambda_{1,4}$ |
|---|---|---|---|---|---|
| $RV$ | -0.55*** | 0.19 | -0.84*** | -0.25** | -0.17 |
| | [-5.44] | [0.89] | [-3.83] | [-2.10] | [-1.14] |
| $PC1G$ | -0.21 | 0.63* | -0.84** | -0.11 | -0.28 |
| | [-1.48] | [1.80] | [-2.03] | [-0.91] | [-1.33] |
| $PC1L$ | 0.61 | -3.35 | 3.74 | 0.40 | 2.02 |
| | [0.59] | [-1.23] | [1.15] | [0.43] | [1.24] |
| $PC2$ | -0.14 | 0.00 | -0.09 | 0.07 | -0.20 |
| | [-1.00] | [0.00] | [-0.25] | [0.42] | [-0.97] |

Panel B.III: STOXX 01/09-08/15 ($T = 80$)

| | $\Lambda_0$ | $\Lambda_{1,1}$ | $\Lambda_{1,2}$ | $\Lambda_{1,3}$ | $\Lambda_{1,4}$ |
|---|---|---|---|---|---|
| $RV$ | -0.36*** | 0.11 | -0.43*** | -0.11 | 0.06 |
| | [-3.98] | [0.77] | [-3.27] | [-1.06] | [0.54] |
| $PC1G$ | -0.09 | 0.24* | -0.31** | -0.05 | -0.03 |
| | [-1.18] | [1.68] | [-2.47] | [-0.52] | [-0.22] |
| $PC1L$ | 0.04 | -0.03 | -0.03 | -0.29 | -0.03 |
| | [0.17] | [-0.07] | [-0.09] | [-0.92] | [-0.09] |
| $PC2$ | -0.08 | 0.31* | -0.43*** | -0.22* | -0.21 |
| | [-0.73] | [1.87] | [-3.04] | [-1.74] | [-1.52] |

Panel B.IV: NKY 01/09-08/15 ($T = 80$)

| | $\Lambda_0$ | $\Lambda_{1,1}$ | $\Lambda_{1,2}$ | $\Lambda_{1,3}$ | $\Lambda_{1,4}$ |
|---|---|---|---|---|---|
| $RV$ | -0.40*** | -0.36* | -0.53*** | -0.17 | 0.11 |
| | [-3.47] | [-1.94] | [-4.36] | [-1.29] | [0.51] |
| $PC1G$ | -0.04 | 0.00 | -0.20* | -0.05 | 0.04 |
| | [-0.50] | [0.03] | [-1.95] | [-0.49] | [0.24] |
| $PC1L$ | -0.20* | -0.23 | -0.25 | -0.25 | 0.49* |
| | [-1.65] | [-1.05] | [-1.55] | [-1.48] | [1.82] |
| $PC2$ | -0.16 | 0.15 | 0.02 | 0.32** | -0.54* |
| | [-1.11] | [0.66] | [0.16] | [2.13] | [-1.82] |

**Table 9: Pricing Factor Interpretations**

This table reports regressions of variance swap principal components on measures of intermediary risk appetite. Hedge fund (HF) risk appetite is measured by the drawdown on the broad HFRI Fund Weighted Composite index. Broker-dealer (BD) risk appetite is measured by broker-dealer leverage obtained from the Flow of Funds. Mutual fund (MF) risk appetite is measured by the ratio of equity mutual fund redemptions to total fund size. Newey-West $t$-statistics with 12 lags are reported in brackets. The sample consists of monthly observations from 1996m1-2015m8 for the US using our synthetic variance swap rates and monthly observations from 2006m9-2015m8 for the global analysis using Markit variance swap rates.

| | US PC1 | | | | US PC2 | | | |
|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (1) | (2) | (3) | (4) |
| HF Drawdown | -0.77*** | | | -0.63*** | 0.03 | | | 0.16* |
| | [-10.01] | | | [-11.30] | [0.20] | | | [1.89] |
| BD Leverage | | 0.17 | | -0.05 | | 0.46*** | | 0.45** |
| | | [1.22] | | [-0.79] | | [3.18] | | [2.59] |
| MF Redemption | | | 0.61*** | 0.36*** | | | 0.18 | 0.12 |
| | | | [4.51] | [5.12] | | | [0.94] | [0.71] |
| $R^2$ | 0.59 | 0.02 | 0.36 | 0.69 | -0.00 | 0.21 | 0.03 | 0.22 |

| | Global PC1 | | | | Global PC2 | | | |
|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (1) | (2) | (3) | (4) |
| HF Drawdown | -0.82*** | | | -0.62*** | 0.01 | | | 0.06 |
| | [-7.87] | | | [-9.27] | [0.10] | | | [0.54] |
| BD Leverage | | 0.18 | | -0.06 | | 0.50*** | | 0.53*** |
| | | [1.19] | | [-0.98] | | [3.49] | | [3.79] |
| MF Redemption | | | 0.71*** | 0.44*** | | | 0.09 | -0.06 |
| | | | [8.65] | [8.07] | | | [0.79] | [-0.44] |
| $R^2$ | 0.67 | 0.02 | 0.50 | 0.81 | -0.01 | 0.24 | -0.00 | 0.24 |

**Table 10: Intermediary Risk Appetite and Variance Term Premia**

This table reports time-series regressions of variance term premia on contemporaneous measures of intermediary marginal values of wealth. The four panels use the 1, 6, 12, and 24 month variance term premium as the LHS variable, where the VTP is obtained from our term structure model. Hedge fund (HF) risk appetite is measured by the drawdown on the broad HFRI Fund Weighted Composite index. Broker-dealer (BD) risk appetite is measured by broker-dealer leverage obtained from the Flow of Funds. Mutual fund (MF) risk appetite is measured by the ratio of equity mutual fund redemptions to total fund size. Newey-West $t$-statistics with 12 lags are reported in brackets. The sample consists of monthly observations from 1996m1-2015m8 for the US using our synthetic variance swap rates and monthly observations from 2006m9-2015m8 for the global analysis using Markit variance swap rates.

| | VTP(1) | | | | VTP(6) | | | |
|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (1) | (2) | (3) | (4) |
| HF Drawdown | -0.32*** | | | -0.45*** | -0.60*** | | | -0.64*** |
| | [-3.02] | | | [-9.35] | [-7.34] | | | [-12.75] |
| BD Leverage | | -0.32*** | | -0.36*** | | -0.24** | | -0.38*** |
| | | [-3.30] | | [-4.25] | | [-2.54] | | [-3.99] |
| MF Redemption | | | -0.06 | -0.14 | | | 0.23* | 0.07 |
| | | | [-0.37] | [-1.08] | | | [1.78] | [0.62] |
| $R^2$ | 0.10 | 0.10 | -0.00 | 0.26 | 0.36 | 0.05 | 0.05 | 0.48 |

| | VTP(12) | | | | VTP(24) | | | |
|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (1) | (2) | (3) | (4) |
| HF Drawdown | -0.69*** | | | -0.67*** | -0.72*** | | | -0.66*** |
| | [-10.84] | | | [-11.89] | [-11.79] | | | [-11.09] |
| BD Leverage | | -0.12 | | -0.30*** | | -0.04 | | -0.24** |
| | | [-0.97] | | [-2.88] | | [-0.28] | | [-2.16] |
| MF Redemption | | | 0.38*** | 0.19* | | | 0.47*** | 0.26*** |
| | | | [3.33] | [1.86] | | | [3.98] | [2.67] |
| $R^2$ | 0.48 | 0.01 | 0.14 | 0.57 | 0.52 | -0.00 | 0.22 | 0.61 |

**Figure 1: Synthetic Variance Swap Rates**

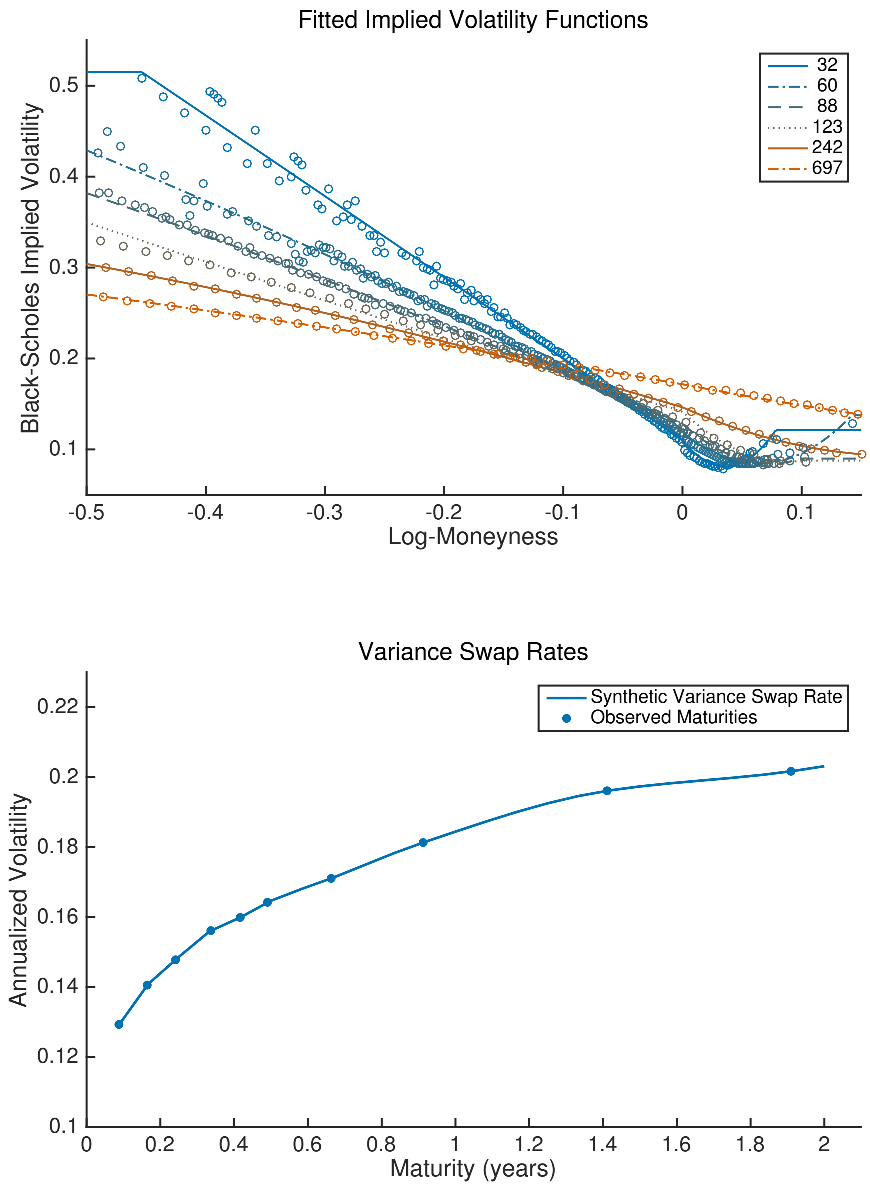


The top figure plots our fitted implied volatility functions on July 20th, 2015 against log-moneyness for six maturities whose time-to-expiration is in the legend. The bottom plot reports the resulting term structure of synthetic variance swaps. We compute synthetic variance swap rates by integrating over the moneyness dimension for each maturity using a weighted average of implied variances, which follows Carr and Lee (2009). We then interpolate between the observed maturities onto a monthly grid from one-month to two-years. Synthetic variance swaps are an approximation of the risk-neutral expected value of the realized variance for an underlying index from the current date until the option's expiration date, thus providing a term structure of equity risk. The underlying index in this example is the S&P 500 Index.

**Figure 2: Variance Swap Term Structure**

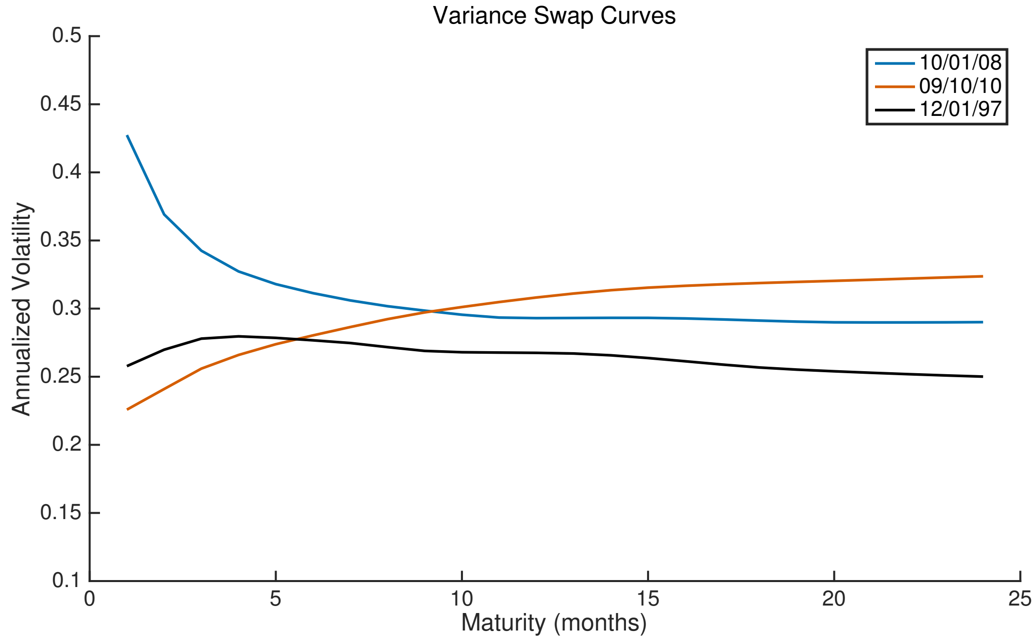


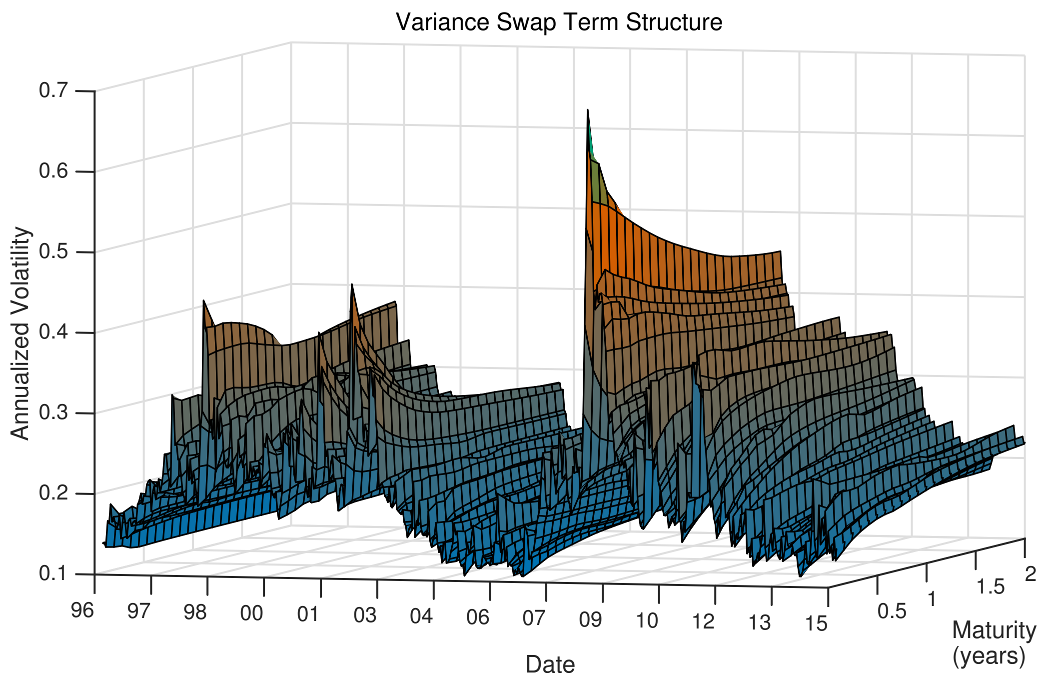


This figure highlights the different shapes of the variance swap curve and plots the variance swap term structure from 1996 to 2015. While the curve is upward sloping on average, it can also invert during periods of market distress such as the financial crisis or take on hump-shaped patterns when uncertainty is large at an intermediate horizon.

**Figure 3: Variance Swap Principal Components**

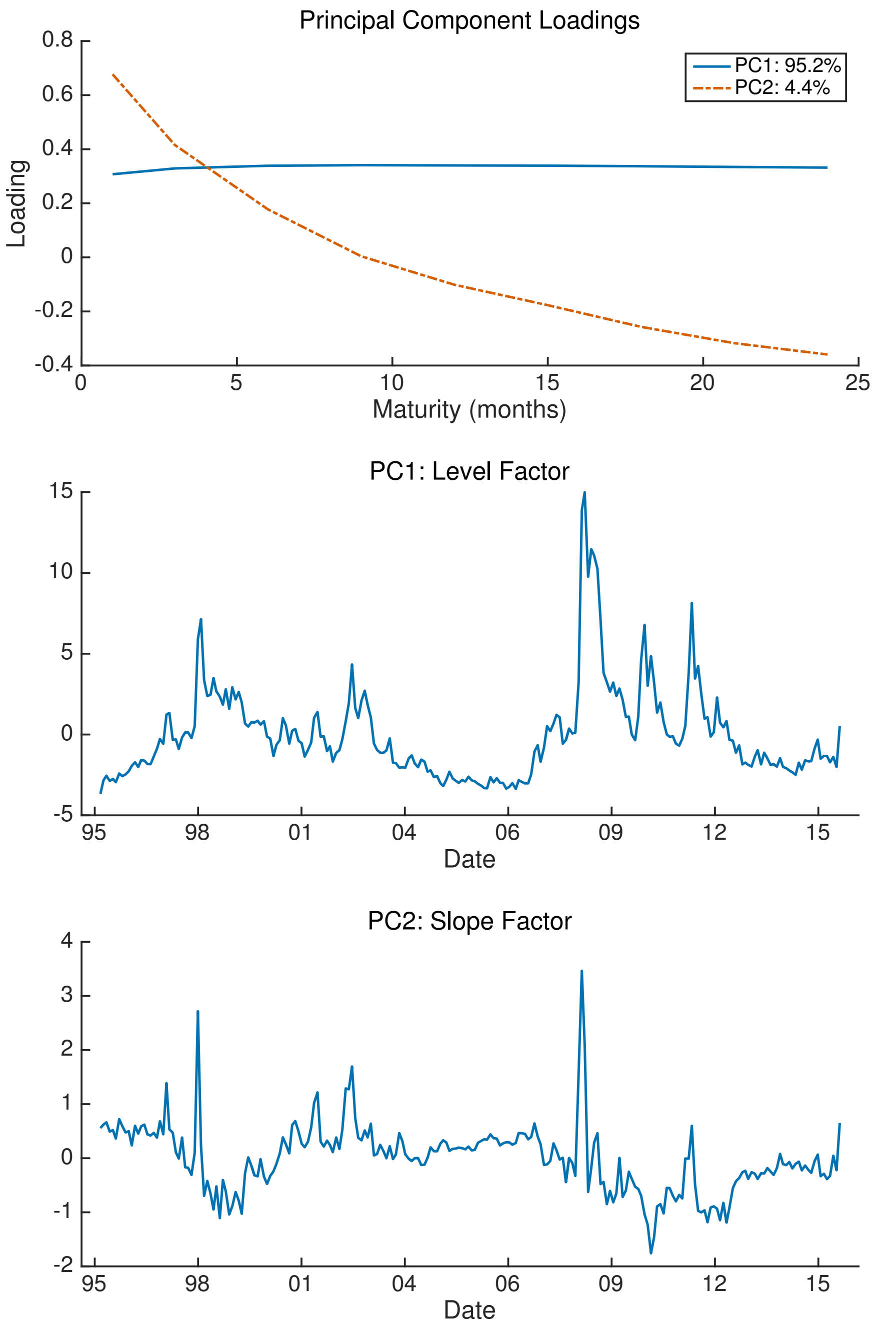


This figure plots the loadings and time-series of the first two principal components of standardized variance swap rates. As the top plot indicates, these factors can be interpreted as level and slope factors that explain over 99% of the variation in the curve. The level factor is positively correlated with the VIX (91%) and realized variance (74%).

**Figure 4: Variance Swap Pricing**

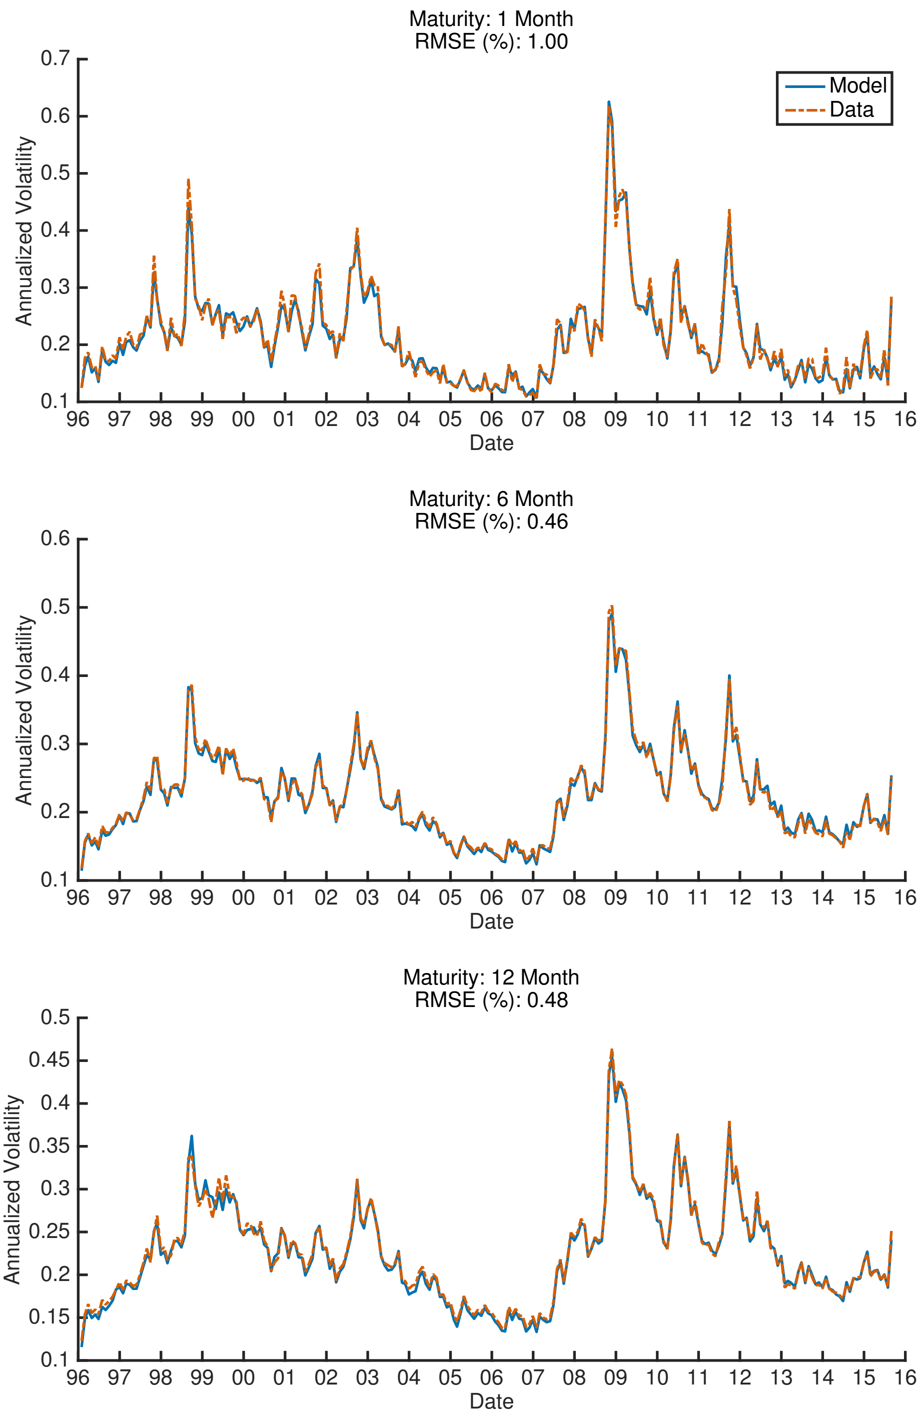


This figure plots variance swap rates $VS_{n,t} = A_n + B_n' X_t$ in the estimated model for $n = 1$, 6, and 12-month maturities against the synthetic variance swap rates during our 1996-2012 sample period. As the plot indicates, the model explains the cross-section of variance swap rates with small pricing errors. For each of the maturities, the root-mean-squared pricing error is less than 1% in annualized volatility units.

**Figure 5: Variance Term Premia**

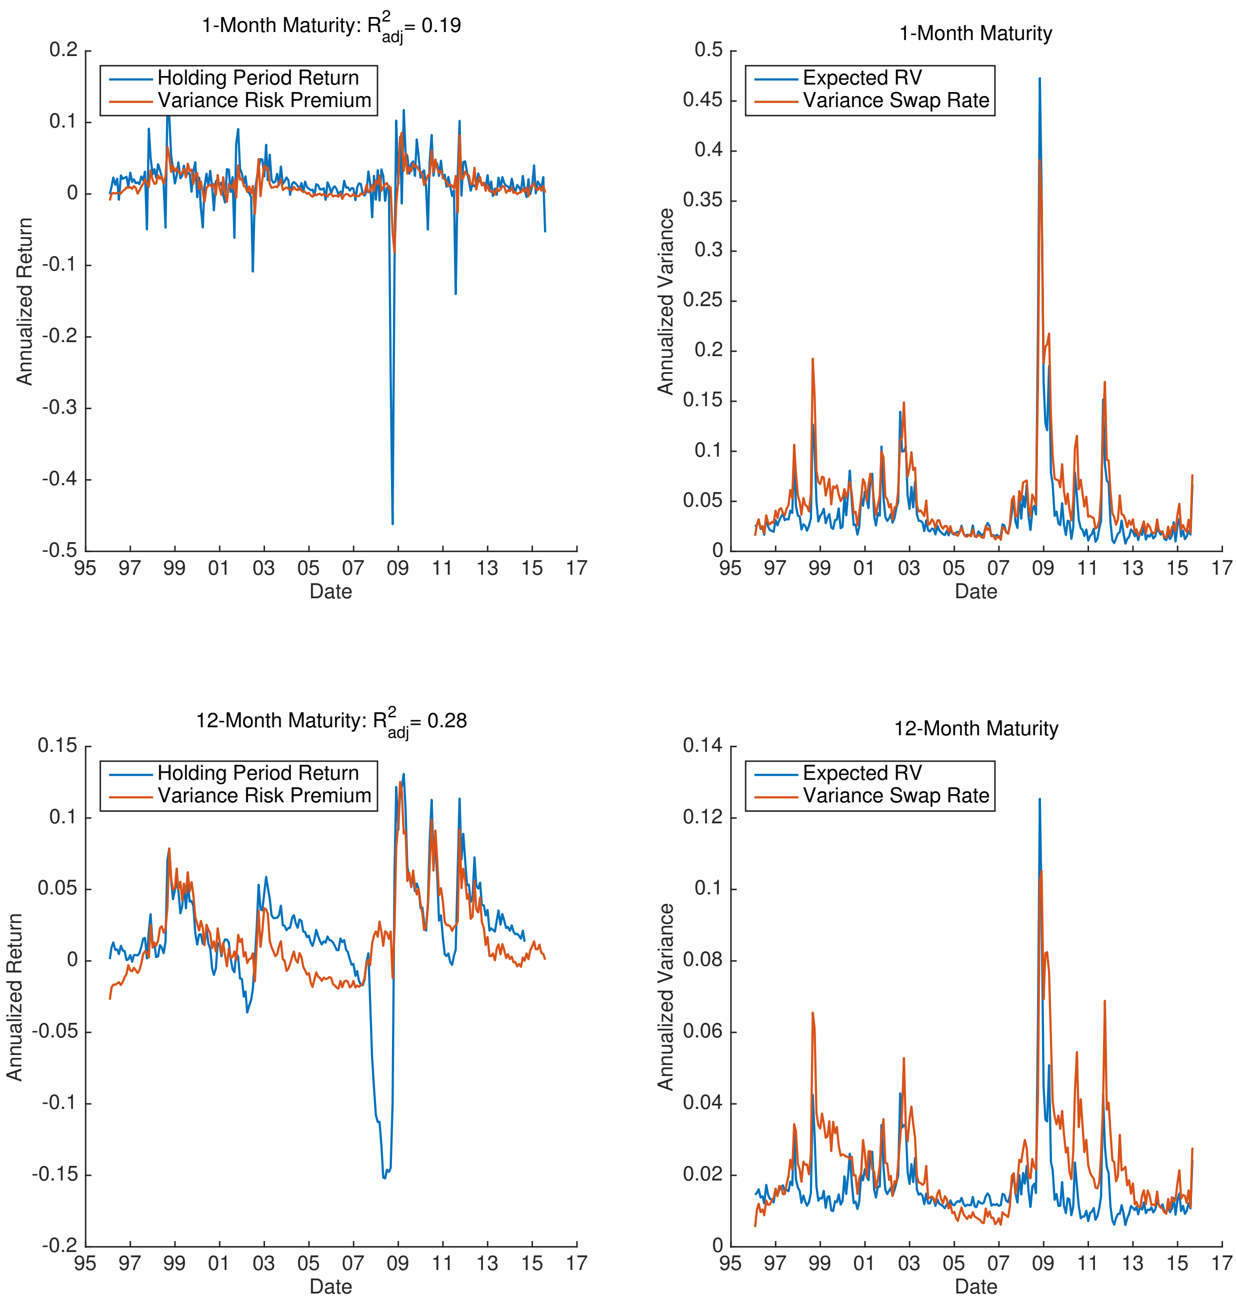


This figure plots the annualized variance term premium at a 1-month and 12-month horizon against realized holding period returns from receiving fixed in variance swaps during our 1996-2015 sample period. The annualized variance term premium is defined as $VTP_{n,t} = 12/n \cdot (E_t^{\mathbb{Q}}[\sum_{i=1}^{n} RV_{t+i}] - E_t^{\mathbb{P}}[\sum_{i=1}^{n} RV_{t+i}])$. The plots on the right decompose the variance term premium into the variance swap rate $VS_{n,t} = E_t^{\mathbb{Q}}[\sum_{i=1}^{n} RV_{t+i}]$, which is observable and is priced with a small error, and the expected amount of realized variance $E_t^{\mathbb{P}}[\sum_{i=1}^{n} RV_{t+i}]$, which is unobservable and is computed from the model's estimated vector-autoregression.

**Figure 6: Variance Term Premia Comparative Statics**

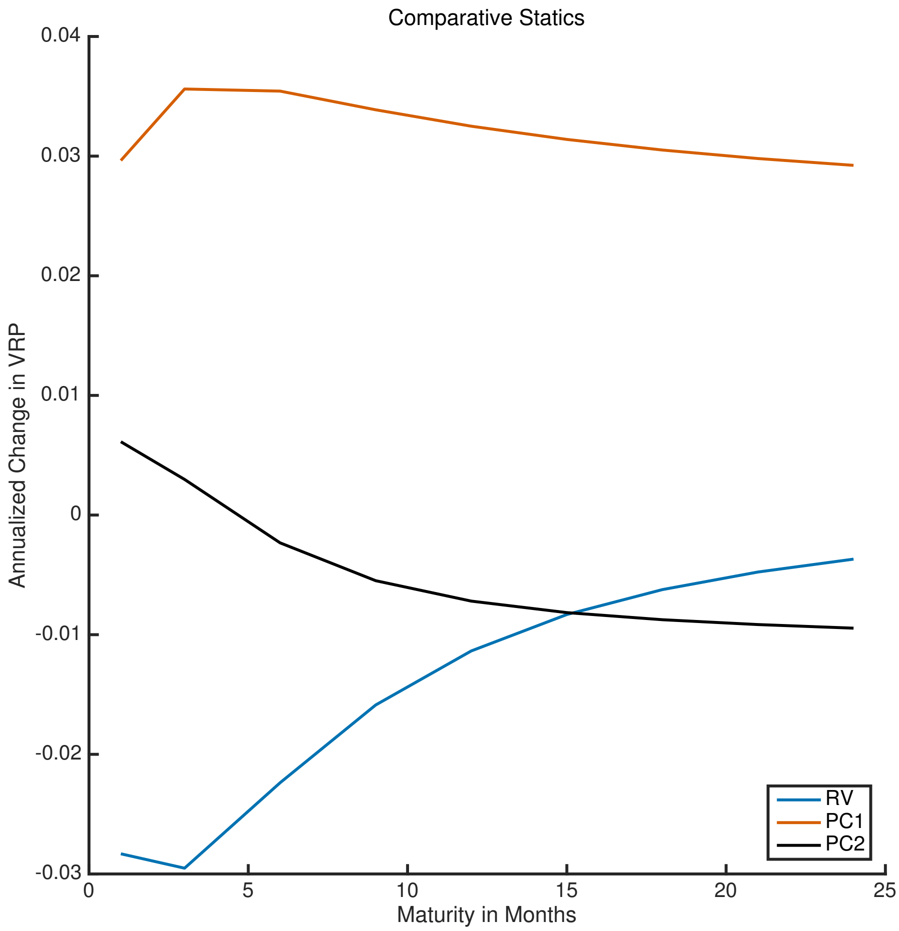


This figure plots the annualized change in the estimated variance term premium for a one-standard deviation move in the state variables at different horizons. The results illustrate how the estimated prices of risk translate into time-variation in the variance term premium. In addition, the results are consistent with the intuition that emerged from the unconstrained forecasting regressions in Table 2. In particular, the plot indicates that an increase in realized variance results in a decrease in the variance term premium that is strongest in magnitude at the short-end of the curve. Meanwhile, an increase in realized variance has a negative but more muted impact on the variance term premium at the long-end of the curve. This distinction stems from the fact that realized variance mean reverts quickly in the model's estimated vector-autoregression. By way of contrast, the level factor is more persistent than realized variance. As a result, an increase in the level factor results in an increase in the variance term premium that is fairly similar across the curve. Last, we can note that the slope factor has a differential impact on the variance term premium at the short-end and long-end of the curve. However, we can also see that the impact of the slope factor is relatively small in comparison to the other factors.

**Figure 7: Global Variance Swap Rates and Term Premia**

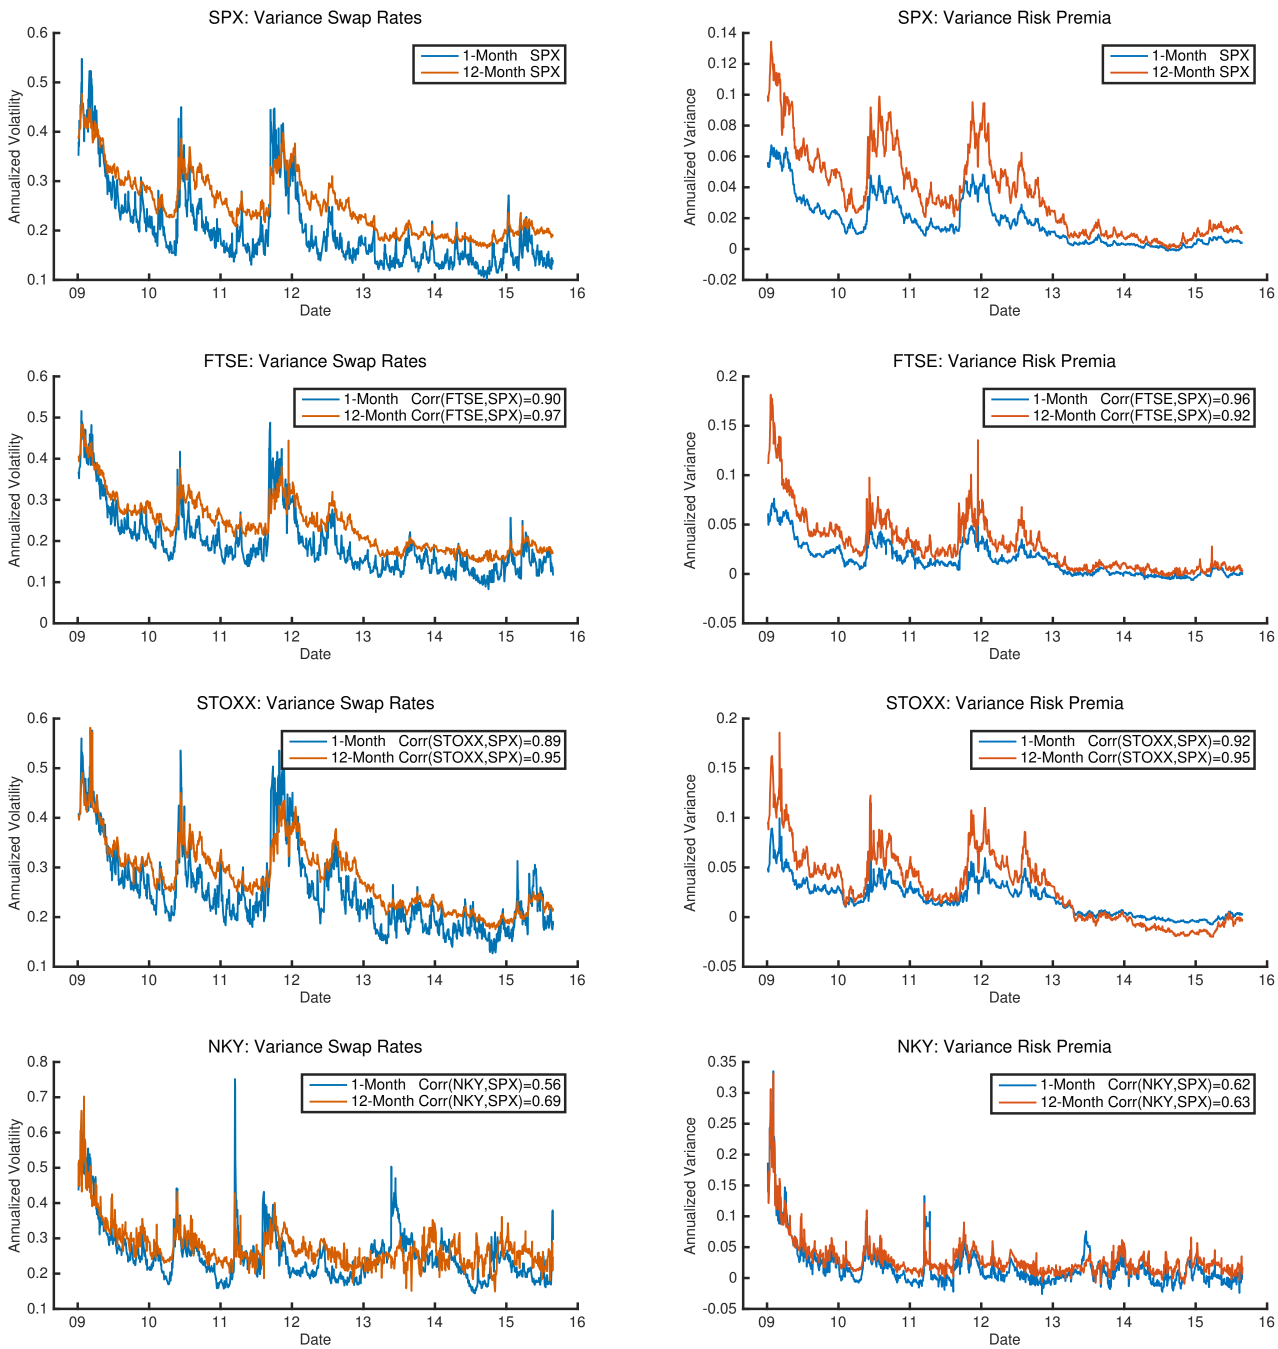


This figure plots variance swap rates and the estimated variance risk premia for the international Bloomberg data from 2009 to 2015. The annualized variance term premium for maturity $n$ is defined as $VRP_{n,t} = 12/n \cdot (E_t^{\mathbb{Q}}[\sum_{i=1}^{n} RV_{t+i}] - E_t^{\mathbb{P}}[\sum_{i=1}^{n} RV_{t+i}])$. The plots on the left report variance swap rate in volatility units. The plots on the right report the 1-month and 12-month variance term premium. The legend shows the correlation of the global data with the corresponding US data.

**Figure 8: VIX Futures Traders**

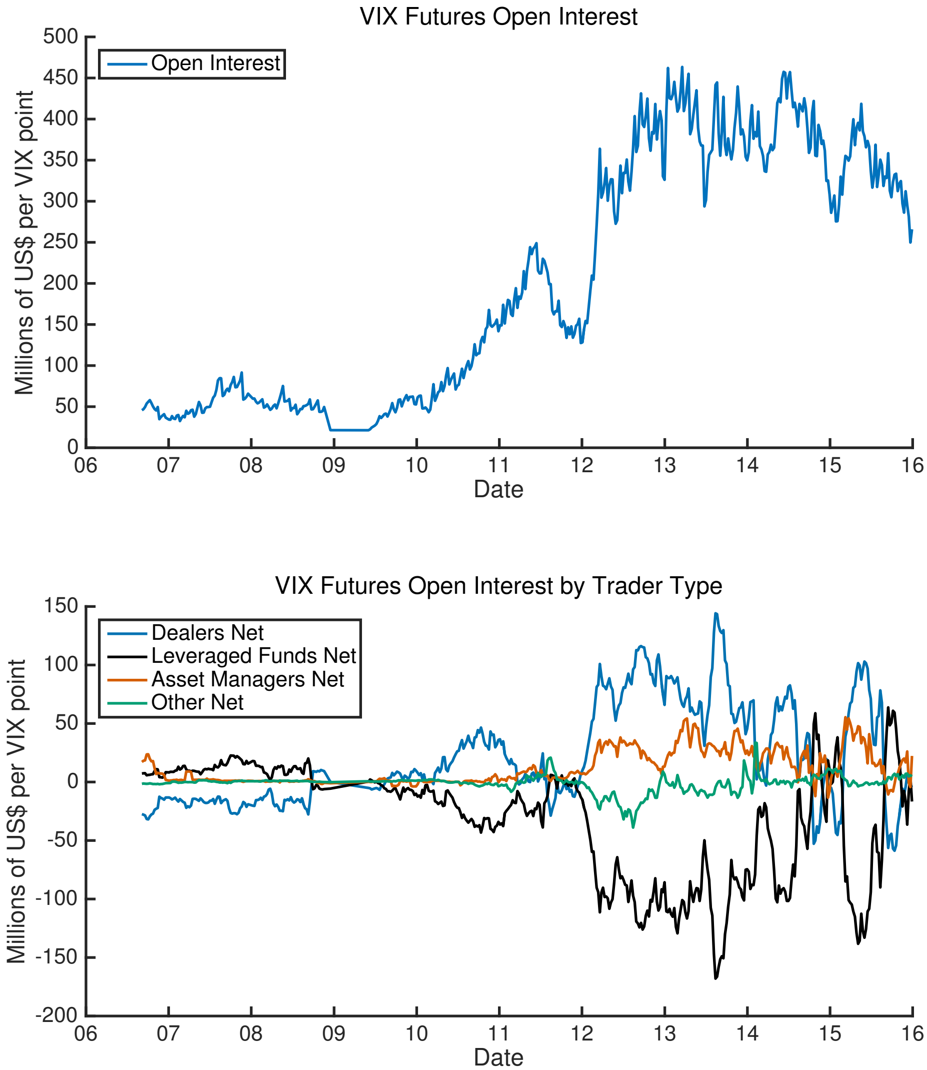


This figure plots the activity of different traders in VIX futures as reported by the Commodity Futures Trading Commission (CFTC). The CFTC provides a disaggregated report for traders in financial futures which includes the four categories: Dealer/Intermediary, Leveraged Funds, Asset Manager/Institutional and Other Reportables. Dealers represent the sell-side market participants such as banks. Leveraged funds and asset managers represent the buy-side. In particular, leveraged funds include hedge funds and commodity trading advisors whereas asset managers include pension funds, endowments, insurance companies, and mutual funds. Additional traders that report to the CFTC are included in the other category. From the plot, it is clear that VIX futures trading has increased over the sample period. Moreover, dealers and leveraged funds have been the most active market participants with largely offsetting positions. Of course, net open interest does not necessarily reflect overall volatility positions. Banks and hedge funds can also gain volatility exposure from variance swaps, options, convertibles, mortgages, and other over-the-counter derivatives. Nonetheless, the figure does highlight the important role that hedge funds, broker-dealers, and asset managers play in the volatility market as compared to other market participants.

**Figure 9: The U.S. Level Factor and Hedge Fund Drawdowns**

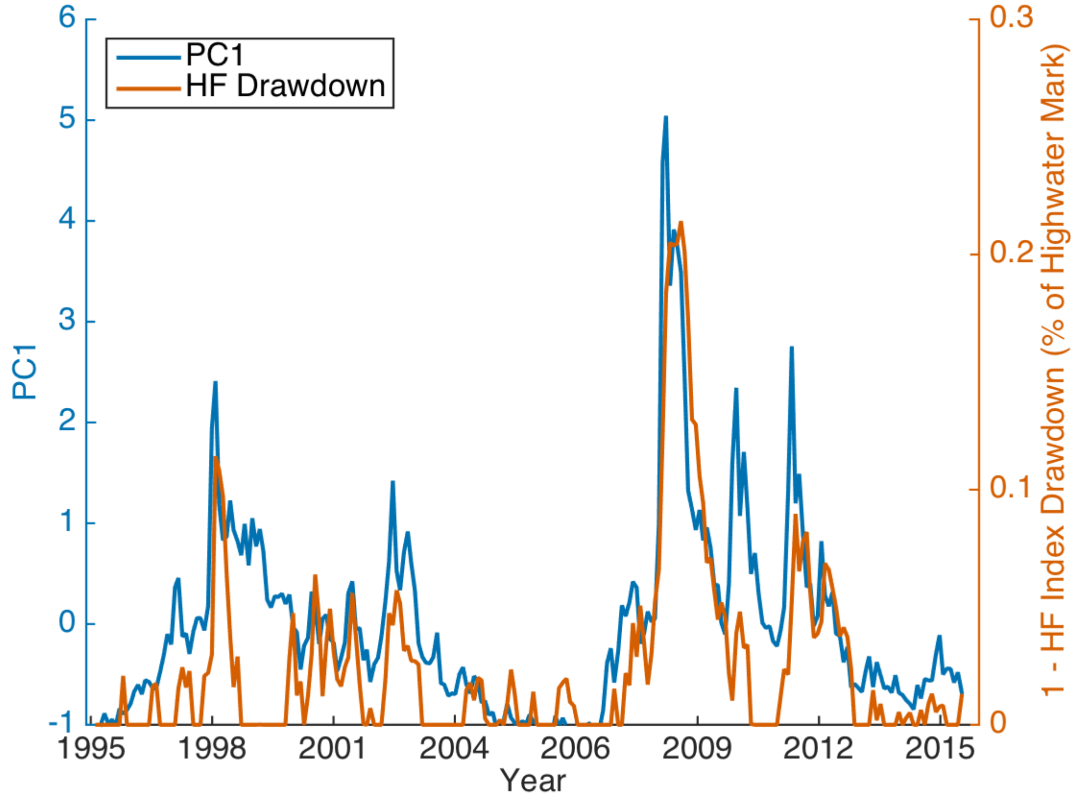


This figure plots the U.S. PC1 against hedge fund drawdowns, where the latter is defined as the percent decline of aggregate hedge fund wealth $W_t$ relative to the previous high-water mark $H_t = \max_{s \leq t} W_s$, or $1 - (W_t/H_t)$. Aggregate hedge fund wealth is measured by the HFRI Fund Weighted Composite Index considered in Jurek and Stafford (2015). The correlation between the two series is 77% using monthly observations from 1996:1 to 2015:8.

# A Appendix

## A.1 Option Summary Statistics

Table A.1 reports summary statistics for the option data used to compute our synthetic variance swap rates which are described in Section 4. For different maturity buckets, we report the average maturity, number of maturities within the bucket, number of options observed per maturity, minimum and maximum observed Black-Scholes delta and log-moneyness, root-mean-squared implied volatility fitting errors, and vega in millions. As the results indicate, the range of observed strikes is large. For maturities less than three-months we observe option prices with Black-Scholes deltas ranging from 2% to 99%. While this range decreases slightly for longer-dated maturities, we still observe 5% to 95% delta options on average. We can also see that the implied volatility fitting errors are small. Across all maturities during the early and latter halves of the sample the root-meqn-squared fitting errors are less than 1%. This indicates that our estimated implied volatility functions will provide an accurate estimate of the risk-neutral distribution. Last, we can observe that the index option market is a deep and liquid market. Toward the end of the sample we observe an average of 8.5 maturities each day that have more than 50 end-of-day quotes for out-of-the money options with over 1 billion in outstanding vega.[15]

## A.2 Comparison to Bloomberg and Markit Totem Data

Our synthetic variance swap rates closely match Bloomberg's synthetic variance swap rates and Markit Totem's over-the-counter variance swap rates.[16] To illustrate this, the top plot in Figure A.1 shows the time-series dynamics of the different variance swap rates. Throughout the sample, the correlation of these serries is over 99% for each of the twenty-four maturities. In addition, the bottom plot in Figure A.1 reports the average variance swap curves alongside the median bid-offer spread in dashed lines, which is computed from Markit's bid-offer spread data and is centered around our synthetic variance swap curve. The different curves closely align and are largely within the typical bid-offer spreads in the variance swap market. The most notable differences occur at the short-end and long-end of the curve. At the short-end, the difference likely stems from our application of the Carr and Wu (2009) approach. In contrast to the VIX construction which involves strike truncation and discretization error, we use all observed strikes to estimate an implied volatility function with log-normal tails

[15] Recall that we only use options with traditional AM settlement on the third Friday of the month. As a result, the reported outstanding vega is a lower bound for the size of the S&P 500 index option market.

[16] Markit Totem provides over-the-counter variance swap rates that are an average of dealer quotes. To obtain the Markit data on a fixed monthly grid, we interpolate between the observed maturities.

appended at the lowest and highest strike. As a result, our approach will tend to increase short-dated synthetic variance swap rates. Beyond these results, Tables A.2 and A.3 report summary statistics for the differences between our synthetic rates and the alternative data sources. The tables indicate that the maximum differences and standard deviation between the various rates are small throughout the sample.[17] Finally, as a robustness check, we report the estimated prices of risk using our synthetic data, the Bloomberg data, and the Markit data during the subsample when all of these datasets are available. Table A.4 shows that the different datasets lead to similar conclusions about which risks are priced. In addition, we can note that the primary difference between the full sample and more recent period is that the slope factor is not priced in the 2009 to 2015 subsample. In part this might be attributed to the fact that the more recent period is relatively short sample with no periods of elevated volatlity like the LTCM or financial crises, when one might expect an inverted variance swap curve to be informative about expected returns. At the same time, we do see that realized variance is priced during both periods, and that increases in realized variance decrease expected returns while increases in the level factor increase expected returns.

## A.3 Maximum Likelihood Estimation

Model estimation via three-step linear regression is asymptotically consistent and computationally efficient. It does not, however, enforce the nonlinear no-arbitrage pricing recrusions for the factor loadings, $\tilde{\beta}_n = -\tilde{B}_{n-1}$, or the internal consistency constraint that the model price the variance swap principal components exactly so that the state vector can be inferred from observed principal components. Empirically, we find that these constraints are nearly satisfied despite the fact that they are not imposed in estimation. This suggests that the regression based estimate will be similar to the maximum likelihood estimate with the constraints imposed. We verify this conjecture below.

In particular, to pursue a likelihood based approach, we assume that the variance swap return pricing errors $e_{t+1}$ are independent and normally distributed with the covariance matrix $\Sigma_e$. Using our regression based estimate as an initial condition, we then maximize the likelihood function over the risk-neutral parameters subject to the no-arbitrage pricing recursions and the internal-consistency constraint that the model price realized variance and the first two principal components of variance swaps exactly. Denoting the risk-neutral parameters as $\Theta^{\mathbb{Q}} \equiv (\mu^{\mathbb{Q}}, \Phi^{\mathbb{Q}}, \Sigma_e)$ and the physical parameters as $\Theta^{\mathbb{P}} \equiv (\mu^{\mathbb{P}}, \Phi^{\mathbb{P}}, \Sigma_v)$, Joslin

[17] The largest differences occur during the financial crisis and after the summer of 2011 when concerns about European sovereign debt led to broad declines in global equity markets and a subsequent rebound. Excluding November 2008, March 2009, September 2011, and October 2011 from the sample, the maximum difference for the one-month contract declines from 5.19% to 2.82%.

et al. (2011) show that the likelihood factors into,

$$f(R_{t+1}, X_{t+1}|X_t, \Theta) \quad = \quad f(R_{t+1}|X_{t+1}, X_t, \Theta^{\mathbb{Q}})f(X_t|X_{t-1}, \Theta^{\mathbb{P}}). \tag{32}$$

As a result of this factorization, the physical parameters can be estimated separately from the risk-neutral parameters. Since the state variables follow a vector-autoregression, it follows that the maximum likelihood estimate for the physical parameters are equal to the ordinary least squares estimate from the vector-autoregression, which implies that $\hat{\Theta}^{\mathbb{P}}$ is the same using either a regression or likelihood-based approach. To estimate the risk-neutral parameters, we maximize the likelihood function numerically and find that it converges rapidly to an estimate that is close to the initial condition from the regression-based approach.

Table A.6 reports the results. A likelihood or regression-based approach produces similar point estimates for the prices of risk. To be specific, we can only reject the null hypothesis that the regression based estimates are not equal to the maximum likelihood estimates for the third row of $\Lambda_1$ which multiplies the slope factor loadings. Moreover, these differences are not that strong, with only one parameter being significant at a 5% level and two parameters at a 10% level. From an economic perspective, we also find in unreported results that using either set of parameters produces similar fitting errors for variance swaps and variance swap returns to those reported in Table 3. Last, we confirm that the likelihood based betas are close to the regression based betas. In particular, Figure A.2 shows that the maximum likelihood betas, which impose the no-arbitrage constraints, are nearly all within the 95% confidence interval for the estimated betas from the regression-based approach.

**Table A.1: Option Summary Statistics**

This table reports summary statistics for the S&P 500 index options used to construct our synthetic variance swap rates by maturity bucket including the average maturity, number of maturities per bucket, number of options per maturity, minimum and maximum Black-Scholes delta and log-moneyness, root-mean-squared implied volatility fitting errors, and Black-Scholes vega in millions.

Full Sample

| 1996-2015 | $\tau$ | $N_\tau$ | $N_{opt}$ | $x_{min}$ | $x_{max}$ | $\Delta_{min}$ | $\Delta_{max}$ | RMSE | Vega |
|---|---|---|---|---|---|---|---|---|---|
| $\tau \in [0,3)$ | 1.63 | 2.63 | 67.70 | -0.44 | 0.13 | 0.02 | 0.99 | 0.49 | 67.56 |
| $\tau \in [3,6)$ | 4.27 | 1.32 | 47.87 | -0.74 | 0.21 | 0.02 | 0.99 | 0.40 | 95.20 |
| $\tau \in [6,12)$ | 8.92 | 2.06 | 37.79 | -0.87 | 0.27 | 0.05 | 0.97 | 0.32 | 88.31 |
| $\tau \in [12,24]$ | 17.63 | 2.00 | 35.80 | -1.08 | 0.36 | 0.08 | 0.94 | 0.30 | 78.56 |

Early Sample

| 1996-2005 | $\tau$ | $N_\tau$ | $N_{opt}$ | $x_{min}$ | $x_{max}$ | $\Delta_{min}$ | $\Delta_{max}$ | RMSE | Vega |
|---|---|---|---|---|---|---|---|---|---|
| $\tau \in [0,3)$ | 1.62 | 2.62 | 31.07 | -0.35 | 0.13 | 0.03 | 0.98 | 0.44 | 20.10 |
| $\tau \in [3,6)$ | 4.44 | 1.05 | 30.32 | -0.57 | 0.21 | 0.04 | 0.98 | 0.33 | 41.93 |
| $\tau \in [6,12)$ | 8.89 | 1.98 | 26.46 | -0.56 | 0.25 | 0.07 | 0.95 | 0.23 | 35.23 |
| $\tau \in [12,24]$ | 17.61 | 1.88 | 22.46 | -0.56 | 0.28 | 0.14 | 0.92 | 0.19 | 25.45 |

Late Sample

| 2006-2015 | $\tau$ | $N_\tau$ | $N_{opt}$ | $x_{min}$ | $x_{max}$ | $\Delta_{min}$ | $\Delta_{max}$ | RMSE | Vega |
|---|---|---|---|---|---|---|---|---|---|
| $\tau \in [0,3)$ | 1.63 | 2.65 | 105.05 | -0.53 | 0.14 | 0.01 | 0.99 | 0.55 | 115.95 |
| $\tau \in [3,6)$ | 4.16 | 1.60 | 59.66 | -0.86 | 0.21 | 0.01 | 0.99 | 0.45 | 130.97 |
| $\tau \in [6,12)$ | 8.94 | 2.14 | 48.52 | -1.16 | 0.29 | 0.03 | 0.98 | 0.42 | 138.58 |
| $\tau \in [12,24]$ | 17.64 | 2.13 | 47.65 | -1.54 | 0.43 | 0.03 | 0.97 | 0.40 | 125.76 |

**Table A.2: Synthetic Variance Swap Rates vs. Bloomberg Data**

This table reports summary statistics for the difference between our synthetic variance swap rates and Bloomberg's synthetic variance swap rates from November 2008 to August 2015 in annualized volatility percentage points.

Synthetic Variance Swap Rate - Bloomberg Rate (% annualized volatility)

| Maturity | 1 | 3 | 6 | 9 | 12 | 18 | 24 |
|---|---|---|---|---|---|---|---|
| Mean | 1.058 | 0.223 | 0.076 | 0.098 | 0.064 | -0.104 | -0.277 |
| Std. Deviation | 0.908 | 0.304 | 0.304 | 0.319 | 0.341 | 0.392 | 0.431 |
| Minimum | -0.126 | -0.429 | -1.307 | -1.092 | -1.130 | -0.970 | -1.497 |
| Median | 0.803 | 0.182 | 0.081 | 0.089 | 0.083 | -0.064 | -0.247 |
| Maximum | 5.191 | 1.423 | 1.021 | 0.884 | 0.854 | 0.777 | 0.849 |
| Skewness | 1.704 | 0.965 | -0.751 | -0.570 | -0.568 | -0.356 | -0.215 |
| Kurtosis | 7.263 | 5.252 | 8.303 | 5.141 | 4.432 | 2.707 | 3.283 |
| Correlation | 0.998 | 0.999 | 0.999 | 0.999 | 0.999 | 0.999 | 0.998 |

## Table A.3: Synthetic Variance Swap Rates vs. Markit Data

This table reports summary statistics for the difference between our synthetic variance swap rates and Markit's over-the-counter variance swap rates from September 2006 to August 2015 in annualized volatility percentage points.

Synthetic Variance Swap Rate - Markit Rate (% annualized volatility)

| Maturity | 1 | 3 | 6 | 9 | 12 | 18 | 24 |
|---|---|---|---|---|---|---|---|
| Mean | 0.897 | 0.240 | 0.187 | 0.164 | 0.195 | 0.359 | 0.461 |
| Std. Deviation | 0.667 | 0.366 | 0.395 | 0.380 | 0.381 | 0.424 | 0.512 |
| Minimum | -0.733 | -0.499 | -0.721 | -0.827 | -0.748 | -0.830 | -0.672 |
| Median | 0.886 | 0.216 | 0.148 | 0.148 | 0.171 | 0.378 | 0.480 |
| Maximum | 3.476 | 1.568 | 1.596 | 1.318 | 1.136 | 1.706 | 1.947 |
| Skewness | 0.970 | 0.645 | 0.481 | -0.037 | 0.112 | -0.336 | 0.068 |
| Kurtosis | 5.019 | 3.790 | 3.712 | 3.113 | 3.251 | 4.404 | 2.836 |
| Correlation | 0.998 | 0.999 | 0.999 | 0.999 | 0.999 | 0.998 | 0.997 |

## Table A.4: Price of Risk Estimates

This table compares price of risk estimates from our synthetic variance swap data with Bloomberg and Markit data during the recovery period from 2009 to 2015 when each of the datasets is available. As the table indicates, the price of risk estimates are similar across the different datasets. In addition, we can observe that the estimated prices of risk are relatively similar during the recovery and full sample periods.

Panel A: SVS 01/96-08/15 ($T = 236$)

| | $\Lambda_0$ | $\Lambda_{1,1}$ | $\Lambda_{1,2}$ | $\Lambda_{1,3}$ |
|---|---|---|---|---|
| $RV$ | -0.21*** | 0.45*** | -0.16*** | -0.15 |
| | [-4.29] | [4.50] | [-5.18] | [-1.40] |
| $PC1$ | -0.05 | 0.99*** | -0.31*** | -0.07 |
| | [-0.62] | [5.78] | [-5.88] | [-0.37] |
| $PC2$ | -0.02 | 0.37*** | -0.11*** | -0.29*** |
| | [-0.57] | [4.34] | [-4.19] | [-3.10] |

Panel B: SVS 01/09-08/15 ($T = 80$)

| | $\Lambda_0$ | $\Lambda_{1,1}$ | $\Lambda_{1,2}$ | $\Lambda_{1,3}$ |
|---|---|---|---|---|
| $RV$ | -0.52*** | -0.14 | -0.11** | 0.15 |
| | [-5.65] | [-0.93] | [-2.28] | [0.68] |
| $PC1$ | -0.28** | 0.93*** | -0.37*** | -0.16 |
| | [-2.00] | [3.74] | [-4.62] | [-0.44] |
| $PC2$ | 0.01 | 0.13 | -0.03 | -0.06 |
| | [0.11] | [1.23] | [-0.73] | [-0.41] |

Panel C: Bloomberg 01/09-08/15 ($T = 80$)

| | $\Lambda_0$ | $\Lambda_{1,1}$ | $\Lambda_{1,2}$ | $\Lambda_{1,3}$ |
|---|---|---|---|---|
| $RV$ | -0.39*** | -0.08 | -0.10* | 0.16 |
| | [-4.66] | [-0.54] | [-1.94] | [0.77] |
| $PC1$ | -0.36*** | 0.87*** | -0.38*** | -0.21 |
| | [-2.67] | [3.42] | [-4.63] | [-0.61] |
| $PC2$ | -0.05 | 0.11 | -0.04 | -0.06 |
| | [-0.71] | [1.01] | [-1.03] | [-0.46] |

robust t-statistics: *p<0.10, **p<0.05, ***p<0.01

Panel D: Markit 01/09-08/15 ($T = 80$)

| | $\Lambda_0$ | $\Lambda_{1,1}$ | $\Lambda_{1,2}$ | $\Lambda_{1,3}$ |
|---|---|---|---|---|
| $RV$ | -0.41*** | -0.05 | -0.11** | 0.12 |
| | [-5.42] | [-0.38] | [-2.44] | [0.60] |
| $PC1$ | -0.33** | 0.90*** | -0.37*** | -0.12 |
| | [-2.50] | [3.73] | [-4.79] | [-0.35] |
| $PC2$ | -0.06 | 0.08 | -0.03 | -0.05 |
| | [-1.38] | [1.03] | [-1.16] | [-0.46] |

robust t-statistics: *p<0.10, **p<0.05, ***p<0.01

**Table A.5: Variance Swap Return Forecasting Regressions**

This table reports analogous variance swap return forecasting regressions to Table 2 using Ibragimov and Müller (2010) t-statistics and point estimates. As before, we see that high levels of variance swap rates forecast high returns from selling variance swaps at all horizons. In unreported results, we confirm this continues to hold during the first half of the sample 1996 to 2007 and during the second half of the sample 2009 to 2015, where both subsamples exclude the financial crisis.

Panel A: Principal Components

| | $h = 1$ | $h = 3$ | $h = 6$ | $h = 9$ | $h = 12$ |
|---|---|---|---|---|---|
| $PC1$ | 0.076** | 0.057*** | 0.041** | 0.036*** | 0.032*** |
| | [3.49] | [3.91] | [3.34] | [4.62] | [4.94] |
| $PC2$ | -0.029 | -0.019 | -0.015* | -0.012 | -0.011 |
| | [-1.77] | [-1.59] | [-2.24] | [-1.71] | [-1.81] |
| $PC3$ | 0.013 | 0.009 | -0.004 | 0.003 | 0.001 |
| | [0.55] | [1.48] | [-0.85] | [1.54] | [0.40] |
| $PC4$ | -0.001 | -0.003 | 0.007 | 0.000 | 0.002 |
| | [-0.10] | [-0.63] | [1.72] | [0.05] | [0.67] |
| $PC5$ | -0.022 | -0.001 | 0.000 | -0.001 | 0.005 |
| | [-1.03] | [-0.21] | [0.01] | [-0.25] | [1.46] |

Ibragimov-Müller point estimates and t-statistics $q = 8$

Panel B: Principal Components and Forecasting Variables

| | $h = 1$ | $h = 3$ | $h = 6$ | $h = 9$ | $h = 12$ |
|---|---|---|---|---|---|
| $PC1$ | 0.151** | 0.081*** | 0.048*** | 0.021** | 0.015** |
| | [2.91] | [3.95] | [5.83] | [2.85] | [3.39] |
| $PC2$ | -0.031* | -0.000 | -0.015* | -0.010 | 0.001 |
| | [-1.90] | [-0.01] | [-2.00] | [-1.19] | [0.28] |
| $RV$ | -0.012 | -0.018 | -0.001 | 0.002 | -0.002 |
| | [-0.38] | [-1.40] | [-0.15] | [0.25] | [-0.68] |
| Term Slope | -0.041 | 0.016 | -0.015 | -0.005 | 0.005 |
| | [-1.55] | [1.19] | [-0.99] | [-0.65] | [0.75] |
| Credit | 0.009 | 0.024 | -0.007 | -0.012 | 0.008 |
| | [0.18] | [1.39] | [-0.68] | [-1.05] | [1.50] |
| DP | 0.047 | 0.025 | 0.038 | 0.039** | 0.020** |
| | [1.13] | [0.93] | [1.52] | [2.54] | [2.38] |
| IP Growth | 0.026 | -0.000 | 0.014 | 0.007 | 0.005 |
| | [0.57] | [-0.01] | [1.13] | [1.05] | [1.41] |

Ibragimov-Müller point estimates and t-statistics $q = 8$

**Table A.6: Maximum Likelihood Estimation**

This table compares maximum likelihood estimates and regression-based estimates for the prices of variance risk. As the results indicate, the estimated parameters are similar using either approach. The null hypothesis that the regression based estimates are not equal to the maximum likelihood estimates can only be rejected for one parameter at a 5% significance level and for two parameters at a 10% significance level. In terms of the economic difference, we find the model has similar fitting errors for variance swaps and variance swap returns using either approach.

Panel A: Regression-Based Prices of Risk

| | $\Lambda_0$ | $\Lambda_{1,1}$ | $\Lambda_{1,2}$ | $\Lambda_{1,3}$ |
|---|---|---|---|---|
| $RV$ | -0.2050 | 0.4493 | -0.1607 | -0.1537 |
| $PC1$ | -0.0524 | 0.9889 | -0.3075 | -0.0690 |
| $PC2$ | -0.0228 | 0.3717 | -0.1120 | -0.2885 |

Panel B: Maximum-Likelihood Prices of Risk

| | $\mu$ | $\Phi_{1,1}$ | $\Phi_{1,2}$ | $\Phi_{1,3}$ |
|---|---|---|---|---|
| $RV$ | -0.2185 | 0.6009 | -0.2005 | -0.3167 |
| $PC1$ | -0.0414 | 0.8324 | -0.2680 | 0.0784 |
| $PC2$ | -0.0189 | 0.1858 | -0.0665 | -0.1259 |

Panel C: Difference in Prices of Risk

| | $\mu$ | $\Phi_{1,1}$ | $\Phi_{1,2}$ | $\Phi_{1,3}$ |
|---|---|---|---|---|
| $RV$ | -0.0135 | 0.1516 | -0.0398 | -0.1629 |
| | [-0.28] | [1.51] | [-1.28] | [-1.48] |
| $PC1$ | 0.0111 | -0.1565 | 0.0395 | 0.1474 |
| | [0.13] | [-0.91] | [0.75] | [0.80] |
| $PC2$ | 0.0040 | -0.1859** | 0.0455* | 0.1626* |
| | [0.10] | [-2.17] | [1.70] | [1.74] |

robust t-statistics: *p<0.10, **p<0.05, ***p<0.01

**Table A.7: Robustness: Quarterly Pricing Factor Interpretations**

This table reports regressions of variance swap principal components on measures of intermediary marginal values of wealth. Hedge fund (HF) risk appetite is measured by the drawdown on the broad HFRI Fund Weighted Composite index. Broker-dealer (BD) risk appetite is measured by broker-dealer leverage obtained from the Flow of Funds. For mutual funds (MF), the ratio of equity mutual fund redemptions to total fund size measures their risk appetite. Newey-West $t$-statistics with 4 lags are reported in brackets. The sample consists of quarterly observations from 1996q1 to 2015q8.

| | US PC1 | | | | US PC2 | | | |
|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (1) | (2) | (3) | (4) |
| HF Drawdown | -0.79*** | | | -0.57*** | 0.13 | | | 0.23 |
| | [-12.58] | | | [-6.65] | [1.42] | | | [1.22] |
| BD Leverage | | 0.07 | | -0.09 | | 0.45*** | | 0.46*** |
| | | [1.01] | | [-1.49] | | [3.24] | | [3.06] |
| MF Redemption | | | 0.71*** | 0.40*** | | | 0.04 | 0.07 |
| | | | [4.31] | [3.12] | | | [0.17] | [0.24] |
| $R^2$ | 0.62 | -0.01 | 0.50 | 0.72 | 0.00 | 0.19 | -0.01 | 0.21 |

| | Global PC1 | | | | Global PC2 | | | |
|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (1) | (2) | (3) | (4) |
| HF Drawdown | -0.85*** | | | -0.42*** | 0.06 | | | 0.15 |
| | [-12.64] | | | [-4.27] | [0.37] | | | [0.88] |
| BD Leverage | | 0.03 | | -0.15* | | 0.63*** | | 0.63*** |
| | | [0.28] | | [-1.80] | | [6.80] | | [5.65] |
| MF Redemption | | | 0.88*** | 0.62*** | | | 0.11 | 0.07 |
| | | | [10.55] | [3.48] | | | [0.58] | [0.24] |
| $R^2$ | 0.72 | -0.03 | 0.77 | 0.89 | -0.03 | 0.38 | -0.02 | 0.36 |

**Table A.8: Robustness: Quarterly Intermediary Risk Appetite and VTP**

This table reports time-series regressions of variance term premia on contemporaneous measures of intermediary risk appetite. The four panels use the 1, 6, 12, and 24 month variance term premium as the LHS variable, where the VTP is obtained from our term structure model. Hedge fund (HF) risk appetite is measured by the drawdown on the broad HFRI Fund Weighted Composite index. Broker-dealer (BD) risk appetite is measured by broker-dealer leverage obtained from the Flow of Funds. For mutual funds (MF), the ratio of equity mutual fund redemptions to total fund size measures their risk appetite. Newey-West $t$-statistics with 4 lags are reported in brackets. The sample consists of quarterly observations from 1996q1 to 2015q2.

| | VTP(1) | | | | VTP(6) | | | |
|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (1) | (2) | (3) | (4) |
| HF Drawdown | -0.26* | | | -0.19 | -0.58*** | | | -0.46*** |
| | [-1.81] | | | [-1.26] | [-4.92] | | | [-4.05] |
| BD Leverage | | -0.44*** | | -0.52*** | | -0.34*** | | -0.46*** |
| | | [-5.35] | | [-6.57] | | [-3.84] | | [-5.87] |
| MF Redemption | | | 0.24 | 0.25 | | | 0.48*** | 0.32** |
| | | | [1.17] | [1.13] | | | [3.80] | [2.33] |
| $R^2$ | 0.06 | 0.19 | 0.05 | 0.31 | 0.33 | 0.10 | 0.22 | 0.56 |

| | VTP(12) | | | | VTP(24) | | | |
|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (1) | (2) | (3) | (4) |
| HF Drawdown | -0.69*** | | | -0.54*** | -0.73*** | | | -0.57*** |
| | [-7.18] | | | [-5.22] | [-8.91] | | | [-5.79] |
| BD Leverage | | -0.21* | | -0.35*** | | -0.13 | | -0.27** |
| | | [-1.88] | | [-3.58] | | [-0.97] | | [-2.44] |
| MF Redemption | | | 0.57*** | 0.34** | | | 0.61*** | 0.34** |
| | | | [4.56] | [2.63] | | | [4.41] | [2.60] |
| $R^2$ | 0.46 | 0.03 | 0.31 | 0.62 | 0.53 | 0.00 | 0.36 | 0.65 |

**Figure A.1: Synthetic Variance Swap Rates vs. Bloomberg and Markit Data**

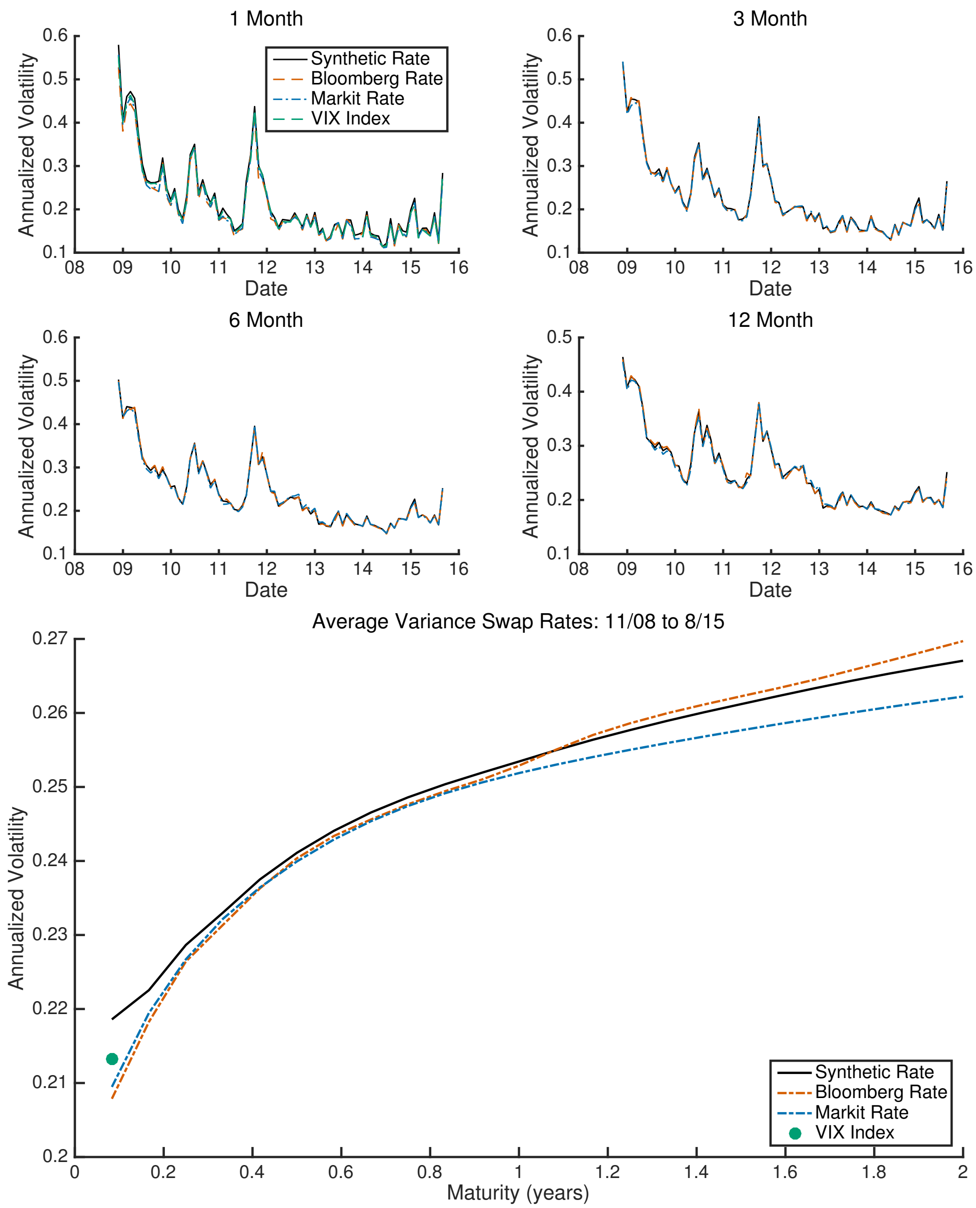


This figure plots our synthetic variance swap rates against Bloomberg synthetic rates, Markit over-the-counter rates, and the VIX index. The top plot highlights the time series dynamics. Throughout the sample, our synthetic rates move with a nearly perfect correlation to the other rates. The bottom figure plots the average variance swap curves. The different rates are close and largely within the typical bid-offer spreads in the variance swap market. In addition, we can note that our synthetic rates are slightly higher than the other rates at the short-end of the curve. This likely stems from our application of the Carr and Wu (2009) approach which uses all observed strikes to estimate an implied volatility function with log-normal tails appended at the lowest and highest strike. By way of contrast, the VIX methodology introduces strike truncation and discretization error which will tend to lower variance swap rates when a volatility smile is present.

**Figure A.2: Factor Loadings**

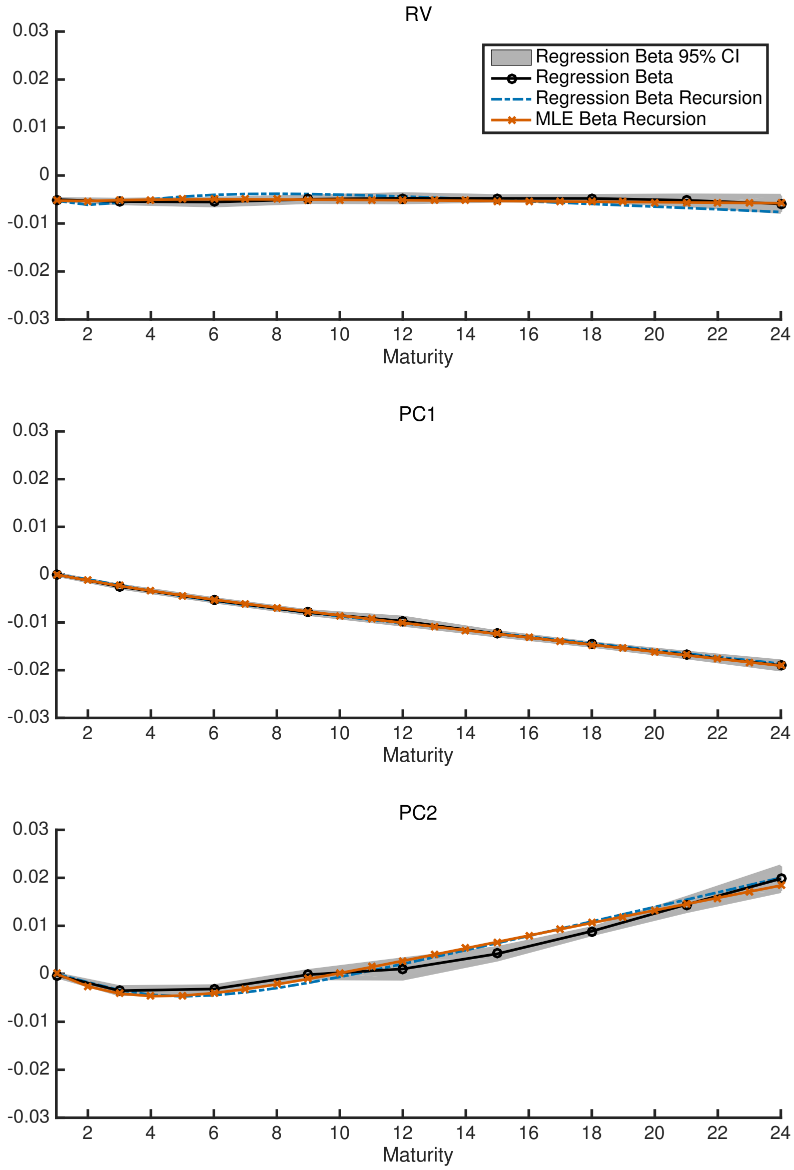


This figure plots the factor loadings estimated by regression and maximum likelihood. The grey band is a 95% pointwise confidence interval for the regression betas which are denoted with black circles. The plot then reports the recursion betas from the estimated risk-neutral parameters using the regression-based and likelihood-based approaches. The plot indicates that the recursion betas align closely with the regression betas. At nearly all of the maturities, the recursion betas are within the confidence interval indicating that we cannot reject the null hypothesis that the regression betas and recursion betas are equal.