

Federal Reserve Bank of New York
Staff Reports

The Term Structure of the Price of Variance Risk

Marianne Andries
Thomas Eisenbach
Martin Schmalz
Yichuan Wang

Staff Report No. 736
August 2015



This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

The Term Structure of the Price of Variance Risk

Marianne Andries, Thomas Eisenbach, Martin Schmalz, and Yichuan Wang

Federal Reserve Bank of New York Staff Reports, no. 736

August 2015

JEL classification: G12, G13

Abstract

We estimate the term structure of the price of variance risk (PVR), which helps distinguish between competing asset-pricing theories. First, we measure the PVR as proportional to the Sharpe ratio of short-term holding returns of delta-neutral index straddles; second, we estimate the PVR in a Heston (1993) stochastic-volatility model. In both cases, the estimation is performed separately for different maturities. We find the PVR is negative and decreases in absolute value with maturity; it is more negative and its term structure is steeper when volatility is high. These findings are inconsistent with calibrations of established asset-pricing models that assume constant risk aversion across maturities.

Key words: volatility risk, option returns, straddle, term structure

Andries: Toulouse School of Economics (e-mail: marianne.andries@tse-fr.eu). Eisenbach: Federal Reserve Bank of New York (e-mail: thomas.eisenbach@ny.frb.org). Schmalz: University of Michigan (e-mail: schmalz@umich.edu). Wang: University of Michigan (e-mail: yichuanw@umich.edu). For helpful comments and discussions, the authors thank Yakov Amihud, Markus Brunnermeier, Giovanni Barone-Adesi, Ing-Haw Cheng, Max Croce, Ian Dew-Becker, Robert Dittmar, Stefano Giglio, Ralph Koijen, Owen Lamont, Gordon Lawrence, Erik Loualiche, Thomas Mariotti, Stefan Nagel, Christian Schlag, Tyler Shumway, Adrien Verdelhan, and Grigory Vilkov, as well as seminar participants at the Tepper School of Business at Carnegie Mellon University, Maastricht University, Toulouse School of Economics, and the University of Michigan. Schmalz is grateful for generous financial support through an NTT Fellowship from the Mitsui Life Financial Research Center. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

1 Introduction

A fundamental debate in asset pricing has arisen concerning the term structure of risk premia. Well-established theoretical asset-pricing models such as [Campbell and Cochrane \(1999\)](#) and [Bansal and Yaron \(2004\)](#) predict a flat or upward-sloping term structure of excess returns; similarly, the price of variance risk is constant across maturities in standard option pricing models such as [Heston \(1993\)](#). However, [van Binsbergen, Brandt, and Koijen \(2012\)](#) and [van Binsbergen and Koijen \(2014\)](#) find that, in the data, one-period returns in equity and equity derivatives markets are actually higher for shorter maturities. Similarly, [Giglio, Maggiori, and Stroebe \(2013\)](#) show that very long-run risk premia in housing markets are low compared to observed risk prices for shorter maturities.

In response to these findings, several new asset pricing models have been developed that generate a downward-sloping term structure of equity risk premia. Most of these models enrich the underlying production economy and thus affect the expected *quantity* of risk (under the physical measure) at various horizons.¹ By contrast, [Andries, Eisenbach, and Schmalz \(2014\)](#) maintain the long-run risk endowment economy of [Bansal and Yaron \(2004\)](#) but generalize the agents' [Epstein and Zin \(1989\)](#) preferences to allow for horizon-dependent risk aversion. This framework predicts negative variance risk premia with a declining term structure (in absolute value) as a driver of the downward-sloping term structure of equity risk premia—both of which are amplified in times of high volatility. Importantly, the driver is a term structure in the *price* of variance risk. The present paper helps inform this fundamental debate by empirically investigating whether the price of variance risk has indeed a non-trivial term structure.

To investigate the price of variance risk (PVR) and its term structure, we use standard data on S&P 500 index options from February 1996 to April 2011 and estimate the PVR separately for different maturities, ranging from 11 to 252 days. We first measure Sharpe ratios of delta-neutral straddles with different maturities which are a valid qualitative measure of the PVR. We find that Sharpe ratios are negative and large (in absolute value) for short maturities, but they are much closer to zero at longer maturities. This finding indicates a sharply decreasing term structure for the price of variance risk (in absolute value).

For an estimation that enables a cleaner and more robust interpretation—in particular in light of potentially time-changing prices of risk—we then adapt the maximum-likelihood approach of [Christoffersen, Heston, and Jacobs \(2013\)](#) to estimate the PVR parameter separately for options of different maturities and find results consistent with our non-parametric

¹See the literature review below. [Van Binsbergen and Koijen \(2015\)](#) give a comprehensive review of the empirical and theoretical research on the term structure of risk premia.

Sharpe-ratio analysis. From the shortest maturities, between 11 and 30 days, to the longest maturities, between 230 and 250 days, the PVR drops by 44 percent and over half of that drop occurs going from the 11–30 day bucket to the 30–50 day bucket. Furthermore, higher levels of volatility are associated with more negative prices of variance risk—especially at shorter maturities, resulting in a steeper term structure of the PVR. Our findings thus suggest that the known fact of a negative overall PVR is predominantly driven by short maturities and by periods of high market volatility.

The present paper contributes to the literature as follows. Possibly guided by the predictions of existing option-pricing models such as [Heston \(1993\)](#), which predict a constant price of variance risk across maturities, no paper to date in the options literature has investigated if variance risk prices have a non-trivial term structure. For example, work by [Coval and Shumway \(2001\)](#) or [Carr and Wu \(2009\)](#) measures variance risk premia for options with a single maturity; [Christoffersen et al. \(2013\)](#) pool all maturities to estimate the price of variance risk. [Choi, Mueller, and Vedolin \(2015\)](#) find a negative and upward-sloping term structure of variance premia in the Treasury futures market. Given our finding that estimating a [Heston \(1993\)](#) model separately for different maturities rejects its assumption of a flat term structure, the next generation of option-pricing models would benefit from allowing risk prices to vary depending on maturity.

Outside the options literature, other papers have investigated the term structure of variance risk premia and prices, using different data sets and different methodologies than the present paper. Most recently, [Dew-Becker, Giglio, Le, and Rodriguez \(2014\)](#) use proprietary data on variance swaps to estimate term-structure models, similar to [Amengual \(2008\)](#) and [Ait-Sahalia et al. \(2012\)](#), but add realized volatility as a third factor to a standard level-and-slope analysis. They find that only shocks to realized volatility are priced, implying a term structure that is steeply negative at the short end (a one-month horizon) but essentially flat at zero beyond that. Both methodologies we employ, as well as our data, are different from and complementary to [Dew-Becker et al. \(2014\)](#). Given the importance of the empirical question for asset pricing, we find it valuable to provide support from an entirely different and relatively easy-to-understand estimation approach—the one we use is unique to the literature. In terms of results, we also find a strong concavity in the term structure, but we measure a negative price of variance risk for all maturities. In addition, we offer more granular estimates (daily maturity buckets) and include shorter maturities (11 days versus 1 month).

Our conditional results on the relationship between current market volatility and the term structure of risk prices are related to the work of [Cheng \(2014\)](#) who studies the returns of hedging volatility with VIX futures. Cheng documents that hedging is cheaper during

turbulent times, whereas we find that the price of variance risk is more negative and that its term structure is steeper when current volatility is high. [Barras and Malkhozov \(2015\)](#) find differences in estimates of variance risk premia in the equity and option markets that are driven by institutional factors. While this finding suggests a potential explanation for the differences between our results and those of [Dew-Becker et al. \(2014\)](#) as well as those of [Cheng \(2014\)](#), it also emphasizes the value of using different methodological approaches and different data sets to approach an academic understanding of the market for volatility risk.

Our findings have implications for asset pricing models also outside the options literature. In particular, our results suggest a preference-based explanation to the downward-sloping term structure of equity risk premia. While the long-run-risk model of [Bansal et al. \(2013\)](#) as well as the rare-disaster model of [Wachter \(2013\)](#) correctly predict a negative price per unit of variance risk, the models cannot quantitatively match its decline with maturity (in absolute value). Consumption-based asset pricing models with loss aversion, such as [Andries \(2012\)](#) and [Curatola \(2014\)](#), predict a pricing per unit of risk that declines intrinsically (in absolute value) with the quantity of risk, consistent with the evidence on markets where the declines in Sharpe ratios in the term-structure are accompanied by increases in volatility (see [van Binsbergen and Koijen, 2015](#) for examples). However, our results highlight a decline in both the pricing and quantity of risk in the term-structure and cannot be simply rationalized by first-order risk aversion.

The paper proceeds as follows. Section 2 presents the theoretical derivation of the price of variance risk in the [Heston \(1993\)](#) model as well as its relation to the Sharpe ratios of short-term returns of delta-neutral straddles and our parametric estimation procedure. Section 3 gives the empirical results. Section 4 concludes.

2 Hypotheses Development and Empirical Strategy

2.1 Theoretical Background and Empirical Hypotheses

We use the structure of the option-pricing model of [Heston \(1993\)](#) to isolate the role of variance risk. Specifically, we assume stock price S_t and variance v_t satisfy the following physical dynamics:

$$\begin{aligned} dS_t &= \mu S_t dt + \sqrt{v_t} S_t dW_{1t} \\ dv_t &= \kappa (\theta - v_t) dt + \sigma \sqrt{v_t} dW_{2t} \end{aligned} \tag{1}$$

The stock return has drift μ and volatility $\sqrt{v_t}$. The variance v_t itself has long-run mean θ , to which it reverts at speed κ , and volatility $\sigma\sqrt{v_t}$. Both dW_{1t} and dW_{2t} are Brownian motions and ρ denotes the correlation between shocks to the return and variance processes.

To identify the premia for equity risk and variance risk, we can risk-neutralize the dynamics in (1) as follows:

$$\begin{aligned} dS_t &= \mu^* S_t dt + \sqrt{v_t} S_t dW_{1t}^* \\ &= r S_t dt + \sqrt{v_t} S_t dW_{1t}^* \\ dv_t &= \kappa^* (\theta^* - v_t) dt + \sigma \sqrt{v_t} dW_{2t}^* \\ &= (\kappa (\theta - v_t) - \lambda v_t) dt + \sigma \sqrt{v_t} dW_{2t}^* \end{aligned}$$

The standard intuition is that to compensate for equity risk, the stock return under the physical measure has a drift with a premium $\mu - r$ compared to the risk-free rate r . Similarly, to compensate for variance risk, the variance under the physical measure has a drift with a premium λv_t compared to the risk-neutral drift. Alternatively, the physical variance dynamic has a lower long-run mean, $\theta < \theta^*$, and faster mean-reversion, $\kappa > \kappa^*$ for a negative variance risk premium, $\lambda v_t < 0$.

Our main interest is to study if and how the compensation investors demand for variance risk depends on the horizon and what drives this dependence. Since the variance risk premium λv_t depends on current variance v_t —which varies in the time series—we focus our analysis on the parameter λ and refer to it as the ‘price of variance risk’ (PVR). Inspired by the existing evidence on the term-structure of risk premia, we test three hypotheses:

1. *The PVR is negative at all maturities.*
2. *The PVR decreases in absolute value with maturity.*
3. *The PVR is more negative and its term structure is steeper when volatility is high.*

The first prediction is consistent with various established asset pricing models, including [Bansal and Yaron \(2004\)](#). The latter two predictions are specific to the model by [Andries et al. \(2014\)](#). We now explain the two different estimation procedures we use to test these hypotheses: a non-parametric estimation using short-horizon Sharpe ratios and a parametric estimation based on [Christoffersen et al. \(2013\)](#).

2.2 Non-parametric Estimation: Short-horizon Sharpe Ratios

We show how the short-horizon Sharpe ratios of delta-neutral straddles identify the sign of the PVR and the slope of its term structure. In the Heston model, the no-arbitrage price X_t

of any option satisfies the following partial differential equation:

$$\begin{aligned} \frac{1}{2} \frac{\partial^2 X}{\partial S^2} v_t S_t^2 + \frac{\partial^2 X}{\partial S \partial v} \rho \sigma S_t v_t + \frac{1}{2} \frac{\partial^2 X}{\partial v^2} \sigma^2 v_t + \frac{\partial X}{\partial S} r S_t \\ + \frac{\partial X}{\partial v} [\kappa (\theta - v_t) - \lambda v_t] - r X_t + \frac{\partial X}{\partial t} = 0 \end{aligned} \quad (2)$$

The option price X_t therefore follows a dynamic given by:

$$\begin{aligned} dX_t = \left[\frac{\partial X}{\partial v} \lambda v_t + \left(X_t - \frac{\partial X}{\partial S} S_t \right) r + \frac{\partial X}{\partial S} \mu S_t \right] dt \\ + \frac{\partial X}{\partial S} S_t \sqrt{v_t} dW_{1t} + \frac{\partial X}{\partial v} \sigma \sqrt{v_t} dW_{2t} \end{aligned} \quad (3)$$

A complication arises with the measurement of λ because μ is not observable. To address this challenge, we form a measurable portfolio of straddles that are delta neutral so that the portfolio is independent of μ .² To that end, first note that we can use the stock-return dynamic (1) to rewrite the option dynamic (3) as:

$$d \left(X_t - \frac{\partial X}{\partial S} S_t \right) = \left[\left(X_t - \frac{\partial X}{\partial S} S_t \right) r + \frac{\partial X}{\partial v} \lambda v_t \right] dt + \frac{\partial X}{\partial v} \sigma \sqrt{v_t} dW_{2t} \quad (4)$$

Next, we can discretize the dynamic in (4) and rearrange to arrive at

$$\lambda \frac{\sqrt{v_t \Delta t}}{\sigma} + \varepsilon = \frac{\Delta \left(X_t - \frac{\partial X}{\partial S} S_t \right) - \left(X_t - \frac{\partial X}{\partial S} S_t \right) r \Delta t}{\frac{\partial X}{\partial v} \sigma \sqrt{v_t \Delta t}}, \quad (5)$$

where $\varepsilon = \Delta W_{2t} / \sqrt{\Delta t}$ which is zero in expectation. Note that the denominator on the right hand side of equation (5) is just the standard deviation of the process in equation (4). Hence, when X_t is a delta-neutral straddle, we have

$$\begin{aligned} E \left[\lambda \frac{\sqrt{v_t \Delta t}}{\sigma} \right] &\approx \frac{E \left[\Delta \left(X_t - \frac{\partial X}{\partial S} S_t \right) \right] - \left(X_t - \frac{\partial X}{\partial S} S_t \right) r \Delta t}{\sqrt{\text{Var} \left[\Delta \left(X_t - \frac{\partial X}{\partial S} S_t \right) \right]}} \\ &\approx \frac{E[\Delta X_t] - X_t r \Delta t}{\sqrt{\text{Var}[\Delta X_t]}} \\ &= \text{SR}(X_t) \end{aligned} \quad (6)$$

²Delta-neutral straddles are not necessarily at the money. While at-the-money straddles are approximately delta neutral for short maturities, the delta-neutral moneyness increases with maturity; see the decreasing ratios of S_t/K in the delta neutral straddles described in Table 1. Following the literature, we compute the delta-neutral portfolios using the Black and Scholes model.

The expected PVR λ therefore differs from the Sharpe ratio $\text{SR}(X_t)$ of a delta-neutral straddle only by a factor of $\sqrt{v_t \Delta t}/\sigma$.

As a result, the Sharpe ratios of delta-neutral straddles are a qualitatively valid measure of both the sign and relative magnitude of the PVR across maturities, even though they are not quantitatively comparable to the results from the parametric estimation we present in section 2.3.³ In contrast to our approach, [Coval and Shumway \(2001\)](#) look at returns from holding one-month delta-neutral straddles to maturity. The long holding period means they cannot use the discretization necessary for equation (6) to hold. The straddles analyzed by [van Binsbergen and Koijen \(2015\)](#) have deltas that increase with maturity, and thus depart from the delta neutrality required by equation (6).

The instantaneous Sharpe ratio of investing in delta-neutral straddles can be estimated by

$$\text{SR} = \frac{E[\Delta X_t/X_t - r\Delta t]}{\sqrt{\text{Var}[\Delta X_t/X_t - r\Delta t]}},$$

where

$$\frac{\Delta X_t}{X_t} = \frac{X_{t+\Delta t} - X_t}{X_t}.$$

We estimate the Sharpe ratios of options with different maturities ranging from 11 days to 252 days, using daily returns. To estimate the Sharpe ratio SR^τ for options with maturity τ , we use returns from options with maturities in the range $[\tau, \tau + 20)$ and compute the average divided by the standard deviation of such returns. Figure 1 shows that these returns are not auto-correlated over time. Therefore, asymptotic standard errors for the Sharpe ratios can be computed by bootstrapping, treating each return as an independent observation. The results of our analysis are described and discussed in Section 3.

2.3 Parametric Estimation Procedure

The factor $\sqrt{v_t \Delta t}/\sigma$ in Equation (6), while constant in the term-structure, may vary in the time series. These time series variations can be correlated—and we show in Section 3 that they are—with variations in the slope of the PVR term-structure. Such covariation can potentially introduce a bias into the magnitude of the estimated slope in the Sharpe ratio analysis described above. This concern motivates us to also estimate the parameter λ directly in a parametric model, using a discrete-time method based on [Christoffersen, Heston, and Jacobs \(2013, hereafter CHJ\)](#). CHJ estimate their model using a sample of options pooled

³Because the factor $\sqrt{v_t \Delta t}/\sigma$ is guaranteed to be positive, the Sharpe ratio is a robust test of the sign of the PVR. Moreover, the extra factor $\sqrt{v_t \Delta t}/\sigma$ does not change with maturity, so it does not affect the sign of the slope of the term structure of the PVR.

across different maturities and strike prices. We adapt the procedure to subsets of at-the-money options and run the estimation of λ separately for options of different maturities and volatility levels.⁴ We first describe the economic intuition and then explain the formal estimation procedure.

CHJ discretize the continuous-time dynamic of stock return and variance in (1) using approach of [Heston and Nandi \(2000\)](#) where the stock follows a GARCH process and the one-period excess return has variance h_t ; the variance itself follows an ARMA(1, 1) process:⁵

$$\begin{aligned}\log S_t &= \log S_{t-1} + r_t + \left(\eta - \frac{1}{2}\right) h_t + \sqrt{h_t} z_t \\ h_t &= \omega + \beta h_{t-1} + \alpha \left(z_{t-1} - \gamma \sqrt{h_{t-1}}\right)^2\end{aligned}$$

with $z_t \sim \mathcal{N}(0, 1)$. Assuming a pricing kernel with equity risk aversion ϕ and variance risk aversion ξ , CHJ show that the processes can be risk-neutralized as

$$\begin{aligned}\log S_t &= \log S_{t-1} + r_t - \frac{1}{2} h_t^* + \sqrt{h_t^*} z_t^*, \\ h_t^* &= \omega^* + \beta h_{t-1}^* + \alpha^* \left(z_{t-1}^* - \gamma^* \sqrt{h_{t-1}^*}\right)^2,\end{aligned}$$

with

$$\begin{aligned}h_t^* &= \frac{1}{1 - 2\alpha\xi} h_t, & \omega^* &= \frac{1}{1 - 2\alpha\xi} \omega, \\ \alpha^* &= \frac{1}{1 - 2\alpha\xi} \alpha, & \gamma^* &= \gamma - \phi,\end{aligned}$$

and $z_t^* \sim \mathcal{N}(0, 1)$. The difference between physical and risk-neutral processes is intuitively analogous to the continuous-time case. To compensate for variance risk, the physical variance process has a lower long-run mean and lower persistence for $\xi > 0$. The only notable difference is that over a discrete time interval, there is a difference in the contemporaneous levels of physical variance h_t and risk-neutral variance h_t^* while in continuous time there is only one instantaneous variance v_t .

Conditional on the physical GARCH parameters $\Theta = \{\omega, \beta, \alpha, \gamma, \eta\}$, a value of the param-

⁴In principle, the maximum-likelihood estimation (MLE) can be applied to options that are not at the money. However, we restrict the MLE to at-the-money options to limit the effect of jumps on the estimation since jumps have a much larger effect on the price of out-of-the-money options.

⁵[Heston and Nandi \(2000\)](#) show that this discrete time model nests the continuous time model of [Heston \(1993\)](#) as a special case when the number of trading periods per unit of physical time goes to infinity. Therefore this GARCH approach is precisely the discrete time analogue of the continuous time [Heston \(1993\)](#) model.

eter ξ generates risk-neutral volatilities h_t^* that can be used to price options.⁶ We therefore perform the estimation in two stages: In the first stage, we estimate the parameters Θ governing the GARCH process in index returns. In the second stage, we use this set of common GARCH parameters to estimate the PVR separately for subsets of options by maturity and volatility state.^{7,8}

For the first stage, we estimate the GARCH parameters through maximum likelihood. Using daily series on index returns $R_t = \log(S_t/S_{t-1})$ and the risk-free rate r_t , we solve

$$\hat{\Theta} = \underset{\Theta=\{\omega,\beta,\alpha,\gamma,\eta\}}{\operatorname{argmin}} \left\{ -\frac{1}{2} \sum_{t=1}^T \left[\log h_t + \frac{(R_t - r_t - (\eta - \frac{1}{2}) h_t)^2}{h_t} \right] \right\},$$

where

$$\begin{aligned} h_t &= \omega + \beta h_{t-1} + \alpha \left(z_{t-1} - \gamma \sqrt{h_{t-1}} \right)^2, \\ z_t &= \frac{R_t - r_t - (\eta - \frac{1}{2}) h_t}{\sqrt{h_t}}, \\ h_1 &= \frac{\omega + \alpha}{1 - \beta - \alpha \gamma^2}. \end{aligned}$$

For the second stage, given a value of $\hat{\Theta}$ from the first stage and a particular subset of option prices $\{P_i\}_{i=1}^N$, we estimate ξ through maximum-likelihood estimation

$$\hat{\xi} = \underset{\xi}{\operatorname{argmin}} \left\{ -\frac{1}{2} \sum_{i=1}^N \left(\log \hat{s}_\varepsilon^2 + \frac{\varepsilon_i^2}{\hat{s}_\varepsilon^2} \right) \right\},$$

where we treat the Black-Scholes Vega (BSV) weighted pricing errors as Gaussian random

⁶Given the physical parameters Θ and the variance risk aversion ξ , the equity risk aversion ϕ is pinned down as $\phi = -(\eta - \frac{1}{2} + \gamma)(1 - 2\alpha\xi) + \gamma - \frac{1}{2}$. The details of the option-pricing model come from [Heston and Nandi \(2000\)](#) and are replicated in the appendix.

⁷In our approach, we do not smooth the inputs by computing a volatility surface. Instead, we smooth the outputs from the estimation procedure. This ensures that we are basing our estimates on actual observed prices and that we do not inflate our dataset with interpolated values.

⁸CHJ show that a joint maximum-likelihood procedure with both options and returns gives estimates comparable to those of a procedure that estimates the models sequentially with returns first and options second. The sequential procedure is particularly important in our case because the options all derive value from the same underlying time series for stock returns, so it makes sense for them to share the same time-series parameters.

variables, following the method of CHJ:

$$\begin{aligned}\hat{s}_\varepsilon^2 &= \frac{1}{N} \sum_{i=1}^N \varepsilon_i^2 \\ \varepsilon_i &= \frac{P_i^{\text{Mkt}} - P_i^{\text{Mod}}(\xi)}{\text{BSV}_i^{\text{Mkt}}}\end{aligned}$$

We then derive the continuous-time PVR λ following CHJ by calibrating it to obtain the same unconditional variance of stock returns and the same ratio between physical and risk-neutral unconditional variances as in the discrete-time model:

$$\lambda = -\kappa \frac{E^*[h_t^*] - E[h_t]}{E^*[h_t^*]},$$

where

$$\begin{aligned}\kappa &= (1 - \beta - \alpha\gamma^2) \times 252, \\ E[h_t] &= \frac{\omega + \alpha}{1 - \beta - \alpha\gamma^2}, \\ E^*[h_t^*] &= \frac{\omega^* + \alpha^*}{1 - \beta - \alpha^*\gamma^{*2}}.\end{aligned}$$

To test the different hypotheses, we perform the second stage on several subsets of the data:

1. We estimate λ by considering the prices of options in maturity buckets ranging from 11 to 250 to see if the PVR changes across the term structure.
2. We split the options into two regimes for current volatility. Doing so enables a first look into how the term structure of the price of variance risk changes in high-volatility periods (high h_t) and calm periods (low h_t).

The results of our analysis are described and analyzed in Section 3.

3 Data and Empirical Results

3.1 Data Sources and Summary Statistics

We use daily closing data from February 1996 to April 2011 of European SPX index options and SPX index levels from OptionMetrics. Value-weighted S&P 500 returns, excluding dividends, from January 1990 to December 2014 come from CRSP. The three-month risk-free

rate data are taken from FRED. The risk-free rate for a given daily observation is defined as $\log(1 + r_m)/252$, where r_m is the risk-free rate recorded for the last week of the previous month.

We clean the data by removing duplicate observations of calls or puts on the same day that have the same expiration date, strike price, and midprice. Next, we keep only options that have a maturity between 11 and 252 trading days on the day of observation.⁹ We exclude shorter-maturity options to avoid microstructure noise close to expiration affecting our results, and we exclude longer-maturity options because they are thinly traded.

For the non-parametric estimation, for each maturity and strike on a given day, we estimate the Black-Scholes implied volatility by the average of the call and put Black-Scholes implied volatilities. We then use this implied volatility to estimate the Black-Scholes delta of the call and the put at that strike and maturity observation on that day. We then pick the strike and maturity such that the straddle delta, which is the sum of the put and call deltas, is closest to zero. We drop observations that have straddle deltas greater than 0.10 in absolute value and that have bid ask spreads greater than 10 percent of the midprice.¹⁰ As such, the options under consideration are highly liquid and close to delta neutral. We also follow Bakshi et al. (1997) in excluding any options that do not obey the futures arbitrage constraints.¹¹

We further restrict our sample to options that satisfy the delta constraint and have a maturity between 11 and 252 days during the *entire* $[-1, +1]$ -day window relative to the observation date for the Sharpe ratio analysis. If a given option contract violates an arbitrage bound or goes out of the money in the $[-1, +1]$ -day window, then its return is not used in the calculations.¹² Hence, the Sharpe-ratio analysis excludes options in periods when the index changed dramatically in the span of 3 days, thereby excluding crisis periods. We thus ensure that abnormal events do not drive our results. We keep only calls and puts that can be paired into a straddle.

For the parametric estimation, for each year and each maturity bin of 10 days starting at every 10th day, we drop the observations corresponding to the top and bottom 1 percent of residuals in a third-order polynomial regression of the option price against the GARCH

⁹Using trading days to measure maturity is essential. The GARCH estimation treats the index return series as a continuous series without weekends. Thus, to be consistent, the option maturities should also be expressed in trading days.

¹⁰We find that the results are not sensitive to changing the straddle delta threshold to a lower value of 0.05. Sample sizes decrease substantially, however. We therefore don't focus the analysis on that reduced sample.

¹¹For a call with maturity τ , $C(\tau) \geq \max\{0, S_t - X_t e^{-r\tau}\}$, and for a put, $P(\tau) \geq \max\{0, X_t e^{-r\tau} - S_t\}$.

¹²We do not need to make such corrections for the parametric estimation because the parametric estimation fits prices, not returns.

volatility variable h_t .¹³ We do this instead of restricting the options to be delta neutral in the $[-1, +1]$ -day window. The results are not quantitatively sensitive to the exact level of truncation.

We present summary statistics for the sample of 47,416 option-day observations in Table 1. We note in Table 2 that the dollar value of the bid-ask spread increases along the maturity structure but decreases as a percentage of the option price. We view this observation as an indication of good liquidity across the entire term structure—one of the benefits of studying index returns as opposed to individual-name returns.

3.2 Non-parametric Estimation Results

We provide the “model-free” estimation of Sharpe ratios of straddle returns, interpreted as the sign of intercept and slope of the term structure of the PVR in Figure 3. We present the point estimates for maturity buckets of length 20 days in Table 3.

The term structure of Sharpe ratios is concave and trends upwards at almost all maturities. Between the first two maturity buckets, the 11–30 day maturity bucket and the 30–50 day maturity bucket, the Sharpe ratio increases from -1.15 to -0.71 . This sharp increase represents 40 percent of the overall range in Sharpe ratios over the entire term structure, showing that most of the variation stems from the short end. The Sharpe ratio continues to slope upwards, albeit more slowly, for maturities beyond 50 days. It is -0.54 for intermediate maturities 50–70 days, more than three times as negative as for the 230–250 day straddles, -0.16 . It steadily approaches zero for longer maturities.¹⁴

Our findings indicate that existing measures of the negative price of variance risk in the literature, if obtained from a pooled sample, are mainly driven by short maturities. Moreover, our results are qualitatively consistent with those reported in [van Binsbergen and Koijen \(2015\)](#), although their straddles are not necessarily delta-neutral (which can potentially bias the results) and the Sharpe ratios are only reported for a small number maturities.

As noted in Section 2.2, the factor $\sqrt{v_t \Delta t} / \sigma$ is not constant in the time series so interaction with the sample size may introduce bias into the magnitude of the estimated slope. As this issue is not present in the parametric estimation, it can account for different slope estimates across the analyses. Further, differences in liquidity between the delta-neutral straddles and the at-the-money options used in the non-parametric and parametric analysis, respectively, may introduce differences.

¹³Put prices are converted to equivalent call prices by put-call parity.

¹⁴Because shorting short-maturity straddles generates very high expected returns also when the positions are crash-hedged ([Coval and Shumway, 2001](#)), one-sided crash risk is unlikely to be the driver of our Sharpe-ratio results (see also [Constantinides et al., 2013](#)).

3.3 Parametric Estimation Results

We present the parametric estimation of the term structure of the PVR in Figure 4, and a selection of data points grouped by maturity bucket in Table 4.¹⁵ Our results show that the PVR decreases in absolute value with maturity. For example, the point estimate is -0.61 for maturities 11–30, and -0.34 for maturities 230–250.¹⁶ This result confirms our first two key hypotheses: the PVR is negative, and decreases in absolute value with maturity.

Interestingly, we observe a dip in the unconditional term structure at the 70–90 day maturity bucket, consistent with results obtained from the alternative estimation procedure employed in Dew-Becker et al. (2014). We attribute this anomaly to a change in the distribution of option maturities before and after 2007. As seen in Figure 2, prior to 2007 most traded options had maturities between 0 and 60 days. After 2007, however, the maturity range for most traded options increases to 0–90 days. Our results for the pricing of variance risk in the 70–90 day maturity bucket are thus artificially driven by the post 2007 period. Consistent with our hypothesis for why the dip exists, we find that the dip disappears when we split our sample into pre- and post-2007 subsamples (see Figure 5 and Table 5). Once we split by time period, the term structures are smoothly concave and upward sloping at all maturities, in line with our Sharpe ratio results. In both time periods, most of the change in the PVR occurs in the 11–50 day maturity range, with only 30 percent of the overall term-structure variation occurring at the intermediate maturities, between 50 and 250 days.

To test the third hypothesis, we explore how volatility levels affect the term structure of the PVR. We divide all days into two categories of expected future volatility, based on whether h_{t+1} from the GARCH estimation is above or below the sample median, and then run the previous procedure on each subsample. We present our results in Figure 6 and Table 6. Compared to the low volatility state, the PVR in the high volatility state is more negative and the term structure is steeper. The economic magnitude of these differences is substantial. For the shortest maturities, the point estimate is -1.15 in high volatility states, which compares to -0.31 in the low volatility states. For the longest maturities, the point estimates are -0.58 and -0.13 , respectively. To ensure robustness of this result to a procedure that does not “break” options within their maturities, we alternatively split the sample into a period that includes the beginning of the sample until 2007 and a post-2007

¹⁵We only show second-stage results for brevity. The first-stage GARCH estimation yields $\omega = 0$, $\beta = 0.835$, $\alpha = 3.54 \times 10^{-6}$, $\eta = 3.48$, and $\gamma = 191.03$. These results are in line with estimates from Christoffersen et al. (2013).

¹⁶We do not correct for variations in $\hat{\Theta}$ in the asymptotic MLE standard errors: variations in $\hat{\Theta}$ might increase the variance of the estimate $\hat{\xi}$ on any given subset, but do not imply we are overstating differences in ξ between subsets, which is the measure we are interested in. For a more detailed discussion on how we correct the asymptotic MLE standard errors, see the appendix.

subsample, where the latter is a higher-volatility period. We find qualitatively similar results, presented in Figure 7 and Table 7. We conclude that the third hypothesis, a lower PVR and steeper term structure, finds robust support in the data as well.

4 Conclusion

We provide estimates of the price of variance risk at various horizons, first, by measuring model-free Sharpe ratios of straddle returns with varying maturities and, second, by estimating the price of variance risk in a [Heston \(1993\)](#) model, based on the empirical approach developed by [Christoffersen et al. \(2013\)](#). We find the price of insurance against increases in volatilities varies with the horizon of the risk insured: short-term insurance is more expensive than long-term insurance, and this effect is more pronounced in times of higher volatility.

These results extend the accumulating evidence for non-trivial term structures of risk prices to the market for variance risk. A comparative advantage to the literature is a focus on the price of risk as a driver of the term structure of risk premia. The findings thus help motivate a new generation of option pricing models that allow for horizon-dependent risk prices. However, our findings are informative not only for option pricing. Specifically, the results presented in this paper support preference-based rationalizations of the term-structure of expected returns, such as the horizon-dependent risk aversion model of [Andries et al. \(2014\)](#).

The implicit assumption that risk prices are flat across horizons—which is rejected in this paper—would lead market observers to attribute too much of the term structure of risk premia to a term structure in expected volatility. In other words, our results emphasize that the conversion between objective and risk-neutral measures depends on maturity. This finding may help inspire future generations of asset pricing models and econometricians’ interpretation of economic forecasts.

References

- Ait-Sahalia, Y., M. Karaman, and L. Mancini (2012). The term structure of variance swaps, risk premia, and the expectation hypothesis. Working Paper.
- Amengual, D. (2008). The term structure of variance risk premia. Working Paper.
- Andries, M. (2012). Consumption-based asset pricing with loss aversion. Working Paper.
- Andries, M., T. M. Eisenbach, and M. C. Schmalz (2014). Asset pricing with horizon-dependent risk aversion. Working Paper.
- Bakshi, G., C. Cao, and Z. Chen (1997, December). Empirical performance of alternative option pricing models. *Journal of Finance* 52(5), 2003–2049.
- Bansal, R., D. Kiku, I. Shaliastovich, and A. Yaron (2013). Volatility, the macroeconomy, and asset prices. *Journal of Finance* (forthcoming).
- Bansal, R. and A. Yaron (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance* 59(4), 1481–1509.
- Barras, L. and A. Malkhozov (2015). Does variance risk have two prices? Evidence from the equity and option markets. Working Paper.
- 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. and R. Koijen (2014). Real excess volatility. Working Paper.
- van Binsbergen, J. and R. Koijen (2015). The term structure of returns: Facts and theory. Working Paper.
- Campbell, J. Y. and J. H. Cochrane (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy* 107(2), 205–251.
- Carr, P. and L. Wu (2009). Variance risk premiums. *Review of Financial Studies* 22(3), 1311–1341.
- Cheng, I.-H. (2014). The returns to fear. Working Paper.
- Choi, H., P. Mueller, and A. Vedolin (2015). Bond variance risk premiums. Working Paper.

- Christoffersen, P., S. Heston, and K. Jacobs (2013). Capturing option anomalies with a variance-dependent pricing kernel. *Review of Financial Studies* 26(8), 1962–2006.
- Constantinides, G. M., J. C. Jackwerth, and A. Savov (2013). The puzzle of index option returns. *Review of Asset Pricing Studies* 3(2), 229–257.
- Coval, J. D. and T. Shumway (2001). Expected option returns. *Journal of Finance* 56(3), 983–1009.
- Curatola, G. (2014). Loss aversion, habit formation and the term structures of equity and interest rates. Working Paper.
- Dew-Becker, I., S. Giglio, A. Le, and M. Rodriguez (2014). The price of variance risk. Working Paper.
- Epstein, L. G. and S. E. Zin (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, 937–969.
- Giglio, S., M. Maggiori, and J. Stroebe (2013). Very long-run discount rates. Working Paper.
- Heston, S. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies* 6(2), 327–343.
- Heston, S. and S. Nandi (2000). A closed-form GARCH option valuation model. *Review of Financial Studies* 13(3), 585–625.
- Wachter, J. A. (2013, June). Can time-varying risk of rare disasters explain aggregate stock market volatility? *Journal of Finance* 52(3), 987–1035.

Figures




Figure 1: Daily straddle returns used in Sharpe ratio analysis. Daily option and index price data from February 1996 to April 2011 come from OptionMetrics. Each dot represents the net arithmetic return of a straddle-day observation. Each facet of the plot contains options of different maturities, labeled at the top of each facet.




Figure 2: Calendar of options by maturity and date. Daily option and index price data from February 1996 to April 2011 come from OptionMetrics. Each dot represents an option-day observation. The vertical axis is the maturity, and the horizontal axis has the days of the year. Diagonal lines reflect the fact that not every maturity is traded every day, and that certain maturities are only observed on certain calendar days.




Figure 3: Estimates of the Sharpe ratio of delta-neutral straddle returns SR^τ . Daily option and index price data from February 1996 to April 2011 come from OptionMetrics. The Sharpe ratio SR^τ for options with maturity τ is computed by collecting all returns from options with a maturity in the interval $[\tau, \tau + 20)$ and then dividing the sample mean by the sample standard deviation. Dotted lines mark 95 percent confidence intervals formed by the 2.5 and 97.5 percentiles of 10,000 bootstrap estimates, and the solid line is the mean of such estimates.




Figure 4: Parametric estimation of the term structure of the price of variance risk λ^τ . Daily option and index price data from February 1996 to April 2011 come from OptionMetrics. Return data from January 1990 to December 2014 come from CRSP. The first maturity bucket includes options with maturities of 11 through 29 trading days. Each subsequent point at maturity τ represents the estimation results on options with maturity $[\tau, \tau + 20)$. The last bucket contains options with maturities 230 through 250. Dotted lines mark asymptotic 95 percent confidence intervals conditional on the given realization of GARCH parameters from the second-stage maximum-likelihood estimation.




Figure 5: Parametric estimation of the term structure of the price of variance risk λ^τ by time period. Daily option and index price data from February 1996 to April 2011 come from OptionMetrics. Return data from January 1990 to December 2014 come from CRSP. The first maturity bucket includes options with maturities of 11 through 29 trading days. Each subsequent point at maturity τ represents the estimation results on options with maturity $[\tau, \tau + 20)$. The last bucket contains options with maturities 230 through 250. Dotted lines mark asymptotic 95 percent confidence intervals conditional on the given realization of GARCH parameters from the second-stage maximum-likelihood estimation.




Figure 6: Parametric estimates of the price of variance risk λ^τ at different states of forecasted GARCH volatility h_{t+1} . Daily option and index price data from February 1996 to April 2011 come from OptionMetrics. Error bars mark asymptotic 95 percent confidence intervals conditional on the given realization of GARCH parameters from the second-stage maximum-likelihood estimation.




Figure 7: Parametric estimates of the price of variance risk λ^r at different states of forecasted GARCH volatility h_{t+1} and by different time periods. Daily option and index price data from February 1996 to April 2011 come from Option-Metrics. Error bars mark asymptotic 95 percent confidence intervals conditional on the given realization of GARCH parameters from the second-stage maximum-likelihood estimation.

Tables

Table 1: Summary statistics of options used in parametric estimation by maturity. Daily option and index price data from February 1996 to April 2011 come from OptionMetrics. A price on a day is defined as the midprice between the closing best bid and best ask. Maturity is defined as the number of days from the observation date to expiration. All maturity ranges are inclusive on the left and exclusive on the right. Midprice is the average between the best closing bid price and the best closing ask price on a particular day. S_t/K refers to the average ratio of the underlying stock price to the strike price of the option taken within that maturity category in percentage points. Bid-Ask Spread is the difference between the best bid and best offer on a given day. Bid-Ask Ratio is a percentage computed as $\frac{\text{Bid-Ask Spread}}{\text{Midprice}} \times 100$. N refers to the total number of calls and puts at each maturity. All statistics except the observation count are computed as arithmetic means over option-day observations.

Maturity	N	S_t/K	Calls			Puts		
			Midprice (\$)	Bid-Ask Spread (\$)	Bid-Ask Ratio (%)	Midprice (\$)	Bid-Ask Spread (\$)	Bid-Ask Ratio (%)
11–30	3,042	99.9	24.45	1.51	6.30	24.08	1.50	6.17
30–50	3,547	99.7	34.15	1.91	5.63	35.00	1.88	5.7
50–70	3,180	99.4	41.98	2.12	4.87	44.73	2.11	5.19
70–90	1,876	99.1	50.26	2.33	4.19	56.95	2.29	4.71
90–110	1,500	98.6	52.94	2.34	3.81	62.28	2.29	4.46
110–130	1,479	98.4	57.63	2.39	3.53	68.66	2.33	4.17
130–150	1,494	98.1	62.66	2.36	3.21	75.10	2.34	3.91
150–170	1,510	97.8	66.31	2.44	3.06	81.37	2.39	3.75
170–190	1,488	97.5	69.65	2.56	3.00	86.20	2.54	3.77
190–210	1,485	97.2	73.19	2.60	2.87	91.88	2.59	3.67
210–230	1,469	96.9	76.32	2.64	2.78	96.84	2.58	3.54
230–252	1,638	96.6	78.36	2.78	2.75	101.64	2.73	3.62
Total	47,416		52.07	2.2	4.68	98.58	2.23	4.22

Table 2: Liquidity (bid-ask ratio) of options of various maturities and moneyness. All intervals are inclusive on the left and exclusive on the right. Bid-Ask Ratio is computed as $\frac{\text{Bid-Ask Spread}}{\text{Current Price}} \times 100$. Data come from the full universe of OptionMetrics data from February 1996 to April 2011, after cleaning for duplications.

Maturity	Black-Scholes Delta			
	0–1%	1–5%	5–10%	> 10%
0–11	159.78	77.50	26.64	35.90
11–30	53.46	33.78	12.03	17.21
30–50	7.14	10.56	8.57	12.40
50–70	5.25	5.47	6.05	8.51
70–90	4.57	4.53	4.63	6.54
90–110	3.95	4.19	4.31	5.47
110–130	3.92	3.83	3.88	4.86
130–150	3.54	3.55	3.54	4.41
150–170	3.36	3.40	3.41	4.16
170–190	3.34	3.34	3.37	4.04
190–210	3.27	3.24	3.32	3.87
210–230	3.26	3.15	3.18	3.70
230–252	3.23	3.25	3.25	3.64

Table 3: Point estimates of daily expected returns, standard deviation of returns, and Sharpe ratios of straddles of various maturities. The expected value and standard deviation are computed on log arithmetic returns multiplied by 5. Confidence intervals are also included for Sharpe ratios. Daily option and index price data from February 1996 to April 2011 come from OptionMetrics. N refers to the number of straddle-day return observations used to compute the Sharpe ratio. All maturity buckets are inclusive on the left and exclusive on the right. 95 percent confidence intervals are formed by the 2.5 and 97.5 percentiles of 10,000 bootstrap estimates, and the final estimate reported is the mean of such bootstrap trials.

Maturity τ	N	$E[r]$	$\sigma(r)$	SR_τ	95% CI	
					Lower Bound	Upper Bound
11–30	749	-3.23	2.80	-1.15	-1.27	-1.05
30–50	1729	-1.63	2.32	-0.71	-0.77	-0.65
50–70	1865	-1.09	2.02	-0.54	-0.59	-0.49
70–90	1235	-0.78	1.83	-0.43	-0.49	-0.36
90–110	1039	-0.70	1.82	-0.39	-0.45	-0.32
110–130	1159	-0.56	1.66	-0.34	-0.4	-0.27
130–150	1215	-0.40	1.67	-0.24	-0.3	-0.18
150–170	1262	-0.41	1.69	-0.24	-0.3	-0.18
170–190	1282	-0.38	1.49	-0.25	-0.31	-0.19
190–210	1322	-0.33	1.55	-0.21	-0.27	-0.16
210–230	1333	-0.24	1.54	-0.15	-0.21	-0.10
230–250	1359	-0.25	1.53	-0.16	-0.22	-0.11

Table 4: Point estimates and confidence intervals of the unconditional price of variance risk λ^τ from the parametric estimation. Daily option and index price data from February 1996 to April 2011 come from OptionMetrics. An estimate of λ^τ is estimated by maximum likelihood in order to best price the subset of options in a given maturity bucket. N refers to the number of options observations used to compute λ^τ . All maturity buckets are inclusive on the left and exclusive on the right. Standard errors are derived from the delta method applied to the mapping between ξ^τ and λ^τ .

Maturity τ	N	λ^τ	Standard Error
11–30	5932	-0.61	0.03
30–50	6912	-0.47	0.02
50–70	6200	-0.46	0.02
70–90	3658	-0.55	0.02
90–110	2920	-0.40	0.02
110–130	2878	-0.38	0.02
130–150	2908	-0.39	0.02
150–170	2940	-0.37	0.02
170–190	2896	-0.37	0.02
190–210	2890	-0.37	0.02
210–230	2858	-0.36	0.02
230–250	2920	-0.34	0.02

Table 5: Point estimates and confidence intervals of the unconditional price of variance risk λ^τ from the parametric estimation by time period. Daily option and index price data from February 1996 to April 2011 come from OptionMetrics. An estimate of λ^τ is estimated by maximum likelihood in order to best price the subset of options in a given maturity bucket. N refers to the number of options observations used to compute λ^τ . All maturity buckets are inclusive on the left and exclusive on the right. Standard errors are derived from the delta method applied to the mapping between ξ^τ and λ^τ .

Maturity τ	1996–2007			2008–2011		
	N	λ^τ	Standard Error	N	λ^τ	Standard Error
11–30	4554	-0.41	0.03	1378	-1.41	0.11
30–50	5232	-0.30	0.02	1680	-1.09	0.06
50–70	4332	-0.27	0.02	1868	-0.94	0.04
70–90	1900	-0.25	0.02	1758	-0.91	0.04
90–110	1970	-0.21	0.02	950	-0.82	0.04
110–130	1904	-0.20	0.02	974	-0.77	0.04
130–150	1956	-0.19	0.02	952	-0.84	0.05
150–170	1978	-0.19	0.02	962	-0.77	0.04
170–190	1886	-0.18	0.02	1010	-0.73	0.04
190–210	1908	-0.17	0.02	982	-0.77	0.04
210–230	1900	-0.17	0.02	958	-0.78	0.04
230–250	1898	-0.16	0.02	1022	-0.70	0.03

Table 6: Point estimates and confidence intervals of the conditional price of variance risk λ^τ from the parametric estimation. Daily option and index price data from February 1996 to April 2011 come from OptionMetrics. Volatility quartiles are formed based on the first-stage GARCH estimation parameter h . An estimate of λ^τ is estimated by maximum likelihood in order to best price the subset of options in a given maturity bucket. N refers to the number of options observations used to compute λ^τ . All maturity buckets are inclusive on the left and exclusive on the right. Standard errors are derived from the delta method applied to the mapping between ξ^τ and λ^τ .

Maturity τ	Low Volatility (Below Median)			High Volatility (Above Median)		
	N	λ^τ	Standard Error	N	λ^τ	Standard Error
11–30	2991	-0.31	0.03	2941	-1.15	0.07
30–50	3590	-0.19	0.02	3322	-0.93	0.04
50–70	3180	-0.19	0.02	3020	-0.87	0.03
70–90	1681	-0.24	0.02	1977	-0.89	0.04
90–110	1379	-0.13	0.03	1541	-0.69	0.03
110–130	1537	-0.15	0.02	1341	-0.71	0.04
130–150	1438	-0.14	0.02	1470	-0.67	0.03
150–170	1399	-0.11	0.02	1541	-0.65	0.03
170–190	1500	-0.15	0.02	1396	-0.64	0.03
190–210	1448	-0.14	0.02	1442	-0.62	0.03
210–230	1368	-0.10	0.02	1490	-0.62	0.03
230–250	1448	-0.13	0.02	1472	-0.58	0.03

Table 7: Point estimates and confidence intervals of the conditional price of variance risk λ^τ from the parametric estimation by time period. Daily option and index price data from February 1996 to April 2011 come from OptionMetrics. Volatility categories are formed based on whether the first-stage GARCH estimation parameter h_t is above or below its median over the full sample from 1996 to 2011. An estimate of λ^τ is estimated by maximum likelihood in order to best price the subset of options in a given maturity bucket. N refers to the number of options observations used to compute λ^τ . All maturity buckets are inclusive on the left and exclusive on the right. Standard errors are derived from the delta method applied to the mapping between ξ^τ and λ^τ .

Maturity τ	1996–2007				2008–2011			
	Low Volatility		High Volatility		Low Volatility		High Volatility	
	λ^τ	Std Error	λ^τ	Std Error	λ^τ	Std Error	λ^τ	Std Error
11–30	-0.20	0.03	-0.84	0.05	-0.83	0.06	-2.06	0.19
30–50	-0.09	0.02	-0.70	0.03	-0.65	0.04	-1.51	0.09
50–70	-0.06	0.02	-0.62	0.03	-0.54	0.03	-1.34	0.07
70–90	-0.03	0.03	-0.56	0.03	-0.56	0.03	-1.21	0.06
90–110	0.02	0.03	-0.50	0.03	-0.54	0.04	-1.05	0.07
110–130	-0.02	0.02	-0.48	0.04	-0.46	0.03	-1.08	0.06
130–150	-0.01	0.02	-0.44	0.03	-0.54	0.03	-1.06	0.07
150–170	0.03	0.03	-0.46	0.02	-0.50	0.04	-0.95	0.06
170–190	-0.01	0.02	-0.42	0.03	-0.46	0.03	-0.97	0.05
190–210	0.00	0.02	-0.41	0.03	-0.52	0.03	-0.96	0.06
210–230	0.03	0.02	-0.42	0.02	-0.51	0.03	-0.94	0.05
230–250	0.02	0.02	-0.37	0.03	-0.44	0.03	-0.91	0.05

Appendix

Option Pricing in Discrete Time under a GARCH Model of the Underlying

Appendix B of [Christoffersen et al. \(2013\)](#) gives a closed-form solution for the call price. Let t denote the current trading day, and let T denote the future trading day on which the option expires. Let S_t be the current stock price. Observe that the option price depends on both the particular current and future days, and not just the maturity, because there is dependence on the current state of volatility h_{t+1}^* :

$$\begin{aligned}
C_i^{\text{Mod}} &= C_t(S_t, h_{t+1}^*, K, T) \\
&= S_t P_1(t) - K \exp(-r(T-t)) P_2(t) \\
P_1(t) &= \frac{1}{2} + \frac{\exp(-r(T-t))}{\pi} \int_0^\infty \Re \left[\frac{K^{-i\varphi} g_{t,T}^*(i\varphi + 1)}{i\varphi S(t)} \right] d\varphi \\
P_2(t) &= \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \Re \left[\frac{K^{-i\varphi} g_{t,T}^*(i\varphi)}{i\varphi} \right] d\varphi \\
g_{i,T}^*(\varphi) &= \exp \{ \varphi \log S_t + A_{t,T}(\varphi) + B_{t,T}(\varphi) h_{t+1}^* \} \\
A_{t,T} &= A_{t+1,T}(\varphi) + \varphi r + B_{t+1,T}(\varphi) \omega^* - \frac{1}{2} \log(1 - 2B_{t+1,T}(\varphi) \alpha^*) \\
A_{T,T} &= 0 \\
B_{t,T} &= -\frac{1}{2} \varphi + B_{t+1,T}(\varphi) \beta + B_{t+1,T}(\varphi) \alpha^* (\gamma^*)^2 \\
&\quad + \frac{\frac{1}{2} \varphi^2 + 2B_{t+1,T}(\varphi) \alpha^* \gamma^* (B_{t+1,T}(\varphi) \alpha^* \gamma^* - \varphi)}{1 - 2B_{t+1,T}(\varphi) \alpha^*} \\
B_{T,T} &= 0 \\
h_{t+1}^* &= \omega^* + \beta h_t^* + \alpha^* \left(z^*(t) - \gamma^* \sqrt{h_t^*} \right)^2 \\
&= \frac{h_t}{1 - 2\alpha\xi}.
\end{aligned}$$

The put price is computed using put-call parity. Observe that delta is just $P_1(t)$.

In the parametric estimation, we use Black-Scholes vegas, which come from inverting the

Black-Scholes formula:

$$\begin{aligned}\hat{\sigma} &= \arg \min_{\sigma} \left(S_t \Phi(d_1) - K e^{r(T-t)} \Phi(d_2) - C^{\text{Mkt}} \right)^2 \\ d_1 &= \frac{\log \frac{S_t}{K} + \left(r + \frac{1}{2} \sigma^2 \right) \tau}{\sigma \sqrt{\tau}} \\ d_2 &= d_1 - \sigma \sqrt{\tau} \\ \text{BSV}_t^{\text{Mkt}} &= \sqrt{\tau} S_t \Phi'(d_1) \\ \Phi &= \text{CDF of Standard Normal.}\end{aligned}$$

Standard Errors

There are two potential reasons the standard errors computed from the parametric estimation may be too narrow:

1. There is uncertainty in the estimation of the GARCH parameters and the sequential estimation procedure does not propagate this uncertainty into the standard errors for λ .
2. The parametric procedure treats the pricing errors as independent over time, which may not be the case if there are unobserved shocks that cause the prices of all options of a certain expiration date to move together.

There are two reasons why the uncertainty in the GARCH parameters is unlikely to have a quantitatively large effect on our results. First, the time period for the GARCH estimation is very long—it covers over 6,301 trading days—so any uncertainty would be very small. Second, recall that the key mechanism that determines $\hat{\xi}$ is trying to match the wedge between the physical variance implied by the GARCH process and the risk-neutral variance implied by the options prices. Although variation in the GARCH parameters may change the size of the wedge and therefore increase the variance of the estimate $\hat{\xi}$ on any given subset, that variation does not imply we are overstating differences in $\hat{\xi}$ between subsets.

To address the problems due to correlated pricing errors in the time dimension, observe that in a given maturity window of 20 trading days there are up to 20 repeated observation of contracts with a given expiration date. [Christoffersen et al. \(2013\)](#) for example consider weekly options in their estimation of ξ and interpret the resulting log likelihood as the result of i.i.d. pricing errors. If we adopt this standard convention then we divide our number of observations by 4 and therefore the standard error when computed under an assumption of i.i.d. pricing errors should be multiplied by 2. These are the standard errors reported in the

plots and figures for the parametric estimation. Note that no similar correction needs to be made for the Sharpe ratios as those look at returns, which would still be independent if there were a persistent shock that raised prices.