

NO. 736
AUGUST 2015

REVISED
APRIL 2024

The Term Structure of the Price of Variance Risk

Marianne Andries | Thomas M. Eisenbach | R. Jay Kahn |
Martin C. Schmalz

The Term Structure of the Price of Variance Risk

Marianne Andries, Thomas M. Eisenbach, R. Jay Kahn, and Martin C. Schmalz

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

August 2015; revised April 2024

JEL classification: G12, G13

Abstract

We empirically investigate the term structure of variance risk pricing and how it varies over time. We estimate the aversion to variance risk in a stochastic-volatility option pricing model separately for options of different maturities and find that variance risk pricing decreases in absolute value with maturity but remains significantly different from zero up to the nine-month horizon. We find consistent non-parametric results using estimates from Sharpe ratios of delta-neutral straddles. We further show that the term structure is downward sloping both during normal times and in times of stress, when required compensation for variance risk increases and its term structure steepens further.

Key words: volatility risk, option returns, term structure

Eisenbach: Federal Reserve Bank of New York (email: thomas.eisenbach@ny.frb.org). Andries: University of Southern California (email: andries@usc.edu). Kahn: Board of Governors of the Federal Reserve System (email: jay.kahn@frb.gov). Schmalz: University of Oxford, CEPR, and ECGI (email: martin.schmalz@sbs.ox.ac.uk). For helpful comments and discussions, the authors thank Yakov Amihud, Giovanni Barone-Adesi, Markus Brunnermeier, 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 Econometric Society World Congress (Montréal), Toulouse Financial Econometrics Conference, Society for Economic Dynamics Conference, CMU Tepper, Maastricht University, Toulouse School of Economics, and the University of Michigan. They also thank Shruthi Nagar and Lulu Wang for excellent research assistance. Schmalz is grateful for generous financial support through an NTT Fellowship from the Mitsui Life Financial Center.

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 author(s) and do not necessarily reflect the position of the Federal Reserve Bank of New York, the Board of Governors of the Federal Reserve System, or the Federal Reserve System. Any errors or omissions are the responsibility of the author(s).

To view the authors' disclosure statements, visit
https://www.newyorkfed.org/research/staff_reports/sr736.html.

1 Introduction

The risks facing investors have a rich term structure with some risks relevant at shorter maturities and other risks relevant at longer maturities. After mostly focusing on bonds, the analysis of how risks at different maturities are priced has more recently shifted to equities, providing new moments for asset pricing models to match (Cochrane, 2017). In particular, van Binsbergen, Brandt, and Kojien (2012) first document a downward-sloping term structure of equity risk premia, suggesting significantly higher risk premia for short horizons than for long horizons and rejecting the predictions of the workhorse asset pricing models of Campbell and Cochrane (1999) and Bansal and Yaron (2004). While similar patterns have since been found in various asset classes, the evidence is still subject to debate.¹ Using equity dividend strips, Bansal et al. (2021) show evidence that the term structure of equity risk premia is weakly upward sloping during normal times and inverts during recessions. Giglio, Kelly, and Kozak (2023) confirm this using synthetic dividend strips going back to the 1970s. Gormsen (2021) shows that the slope of the term structure is countercyclical to the price-dividend ratio. These results highlight the importance of analyzing both unconditional slopes as well as time-series variations in the term structure of risk premia.

Because variance risk is an important driver of equity risk premia, how it is priced in the term-structure is important to understand the results above. Dew-Becker et al. (2017) show that the market for variance swaps indicates a negative price of risk at short maturities, but allows investors to insure against variance risk for free for longer maturities, suggesting that the compensation for variance risk has an important term structure component. The goal of our paper is to deepen our understanding of the term structure of variance risk pricing and its variation over time with standard index options data, using both the structural framework of a parametric options pricing model as well as the robustness of a non-parametric approach.

For our parametric estimation, we use the stochastic variance model with a variance-dependent stochastic discount factor of Christoffersen, Heston, and Jacobs (2013) which yields a single parameter capturing the pricing of variance risk due to variance risk aversion. Our approach is, first, to discipline the physical parameters of the return process using the underlying index returns and, then, to estimate the parameter governing the pricing of variance risk separately for options in distinct maturity buckets. This allows us to study whether long horizon variance risk is priced and whether the relative pricing at different maturities varies systematically over time.

¹See van Binsbergen and Kojien (2016) and Giglio, Kelly, and Kozak (2023) for overviews.

For our unconditional estimates, we find that the term structure of variance risk pricing is downward sloping (in absolute value), with investors in short-maturity options exhibiting significantly higher variance risk aversion than investors in longer maturity options. Across our maturity buckets, investors display highly significant variance risk aversion at 1 to 3 months (first bucket) and at 3 to 6 months (second bucket). However, variance risk aversion decreases significantly between the first and second bucket and between the second and third bucket such that it is only marginally significant at 6 to 9 months (third bucket) and indistinguishable from zero at 9 to 12 months (fourth bucket). Even when we allow the other model parameters to also vary across maturity buckets, we find that only the variance risk aversion parameter shows a systematic term structure.

In light of the fact that our parametric results contrast with the model's implicit assumption of a unique aversion to variance risk at all maturities, we also provide a non-parametric estimate of the term structure of variance risk pricing. Specifically, we study Sharpe ratios of delta-neutral index straddles which, as we show, proxy for the model's variance risk premium parameter up to a scaling factor. We find that the straddles confirm the results from our parametric estimation: Sharpe ratios decline in absolute value with maturity and are significantly negative for the first two maturity buckets and then indistinguishable from zero starting at 6 to 9 months (bucket 3).

Next, we examine time-series variation in the term structure of variance risk pricing. We find a robust negative slope to the compensation for variance risk across periods with high and low volatility, low and high GDP growth and low and high price-dividend ratios. During "bad times", investors demand significantly more compensation for variance risk across all maturities, such that even variance risk aversion at our longest horizon (9 to 12 months) becomes significantly positive. This result contrasts with the model's implicit assumption that the volatility risk premium scales linearly with current volatility and, instead, suggests a convex relationship. However, the slope of the term-structure of variance risk pricing always remains downward sloping, in good as well as in bad times, and in stark contrast to similar analysis on the term-structure of the equity risk pricing, e.g. in [Gormsen \(2021\)](#), [Bansal et al. \(2021\)](#), and [Giglio, Kelly, and Kozak \(2023\)](#).

Finally, we compare our results directly to those in [Dew-Becker et al. \(2017\)](#). Using proprietary data on variance swap prices, they derive variance forwards as non-parametric estimates of the risk neutral expectation of variance at different horizons. Taking the difference to realized variance (expected variance under the physical measure) to quantify the variance risk premium, [Dew-Becker et al. \(2017\)](#) find non-zero risk pricing only at the front end of the term-structure. Both our parametric and non-parametric estimates show that variance shocks are priced at least up to the medium horizon (3 to 6 months, sec-

ond bucket). For a proper comparison, we derive the pricing of variance forwards in our framework. The expected variance forward at a given horizon is determined by a weighted average of the expected variance under the physical measure and the expected variance under the risk-neutral measure, where the weight on the risk-neutral measure increases with the horizon. Any stochastic volatility model with strictly positive variance risk premium thus implies a gradually increasing term structure, anchored to the low physical mean of variance at the short end and converging to the higher risk-neutral expectation at longer maturities.

Because [Dew-Becker et al. \(2017\)](#) document a flat term-structure, except at very short maturities, they conclude that the aversion to variance shocks, other than transitory shocks, is zero. However, this interpretation assumes that variance risk aversion is the same at all horizons. If, instead, the risk-neutral expected variance is calculated separately for each maturity bucket using the corresponding variance risk aversion parameters obtained in our parametric estimation then we obtain an implied term structure of variance forwards strikingly similar to the one documented by [Dew-Becker et al. \(2017\)](#): a large jump from average realized variance to the average variance forward at the shortest maturity and then essentially flat at longer maturities. Letting variance risk pricing decrease with the horizon may thus allow to reconcile the classical asset pricing stochastic volatility framework with the evidence in [Dew-Becker et al. \(2017\)](#), in the same way that the horizon-dependent risk aversion model of [Andries, Eisenbach, and Schmalz \(2023\)](#) reconciles the long-run risk model of [Bansal and Yaron \(2004\)](#) with the term-structure evidence of [van Binsbergen, Brandt, and Koijen \(2012\)](#).²

Our three main results, (i) that variance risk is priced beyond the short horizon, (ii) that the term-structure of variance risk premia is downward sloping in absolute value, and (iii) that it remains so under all market conditions, suggest the classical option pricing model should be extended to let variance risk aversion vary with the horizon, with potentially important implications for asset pricing models. Our first result is consistent with standard asset pricing models (e.g. [Bansal and Yaron \(2004\)](#); [Bansal et al. \(2012, 2014\)](#)) where shocks to future volatility play a key role in matching asset pricing data such as the equity premium. On the other hand, our second and third results, that the average term-structure of variance risk premia is downward sloping and that it remains so under all market conditions, challenges both the standard models as well as models that capture the upward/downward variation in the slope of equity risk premia via variation in risks, such as [Gormsen \(2021\)](#) and [Bansal et al. \(2021\)](#). Overall, our results show the need

²Alternatively, the physical dynamics of the model could be enriched, e.g. allowing for additional state variables or shocks.

for new models to understand the evidence and match the conditional and unconditional term-structure of risk premia across asset classes.

Related literature. Most of the existing option pricing literature has steered clear of the question whether the variance risk pricing varies across maturities. For example, work by [Coval and Shumway \(2001\)](#) or [Carr and Wu \(2009\)](#) measures variance risk premia for options with a single maturity; [Christoffersen, Heston, and Jacobs \(2013\)](#) pool all maturities when estimating the price of variance risk. Our repeated estimation of the price of variance risk on subsamples of the data differs from their approach. In contrast to [Gruber, Tebaldi, and Trojani \(2021\)](#), and [Bardgett, Gourier, and Leippold \(2019\)](#), we offer non-parametric results that are inconsistent with a constant price of variance risk, but consistent with a horizon-depend price of risk, similar to the model of [Andries, Eisenbach, and Schmalz \(2023\)](#).

Outside the standard options pricing literature, other papers have investigated the term structure of variance risk premia, using different data sets and different methodologies than the present paper. As noted above, [Dew-Becker et al. \(2017\)](#) use proprietary data on variance swaps to estimate term-structure models, similar to [Amengual \(2008\)](#) and [Aït-Sahalia, Karaman, and Mancini \(2020\)](#), but add realized volatility as a third factor in addition to the first two principle components (level and slope). 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 the data (index options as opposed to variance swaps) and methodology (estimation of an options pricing model as opposed to price of variance swaps) we use are different and complementary to [Dew-Becker et al. \(2017\)](#). There is evidence of segmentation between the equity market and the options market, as well as within different options markets such that answers to the same economic question can differ based on which market is used and information from different markets can be complementary (see, e.g., [Barras and Malkhozov, 2016](#), [Hülsbusch and Kraftschik, 2018](#), [Park, 2020](#), [Van Tassel, 2020](#), and [Oh and Park, 2023](#)). Nonetheless, as noted above, our estimates from publicly available option data have implications for variance forwards that match the direct evidence from proprietary variance swap data in [Dew-Becker et al. \(2017\)](#).

One potential explanation for our finding of a non-constant price of variance risk is a risk of jumps that have intensities or prices that vary by horizons. Some recent option pricing models with jumps find a non-constant variance risk pricing in the term-structure ([Gruber, Tebaldi, and Trojani, 2021](#); [Bardgett, Gourier, and Leippold, 2019](#)). However, a distinguishing feature of both papers is a change of slope between high and low volatility

regimes, inconsistent with the results we obtain from the data.

Less closely related are [Choi, Mueller, and Vedolin \(2017\)](#), who find a negative and upward-sloping term structure of variance premia in the Treasury futures market. Our conditional results on the relationship between current market volatility and the term structure of risk prices are related to [Cheng \(2018\)](#) who studies the returns of hedging volatility with VIX futures.

The paper proceeds as follows. Section 2 presents our conceptual framework and estimation procedures. Section 3 describes our data sources and presents our main empirical results. Section 4 derives the term-structure of variance forwards under our horizon-dependent estimates. Section 5 provides robustness checks. Section 6 concludes.

2 Conceptual framework and empirical strategy

Following [Christoffersen, Heston, and Jacobs \(2013, hereafter CHJ\)](#), we base our analysis on the canonical option-pricing model with stochastic variance of [Heston \(1993\)](#) where the stock price S_t and variance v_t satisfy the following physical dynamics:

$$\begin{aligned} dS_t &= (r_t + \mu v_t) S_t dt + \sqrt{v_t} S_t dz_{1t} \\ dv_t &= \kappa (\theta - v_t) dt + \sigma \sqrt{v_t} \left(\rho dz_{1t} + \sqrt{1 - \rho^2} dz_{2t} \right) \end{aligned} \quad (1)$$

The shock processes z_{1t} and z_{2t} are independent Brownian motions where z_{1t} affects both stock return and variance and z_{2t} only affects variance. CHJ then assume an exponential-affine stochastic discount factor that allows for aversion to variance risk, captured by the parameter ζ :

$$\frac{M_t}{M_0} = \left(\frac{S_t}{S_0} \right)^\phi \exp \left(\delta t + \eta \int_0^t v_s ds + \zeta (v_t - v_0) \right) \quad (2)$$

As in the standard case without aversion to variance risk that obtains for $\zeta = 0$ ([Rubinstein, 1976](#)), δ and η govern time preference and $\phi < 0$ aversion to equity risk. With aversion to variance risk, $\zeta > 0$, the stochastic discount factor is directly increasing in variance v_t .

Combining the physical dynamics with the stochastic discount factor results in the following risk-neutral dynamics:

$$\begin{aligned} dS_t &= r_t S_t dt + \sqrt{v_t} S_t dz_{1t}^* \\ dv_t &= (\kappa (\theta - v_t) - \lambda v_t) dt + \sigma \sqrt{v_t} \left(\rho dz_{1t}^* + \sqrt{1 - \rho^2} dz_{2t}^* \right) \end{aligned}$$

In this setting, both the instantaneous equity risk premium μv_t and variance risk premium λv_t scale with instantaneous variance v_t and depend on both equity risk aversion ϕ and variance risk aversion ξ as

$$\begin{aligned}\mu &= -\phi - \xi\sigma\rho, \\ \lambda &= -\rho\sigma\phi - \sigma^2\xi.\end{aligned}$$

Even without variance risk aversion ($\xi = 0$) the model can generate a non-zero variance risk premium solely from equity risk aversion $\phi \neq 0$. However, CHJ show that allowing for distinct variance risk aversion $\xi \neq 0$ is crucial for proper variance risk pricing as it allows for stochastic discount factor that is U-shaped in returns. Consistent with the fact that volatility is high both when returns are high and when returns are low, this significantly improves the fit of the model and reconciles several puzzles in the previous option pricing literature (see also [Heston, Jacobs, and Kim, 2022](#)).

The focus of our analysis is to estimate if the compensation λv_t that investors in options demand to bear variance risk depends on the horizon of the options traded and how such a term structure of variance risk pricing depends on the economic environment. Since the instantaneous variance risk premium λv_t depends on current variance v_t — which varies in the time series — we focus our analysis on the variance risk premium parameter λ and the underlying variance risk aversion parameter ξ that CHJ show is key. Specifically, we assess if options of different maturities imply systematically different pricing of variance risk, as reflected in λ and ξ , while keeping the parameters of the underlying physical process (1) constant for all options.

We now explain the two different estimation procedures we use to measure λ and ξ : a non-parametric estimation using Sharpe ratios of delta-neutral straddles and a parametric estimation based on [Christoffersen et al. \(2013\)](#).

2.1 Non-parametric approach

We first show that Sharpe ratios of delta-neutral straddles provide a non-parametric estimate of the variance risk premium parameter λ , our parameter of interest. In our version of the Heston model, the price X_t of any option follows a physical dynamic given by:

$$\begin{aligned}dX_t &= \left(r_t X_t + \frac{\partial X}{\partial S} \mu v_t S_t + \frac{\partial X}{\partial v} \lambda v_t \right) dt \\ &\quad + \frac{\partial X}{\partial S} \sqrt{v_t} S_t dz_{1t} + \frac{\partial X}{\partial v} \sigma \sqrt{v_t} dz_{2t}\end{aligned}\tag{3}$$

For a delta-neutral straddle, we have $\partial X/\partial S = 0$ such that the only pricing parameter remaining in the dynamic is λ .³ If we discretize the dynamic with sufficiently small Δt and substitute in $\partial X/\partial S = 0$, we obtain

$$\Delta X_t \approx \left(r_t X_t + \frac{\partial X}{\partial v} \lambda v_t \right) \Delta t + \frac{\partial X}{\partial v} \sigma \sqrt{v_t} \Delta z_{2t}, \quad (4)$$

which we can rearrange as

$$\frac{\frac{\Delta X_t}{X_t} - r_t \Delta t}{\frac{1}{X_t} \frac{\partial X}{\partial v} \sigma \sqrt{v_t} \sqrt{\Delta t}} \approx \frac{1}{\sigma} \lambda \sqrt{v_t \Delta t} + \frac{\Delta z_{2t}}{\sqrt{\Delta t}}. \quad (5)$$

Note that the variance of the discrete dynamic (4) is given by

$$\text{var}_t(\Delta X_t) \approx \left(\frac{\partial X}{\partial v} \sigma \sqrt{v_t} \right)^2 \Delta t,$$

which we can rewrite to correspond to the numerator on the left-hand side of equation (5):

$$\sqrt{\text{var}_t \left(\frac{\Delta X_t}{X_t} - r_t \Delta t \right)} \approx \frac{1}{X_t} \frac{\partial X}{\partial v} \sigma \sqrt{v_t} \sqrt{\Delta t} \quad (6)$$

Substituting this into equation (5) and taking expectations, we arrive at an expression for the Sharpe ratio of a delta-neutral straddle which scales linearly with the variance risk premium parameter λ :

$$\frac{E_t \left[\frac{\Delta X_t}{X_t} - r_t \Delta t \right]}{\sqrt{\text{var}_t \left(\frac{\Delta X_t}{X_t} - r_t \Delta t \right)}} \approx \lambda \times \frac{\sqrt{v_t \Delta t}}{\sigma} \quad (7)$$

The Sharpe ratio of a delta-neutral straddle corresponds to our parameter of interest λ scaled by a factor $\sqrt{v_t \Delta t}/\sigma$. Since $\sqrt{v_t \Delta t}/\sigma > 0$ is determined by the underlying physical asset process (1), it is the same at all maturities. As a result, the Sharpe ratios of delta-neutral straddles at different maturities provide a qualitatively valid measure of the term-structure of λ , even though they are not quantitatively comparable to the results from the parametric estimation we present in Section 2.2.⁴

³Delta-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. Following the literature, we compute the delta-neutral portfolios using the Black-Scholes model.

⁴In 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 (7) to hold. The straddles analyzed by [van Binsbergen and Koijen \(2016\)](#) have deltas that increase

We estimate the Sharpe ratios of options with different maturities ranging from 20 days to 252 days, using weekly returns. We group straddles into four maturity categories: 20 to 60 days, 60 to 125 days, 125 to 190 days, and 190 to 252 days. The results of our analysis are described and discussed in Section 3.

2.2 Parametric approach

The factor $\sqrt{v_t \Delta t} / \sigma$ in the Sharpe ratio of a delta-neutral straddle in (7) varies in the time series with instantaneous volatility v_t . These time series variations can be correlated with variations in the slope of the term-structure of variance risk pricing — and we show in Section 3 that they are. 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 study the parameter λ directly in a parametric model.

We follow CHJ and use the discrete-time analog to the continuous time model above, as developed by Heston and Nandi (2000) where the stock price S_t follows a GARCH-in-means process and the one-period excess return has variance h_t ,

$$\begin{aligned} \log S_t &= \log S_{t-1} + r_t + \left(\mu - \frac{1}{2}\right) h_t + \sqrt{h_t} z_t \\ h_t &= \omega + \beta h_{t-1} + \alpha \left(z_{t-1} - \gamma \sqrt{h_{t-1}}\right)^2, \end{aligned} \tag{8}$$

with $z_t \sim \mathcal{N}(0, 1)$. 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.

The discrete-time analog of the stochastic discount factor in (2) is given by

$$\frac{M_t}{M_0} = \left(\frac{S_t}{S_0}\right)^\phi \exp\left(\delta t + \eta \sum_{s=1}^t h_s + \zeta (h_{t+1} - h_1)\right).$$

CHJ show that the risk-neutral counterpart to process (8) is

$$\begin{aligned} \log S_t &= \log S_{t-1} + r_t - \frac{1}{2} h_t^* + \sqrt{h_t^*} z_t^*, \\ h_t^* &= \omega^* + \beta h_{t-1}^* + \alpha^* \left(z_{t-1}^* - \gamma^* \sqrt{h_{t-1}^*}\right)^2, \end{aligned} \tag{9}$$

with maturity, and thus depart from the delta neutrality required by equation (4).

with $z_t^* \sim \mathcal{N}(0, 1)$ and

$$\begin{aligned} h_t^* &= \Xi h_t, & \omega^* &= \Xi \omega, \\ \alpha^* &= \Xi^2 \alpha, & \gamma^* &= \frac{1}{\Xi} \left(\gamma + \mu - \frac{1}{2} \right) + \frac{1}{2}, \end{aligned}$$

where the factor $\Xi \equiv (1 - 2\alpha\tilde{\zeta})^{-1}$, increasing in variance risk aversion $\tilde{\zeta}$, plays the key role in the transformation from physical to risk-neutral terms.⁵ To compensate for variance risk, the risk neutral variance process has a higher long-run mean and higher persistence for $\tilde{\zeta} > 0$ or, equivalently, $\Xi > 1$. The only notable difference between the discrete time and continuous time models 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 . [Heston and Nandi \(2000\)](#) also show by simulation that as the time intervals converge to zero, the GARCH price converges to the continuous time Heston price.

CHJ estimate the GARCH parameters and a common variance risk aversion through the factor $\Xi = (1 - 2\alpha\tilde{\zeta})^{-1}$ jointly with a likelihood that incorporates both returns and option prices. The empirical methodology of CHJ has two important advantages. First, the GARCH discrete-time structure reduces the complexity of the filtering problem inherent in any empirical implementation of the Heston model. Second, using the returns data and the options data jointly means fitting the physical and risk neutral processes with a single set of internally consistent GARCH and preference parameters.

We therefore follow their approach exactly except for two differences.⁶ First, we do not smooth the inputs by computing a volatility surface but, instead, 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. Second, and more importantly, we allow the variance risk aversion factor Ξ to vary by maturity, using four maturity buckets: options with maturities of 20 to 60 days, 60 to 125 days, 125 to 190 days, and 190 to 252 days. In our main results in Section 3, we keep the parameters of the underlying process (8) constant across maturities; we show in Section 5 that our results are robust to letting all parameters vary across maturities.

Using the estimates for Ξ_1, \dots, Ξ_4 , we calibrate the corresponding term structure of continuous-time variance risk premium parameters $\lambda_1, \dots, \lambda_4$ following CHJ to obtain

⁵Given the physical GARCH parameters $\Theta = \{\omega, \beta, \alpha, \mu, \gamma\}$ and the variance risk aversion $\tilde{\zeta}$, the equity risk aversion ϕ is pinned down as $\phi = -\frac{1}{\Xi} \left(\mu - \frac{1}{2} + \gamma \right) + \gamma - \frac{1}{2}$ (see [Heston and Nandi \(2000\)](#) and CHJ for additional details).

⁶Our estimation methodology, adapted from CHJ, is described in more detail in the appendix.

the same unconditional variance of stock returns, the same physical variance persistence, and the same ratio between physical and risk-neutral unconditional variances as in the discrete-time model:

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

where

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

and $\omega_j^*, \alpha_j^*, \gamma_j^*$ are the risk neutralized parameters using the variance aversion factor Ξ_j for each maturity bucket.

2.3 Sample splits

Our interest is in whether estimated variance risk pricing varies with the state of the economy as well as with the horizon of the option. To explore this question, we calculate the likelihoods for several splits of the data by VIX levels, price-dividend ratios and GDP growth, and compare the estimated variance risk pricing across horizons and across splits. In particular, we obtain daily closing VIX index values from the CBOE Indexes data, GDP growth data on a quarterly basis from the BEA, and monthly data on the S&P 500 price-dividend ratio from S&P 500 Ratios via Nasdaq Data Link. For each of these variables, we split our samples into “low” and “high” based on whether the variable is below or above its median value over our sample period. Table 1 reports summary statistics over these subsamples. Times of high VIX, low PD ratios and low GDP growth are associated with greater implied volatility and higher option prices. However, these times are also associated with higher option implied volatility at short maturities than at long maturities. Our parametric analysis is designed to determine how much of these differences are due to the underlying index return process versus the pricing of variance risk.

Table 1: Summary statistics of options pricing data by subsamples. The table shows summary statistics of SPX index options prices in the sample used for this paper from January 1996 to January 2018 for sample splits by high and low VIX, low and high price-dividend (PD) ratio, and low and high GDP growth; and, for implied volatility, by maturity.

	Full sample	VIX split		PD ratio split		GDP growth split	
		High	Low	Low	High	Low	High
Implied volatility	21.58	25.23	15.41	23.34	20.50	24.30	19.46
Mid-price	28.68	31.53	22.72	30.04	27.17	31.23	25.95
Bid-ask spread	1.75	1.88	1.50	1.95	1.62	2.06	1.49
Implied volatility by maturity							
≤30	21.41	25.43	14.45	22.97	20.62	24.59	19.31
30-60	21.46	25.56	14.76	23.48	20.44	24.64	19.33
60-90	22.14	26.12	15.46	24.13	20.41	24.50	19.49
90-120	21.64	25.69	15.68	23.59	20.60	24.36	19.69
120-180	21.51	24.70	16.01	23.35	20.46	24.19	19.57
>180	21.42	24.07	16.37	22.36	20.59	23.47	19.53

3 Data and empirical results

3.1 Data sources and summary statistics

We use daily closing data of European SPX index options and SPX index levels from January 1996 to January 2018 from OptionMetrics. S&P 500 returns, excluding dividends, from January 1990 to January 2018 come from CRSP. The risk-free rate for a given daily return observation is defined as $\log(1 + r_t)/252$, where r_t is the average effective federal funds rate for the 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 20 and 252 trading days, inclusive, on the day of observation. Using trading days to measure maturity is essential. The GARCH estimation treats the index return series as a continuous series without weekends. To be consistent, the option maturities should therefore also be expressed in trading days. We follow [Bakshi, Cao, and Chen \(1997\)](#) in excluding shorter-maturity options to avoid microstructure noise close to expiration, and we exclude longer-maturity options because they are thinly traded. We also follow [Bakshi, Cao, and Chen \(1997\)](#) in excluding any options that have quoted bid prices below \$3/8 to avoid discretization issues or options that do not obey the futures

Table 2: Summary statistics of options pricing data. The table shows summary statistics sorted by moneyness and maturity of SPX index options prices in the full data sample used for this paper from January 1996 to January 2018.

Summary Statistics by Maturity							
Maturity	≤ 30	30-60	60-90	90-120	120-180	> 180	All
Count	3,419	6,810	4,184	2,318	4,357	5,317	26,405
Implied volatility	21.41	21.46	22.14	21.64	21.51	21.42	21.58
Mid-price	11.61	18.62	25.98	30.14	36.79	47.41	28.68
Bid-ask spread	1.03	1.45	1.82	1.87	1.96	2.31	1.75
Summary Statistics by Moneyness							
Moneyness	≤ 0.96	0.96-0.98	0.98-1.02	1.02-1.04	1.04-1.06	> 1.06	All
Count	6,428	2,296	6,252	2,122	1,800	7,507	26,405
Implied volatility	18.82	18.00	18.87	20.36	22.05	27.53	21.58
Mid-price	26.48	39.26	40.03	30.56	26.83	17.80	28.68
Bid-ask spread	1.79	1.91	1.94	1.75	1.73	1.51	1.75

arbitrage constraints: for a call with maturity τ , $C(\tau) \geq \max\{0, S_t - X_t e^{-r_t \tau}\}$, and for a put, $P(\tau) \geq \max\{0, X_t e^{-r_t \tau} - S_t\}$.

We restrict our attention to out of the money options to avoid well known issues with the liquidity of in the money options. On every Wednesday, for each maturity we select the out of the money option from each of the 6 most highly traded strikes: if the strike is greater than the stock price we choose the call; if the strike is less than the stock price we choose the put. We convert all put prices to call prices by put-call parity.⁷

Table 2 presents summary statistics for the sample of 26,405 option-day observations used in the parametric analysis. In this sample, the average implied volatility is increasing with maturity. When sorted by moneyness, we also see evidence of both put skew and the volatility smile; out of the money puts have much higher implied volatility than calls and, in general, out of the money options have higher implied volatility than at the money options. Liquidity improves at longer maturities, with the mean bid-ask spread at around 9% of the average mid-price for short maturities and declining to around 5% at longer maturities.

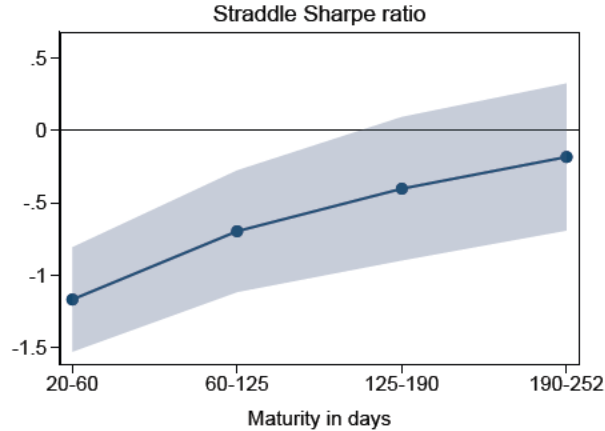


Figure 1: Non-parametric term structure of variance risk pricing. The figure shows Sharpe ratios of delta-neutral straddles by maturity which provide non-parametric estimates of the variance risk premium parameter λ up to a scaling factor. Shaded areas indicate 95% confidence intervals.

3.2 Term structure of variance risk pricing

Figure 1 shows the non-parametric term structure of variance risk pricing using Sharpe ratios of delta-neutral straddles. The Sharpe ratios provide estimates of the variance risk premium parameter λ up to a scaling factor (see Section 2.1). The Sharpe ratios imply a negative variance risk premium at all horizons, significantly negative up to the maturity bucket of 3–6 months (60–125 days) and with a decreasing term structure (in absolute value), indicating that investors require compensation for bearing variance risk and considerably more so at shorter horizons.

Figure 2 shows the parametric term structure of variance risk pricing estimated on the full sample. The left panel shows estimates of the variance risk aversion factor Ξ while the right panel shows the resulting calibrations of the variance risk premium parameter λ (see Section 2.2). The variance risk aversion factor Ξ is significantly greater than 1 (i.e. $\xi > 0$) at all maturities except the longest maturity (190–252 days); the resulting variance risk parameter λ is therefore negative with a decreasing term structure (in absolute value), consistent with its non-parametric estimate in Figure 1. The term structure of Ξ is monotonically decreasing and the declines are statistically significant from the first to the second and from the second to the third maturity bucket.

Table 3, column (1) shows the full estimation results, including the GARCH parameters

⁷We use the dividend yield for the index and an interpolated yield curve from the IvyDB interest rate curve data to construct discounted dividends to then get to call prices.

Table 3: Estimation on full sample and subsamples. The table shows estimates of the model using joint maximum-likelihood estimation for the full sample as well as for sample splits by high and low VIX, low and high price-dividend (PD) ratio, and low and high GDP growth. GARCH processes are held constant over both splits in columns (2)–(7). The top panel shows the estimates for GARCH parameters, the middle panel shows estimates for the variance risk preference parameter, ζ , and the bottom panel shows the likelihood from the fit to returns, Φ_R , the fit to options prices, Φ_O , and the joint likelihood which is the sum of the returns likelihood and the options prices likelihood.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Full sample	VIX split		PD ratio split		GDP growth split	
		High	Low	Low	High	Low	High
β	0.735 (0.004)	0.747 (0.003)		0.725 (0.004)		0.733 (0.003)	
α	1.3×10^{-6} (0.02×10^{-6})	1.1×10^{-6} (0.02×10^{-6})		1.3×10^{-6} (0.03×10^{-6})		1.3×10^{-6} (0.03×10^{-6})	
γ	451.24 (5.72)	463.62 (4.35)		459.38 (6.55)		445.72 (6.14)	
Ξ_{20-60}	1.212 (0.018)	1.896 (0.020)	1.065 (0.019)	1.377 (0.026)	1.166 (0.020)	1.321 (0.026)	1.169 (0.020)
Ξ_{60-125}	1.108 (0.013)	1.621 (0.015)	1.020 (0.014)	1.205 (0.018)	1.082 (0.015)	1.198 (0.015)	1.065 (0.014)
$\Xi_{125-190}$	1.025 (0.010)	1.414 (0.011)	1.000 (0.010)	1.116 (0.014)	1.007 (0.011)	1.105 (0.013)	1.000 (0.013)
$\Xi_{190-252}$	1.000 (0.009)	1.320 (0.010)	1.000 (0.009)	1.056 (0.012)	1.000 (0.009)	1.048 (0.011)	1.000 (0.007)
Φ_R	29,678	29,740		29,683		29,690	
Φ_O	70,951	44,065	30,039	25,481	45,614	30,657	40,397
$\Phi_R + \Phi_O$	100,630	103,844		100,779		100,745	




Figure 2: Parametric term structure of variance risk pricing. The figure shows estimates of the variance risk aversion factor Ξ (left panel) and resulting calibrations of the variance risk premium parameter λ (right panel) by maturity on the full sample (Table 3, column 1). Shaded areas indicate 95% confidence intervals.

α , β and γ . Estimates of the parameters of the GARCH process are similar to those in CHJ who, in their joint estimation, find $\hat{\beta} = 0.756$, $\hat{\alpha} = 1.41 \times 10^{-6}$ and $\hat{\gamma} = 515.57$. The slight differences in our estimated GARCH parameters imply a lower autocorrelation of variance for our GARCH process. This may be due to the longer time series we use, which extends beyond 2009 when their estimation ends and includes the relatively low volatility period from 2013 to 2018 as well as spikes in volatility in 2010 and 2011. We can split our sample into two halves (pre/post 2007) and show in Appendix Figures A1 and A2 that variance risk is priced more in the post-2007 period. Our sample splits into normal and stressed times in Section 3.3 help us understand this in more detail.

3.3 Term structure in different subsamples

To study how the pricing of variance risk and its term structure vary with different states of the economy, we estimate the term structure of variance risk pricing separately on subsamples of the data. Figure 3 shows estimates of the Ξ and the resulting calibrations of λ when splitting the sample into high/low VIX, price-dividend ratio, and GDP growth.⁸ Table 3, columns (2) to (7) show the full parametric estimation results, including the GARCH parameters α , β and γ . Periods of high VIX, low price-dividend ratio and low GDP growth all indicate stressed states of the economy. Irrespective of how we identify stressed times, we

⁸The non-parametric estimates from straddle Sharpe ratios do not show significant differences across sample splits.

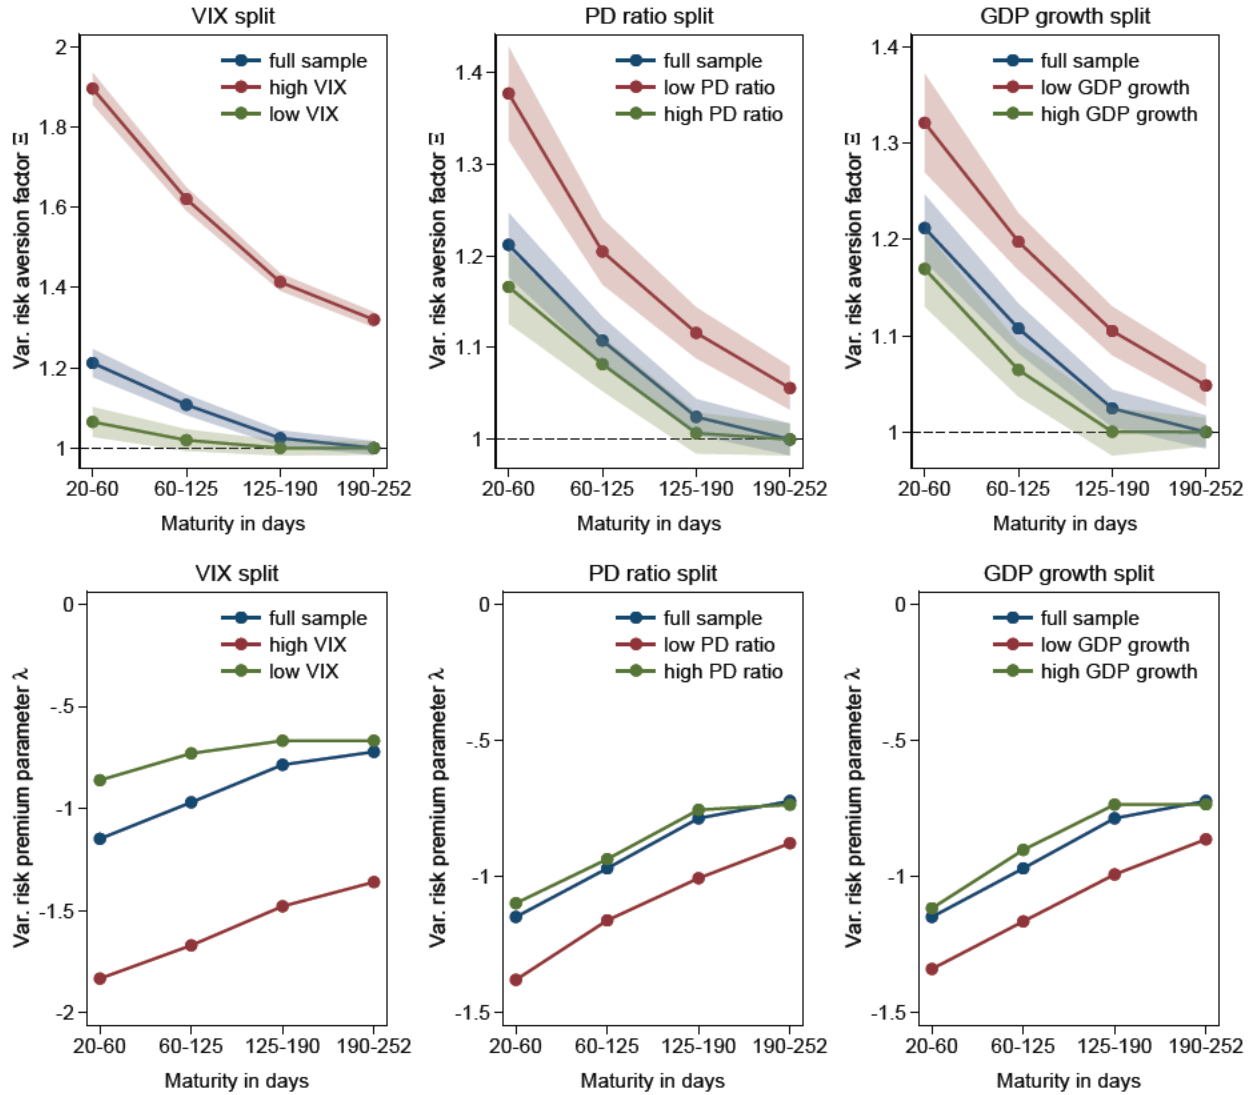


Figure 3: Term structure in different subsamples. The figure shows estimates of Ξ (top panels) and resulting calibrations of λ (bottom panels) on subsamples and compares them to the full-sample versions (Table 3, columns 2–7). Shaded areas indicate 95% confidence intervals.

see that investors require significantly higher compensation for variance risk and the slope of the term structure remains downward sloping. Moreover, across these sub-samples, the variance risk aversion factor Ξ is significantly greater than 1 out to the 60–125 day bin, showing that the result that longer-maturity options display a significant variance risk premium persists even in these subsamples.⁹

The difference between the term structures is starkest for the VIX split where both the level of the curve and the slope (in absolute magnitude) increase the most in stressed times compared to normal times. This is consistent with the VIX being conceptually closest as the relevant measure to split the price of variance risk on. The VIX split also provides considerably more variation and is reasonably orthogonal to the other sample splits. While the VIX split has a correlation of 0.230 with the GDP growth split which only varies at quarterly frequency, it is uncorrelated with the PD ratio split which varies at the same daily frequency (correlation of 0.001). Moreover, across maximum likelihood estimates, high VIX periods display the greatest improvement in the total log likelihood of any sub-sample split, again emphasizing the close relationship between the VIX and the price of variance risk (bottom of Table 3).

4 Implied term structures of variance forwards

Dew-Becker et al. (2017) provide an alternative perspective on the pricing of variance risk at different maturities using proprietary data on variance swap prices to construct “variance forwards.” These represent claims on realized variance at future dates such that the price F_t^n of an n -month variance forward at date t is a model-free measure of the risk neutral expectation of variance in month $t + n$:¹⁰

$$F_t^n = E_{t-1}^*[h_{t+n}^*]$$

In the CHJ model, as in the original Heston model, physical variance and risk neutral variance have conditional term structures since both are expected to revert to their respective unconditional averages. For physical variance h_t we have from the evolution in (8)

⁹Note that the full-sample estimate of Ξ does not necessarily have to be within the range spanned by the subsample estimates because the GARCH parameters of the full-sample estimation can differ from those of the split-sample estimation.

¹⁰Note that in discrete time, the variance h_t of the stock return $\log(S_t/S_{t-1})$ at date t is already known at $t - 1$ after the realization of the shock z_{t-1} .

that the conditional expectation of variance n periods ahead is given by

$$E_{t-1}[h_{t+n}] = (\beta + \alpha\gamma^2)^n h_t + \left(1 - (\beta + \alpha\gamma^2)^n\right) E[h],$$

with the unconditional average given by

$$E[h] = \frac{\omega + \alpha}{1 - \beta - \alpha\gamma^2}.$$

Analogously for risk neutral variance h_t^* , we have from the evolution in (9) that the conditional expectation (under the risk-neutral measure) of risk-neutral variance n periods ahead is given by

$$E_{t-1}^*[h_{t+n}^*] = (\beta + \alpha^*\gamma^{*2})^n h_t^* + \left(1 - (\beta + \alpha^*\gamma^{*2})^n\right) E^*[h^*], \quad (10)$$

with the unconditional average given by

$$E^*[h^*] = \frac{\omega^* + \alpha^*}{1 - \beta - \alpha^*\gamma^{*2}}.$$

While the unconditional term structures of both variances are flat under their respective measures, the unconditional term structure of *risk-neutral* variance under the *physical* measure is not flat. Specifically, substituting $h_t^* = \Xi h_t$ into (10) and taking unconditional expectations under the physical measure, we have

$$\begin{aligned} E[h_{t+n}^*] &\equiv E[E_{t-1}^*[h_{t+n}^*]] \\ &= (\beta + \alpha^*\gamma^{*2})^n \Xi E[h] + \left(1 - (\beta + \alpha^*\gamma^{*2})^n\right) E^*[h^*]. \end{aligned} \quad (11)$$

At the short end, average risk neutral variance $E[h_{t+n}^*]$ is anchored by average realized variance $E[h]$,

$$E[h_{t+0}^*] = \Xi E[h],$$

where the factor Ξ is due to the discrete-time structure. With increasing horizon, this anchor grows weaker and average risk neutral variance converges to its expectation under the risk-neutral measure:

$$\lim_{n \rightarrow \infty} E[h_{t+n}^*] = E^*[h^*]$$

Dew-Becker et al. (2017) show that, in the data, the unconditional average term structure of variance forward prices, i.e. $E[F_t^n] = E[h_{t+n}^*]$, is steeply upward-sloping at the short

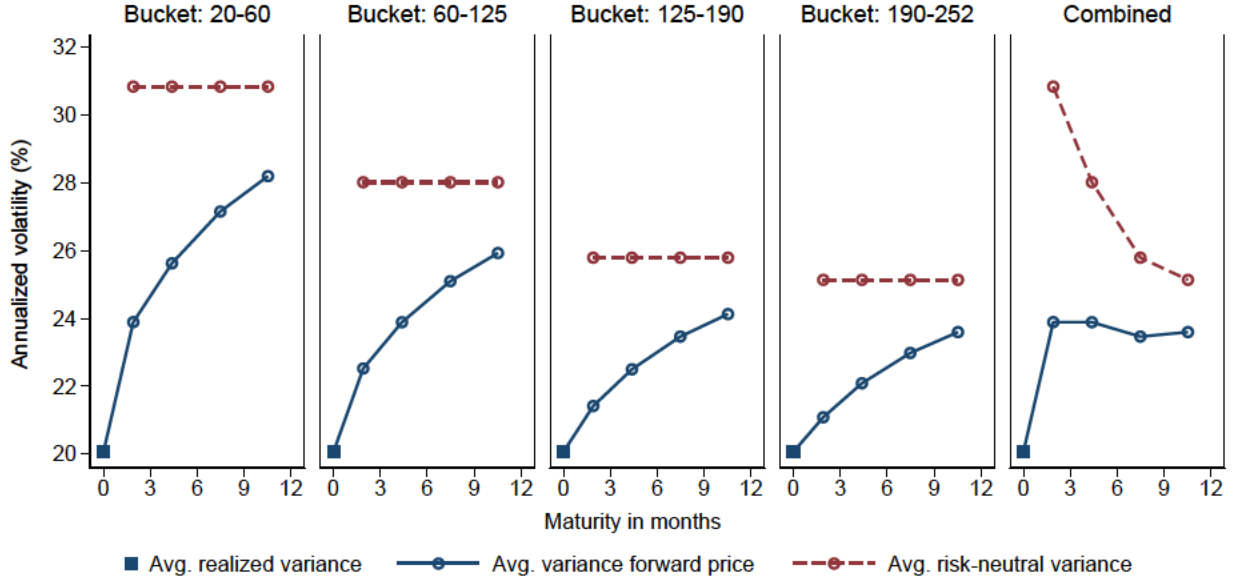


Figure 4: Implied term structures of variance forward prices. The figure shows the unconditional average term structures implied by the parametric estimates of variance risk pricing in each of four maturity buckets (first four panels). The last panel shows the implied term structures using each bucket for the respective maturity.

end but then essentially flat at longer maturities. This is in contrast to the more gradually increasing term structure in (11) that is implied by a [Heston \(1993\)](#) stochastic variance model such as CHJ.

How do our estimates by maturity bucket relate to this evidence? We have a different estimate $\hat{\Xi}_j$ for each option maturity bucket j in 20 to 60 days, 60 to 125 days, 125 to 190 days, and 190 to 252 days. Each of these estimates results in different risk neutral parameters α^* , γ^* and $E^*[h^*]$, and therefore implies a different unconditional average term structure of variance forward prices $E[h_{t+n}^*]$ from equation (11). The first four panels of [Figure 4](#) show these term structures as well as the unconditional average risk neutral variance $E^*[h^*]$ that each term structure converges to.¹¹

However, this approach assumes that variance risk pricing as represented by the parameter Ξ is the same through the entire term structure, conflicting with our result that the estimates $\hat{\Xi}_j$ are monotonically decreasing. The last panel of [Figure 4](#) therefore shows the implied term structure of variance forward prices when using the corresponding estimate $\hat{\Xi}_j$ for each maturity bucket j .¹² The resulting term structure is strikingly similar to that

¹¹For consistency with [Dew-Becker et al. \(2017\)](#), we anchor the variance forward term structure at the average realized (physical) variance for maturity $n = 0$.

¹²We show the term structures at the midpoints of the maturity buckets, i.e. 40 days for bucket 20–60, 92

in Figure 2 of [Dew-Becker et al. \(2017\)](#): a large jump from average realized variance to the average variance forward price at the shortest maturity and then essentially flat at longer maturities.

The structural framework of our analysis provides an interpretation of this result: At each maturity, the average variance forward price is a convex combination of average realized variance and average risk neutral variance with the weight on average risk neutral variance increasing in maturity. While average realized variance cannot vary across maturities, average risk neutral variance decreases in maturity, counteracting the effect of the increasing weight in the convex combination. It turns out that, empirically, the effects of decreasing risk neutral variance and increasing weight almost perfectly offset each other, resulting in the flat term structure of average variance forward prices.

5 Robustness of the parametric estimation

5.1 Joint vs. sequential estimation

Our baseline specification follows CHJ in estimating the variance risk aversion factor Ξ (and therefore the variance risk premium parameter λ) jointly with the GARCH parameters α , β , and γ . Figure 5 shows the coefficient estimates for Ξ if we do the estimation in two sequential steps, first estimating the GARCH parameters without making use of the options data and then estimating the variance risk aversion factor with the options data. Table 4, column (2) shows the full estimation results, including the GARCH parameters α , β and γ while column (1) repeats the joint estimation of the benchmark specification. While the resulting estimates show a level shift and the slope becomes flatter it is still significantly downward sloping (in absolute value).

Comparing the likelihoods, we see that the total likelihood of the joint estimation is considerably higher (bottom of Table 4), consistent with the importance of the options data informing the GARCH parameters (as in CHJ). The difference in option pricing parameters suggest that options prices imply higher volatility and lower persistence of volatility than would be gained by estimation of the GARCH process using returns alone. One possibility is that the period of the sample estimation has seen relatively high volatility compared to the physical process, which option prices foresee but which the GARCH process is unable to pick up.

for 60–125, 157 for 125–190, and 221 for 190–252.

Table 4: Sequential estimation and allowing the GARCH parameters to vary with horizon.

The table shows alternative specifications of the unconditional estimates of the model. Column (1) uses joint maximum-likelihood estimation for the full sample, column (2) uses sequential estimation, first estimating the GARCH process using returns and then estimating Ξ using option prices, and columns (3)–(6) use joint estimation which allow both GARCH parameters and Ξ to vary by maturity. The top panel shows the estimates for GARCH parameters, the middle panel shows estimates for the variance risk preference parameter, ζ , and the bottom panel shows the likelihood from the fit to returns, Φ_R , the fit to options prices, Φ_O , and the joint likelihood which is the sum of the returns likelihood and the options prices likelihood.

	(1)	(2)	(3)	(4)	(5)	(6)
	Joint	Sequential	Joint, by maturity			
β	0.735 (0.004)	0.819 (0.013)	0.714 (0.006)	0.705 (0.007)	0.775 (0.009)	0.787 (0.006)
α	1.3×10^{-6} (0.02×10^{-6})	3.9×10^{-6} (0.0×10^{-6})	1.5×10^{-6} (0.03×10^{-6})	1.2×10^{-6} (0.04×10^{-6})	1.2×10^{-6} (0.05×10^{-6})	1.1×10^{-6} (0.04×10^{-6})
γ	451.24 (5.72)	193.63 (1.94)	429.69 (5.14)	488.67 (10.19)	417.86 (4.42)	436.44 (11.50)
Ξ_{20-60}	1.212 (0.018)	1.240 (0.003)	1.227 (0.015)			
Ξ_{60-125}	1.108 (0.013)	1.217 (0.002)		1.141 (0.019)		
$\Xi_{125-190}$	1.025 (0.010)	1.200 (0.003)			1.000 (0.020)	
$\Xi_{190-252}$	1.000 (0.009)	1.211 (0.005)				1.000 (0.017)
Φ_R	29,678	29,873	29,724	29,669	29,666	29,654
Φ_O	70,951	61,794	25,229	19,304	13,404	12,524
$\Phi_R + \Phi_O$	100,630	91,667	54,952	48,973	43,070	42,178

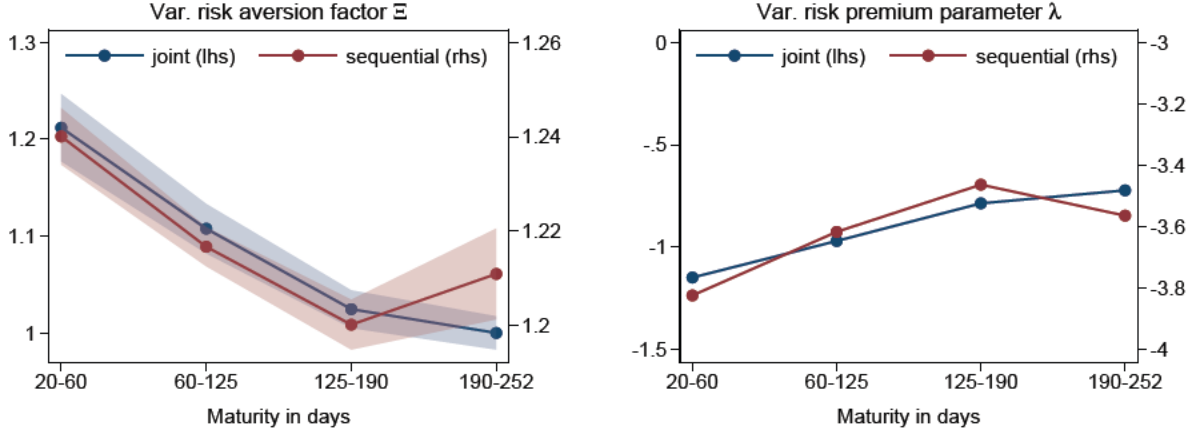


Figure 5: Estimating parameters jointly or sequentially. The figure shows estimates of Ξ (left panel) and resulting calibrations of λ (right panel) by maturity on the full sample, comparing the joint estimation of GARCH parameters and variance pricing parameters (Table 4, column 1) to a sequential estimation, first estimating the GARCH process using returns and then estimating Ξ using option prices (Table 4, column 2). Shaded areas indicate 95% confidence intervals.

5.2 Allowing the GARCH parameters to vary with horizon

Our benchmark specification only allows the variance risk aversion factor Ξ to vary with the horizon but maintains a single set of GARCH parameters α , β , and γ . Figure 6 shows the coefficient estimates if we allow the GARCH parameters to vary with the horizon. Table 4, columns (3) to (6) show the full estimation results. While this additional flexibility results in some variation in the GARCH parameter estimates across horizons, the effect on the estimates of Ξ (and therefore λ) is negligible (top-left panel of Figure 6). Considering the variation in the GARCH parameters, only α and β show a somewhat monotonic term structure. The estimates for α and β from short-horizon options suggest lower persistence of the variance and fatter tails in the distribution of shocks to the variance process, respectively. Both are conceptually consistent with our main finding of a higher price of variance risk (in absolute value) at shorter horizons.

5.3 Allowing GARCH parameters to vary with sample splits

Our benchmark specification estimates one set of GARCH parameters and only estimates the variance risk aversion factor Ξ separately on different subsamples. We could, instead, allow the GARCH parameters to vary with sample splits as well. However, due to the auto-regressive nature of the process, estimating separate GARCH parameters for differ-

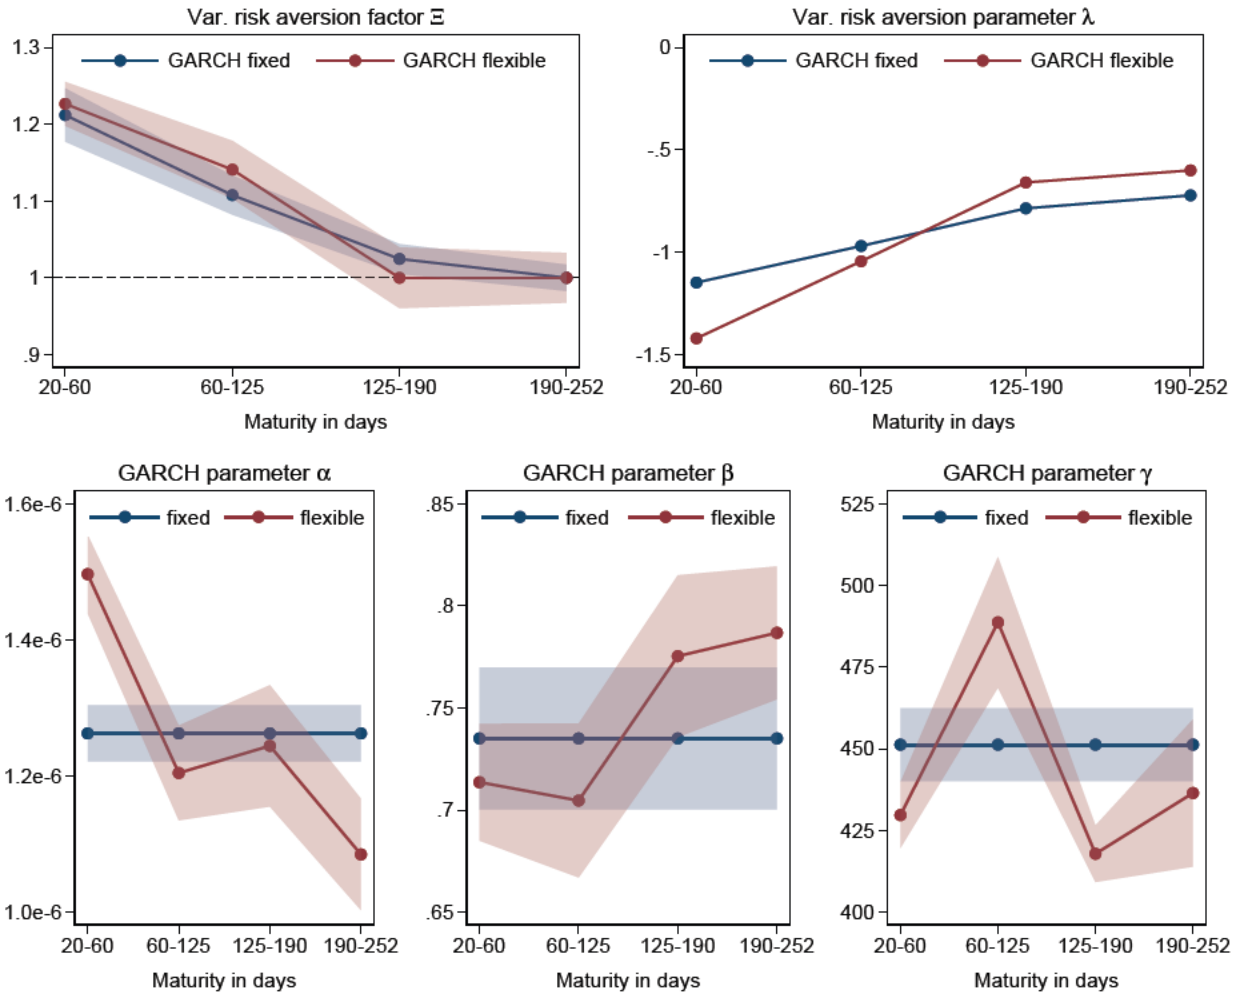


Figure 6: Allowing the GARCH parameters to vary with horizon. The figure shows estimates of Ξ (top left panel) and resulting calibrations of λ (top right panel), and the GARCH parameters (bottom panels), comparing the estimation where GARCH parameters are fixed across maturities (Table 4, column 1) to the estimation which allows both GARCH parameters and Ξ to vary by maturity (Table 4, columns 3–6). Shaded areas indicate 95% confidence intervals.

ent subsamples is conceptually problematic if the subsamples are not sufficiently long. This concern is particularly strong for our VIX split and our PD ratio split which can vary at daily frequency.

Figure 7 shows the coefficient estimates of λ and Ξ if we do allow the GARCH parameters to vary with the sample splits. Table 5, columns (2) to (7) show the full estimation results while column (1) repeats the estimation of the benchmark specification. The effect on the term structures of Ξ (and therefore λ) is negligible if we split the sample on PD ratio and on GDP growth. However, in case of the VIX split, the ordering of the term structures changes as now Ξ is higher in each of the subperiods than in the full sample period and higher in the low-VIX sample than the high-VIX sample. This is due to the fact that the GARCH parameters change considerably in the two subsamples. In the high VIX sample, we estimate considerably larger kurtosis α and smaller correlation with returns γ (and slightly lower persistence β). The reason that the GARCH process impacts the pricing of variance risk is intuitive: Options prices may be higher at longer maturities either because of a lower aversion to variance risk at longer horizons or because variance is expected to be lower in the future. Again, this emphasizes the importance of joint estimation of returns and options prices.

The finding that Ξ is higher in each sub-period than in the full sample deserves some discussion and interpretation. We note that the estimated variance of variance that results from estimating the returns process from either high- or low-variance sub-periods is artificially low, compared to the true returns process, which has a chance to move from a high- to a low-variance regime, reflecting relatively high variance of variance. As a result, the estimated model would have trouble making sense of high option prices in any sub-period, unless it assumes a high price of variance risk (in absolute value). The joint estimation solves this problem by inferring different GARCH parameters for the two periods based on the option prices as well as the returns process, which then leads to different variance risk pricing parameters that could be hard to compare. The fact that the model fit does not greatly differ between the results in which GARCH is separated for the entire sample as opposed to separately estimated in each subsample while the estimated price of variance risk differs quite starkly suggests that the reason for the latter findings is more likely to be a change in the pricing of variance risk rather than a change in the underlying returns process.

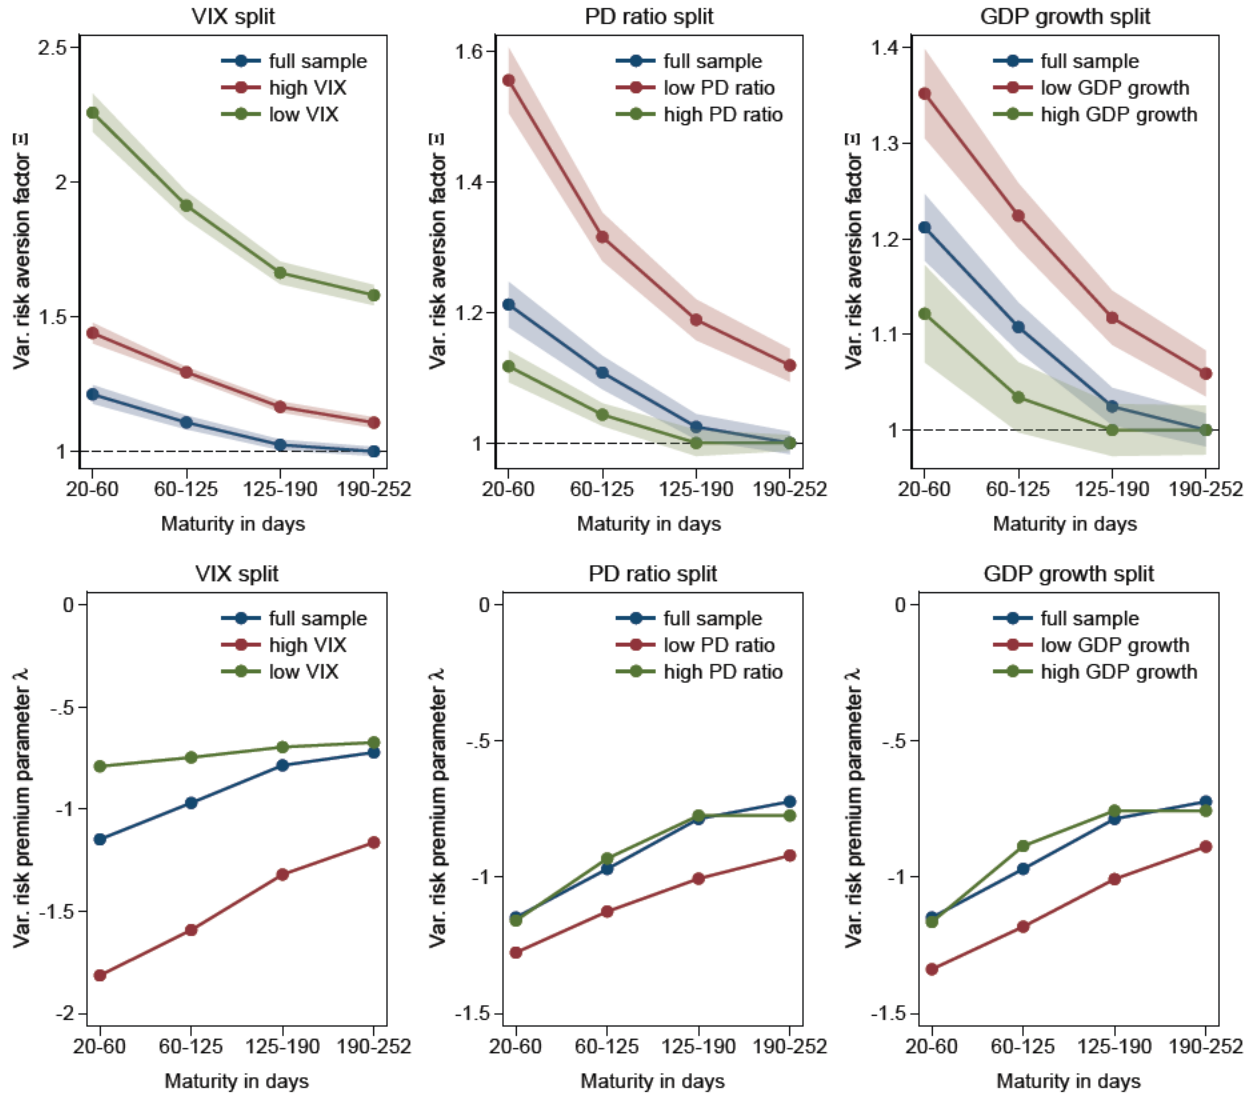


Figure 7: Allowing GARCH parameters to vary with sample splits. The figure shows estimates of Ξ (top panels) and resulting calibrations of λ (bottom panels) on subsamples, allowing for the GARCH parameters to vary with the sample splits (Table 5). Shaded areas indicate 95% confidence intervals.

Table 5: Allowing GARCH parameters to vary with sample splits. The table shows alternative specifications of the sample splits. Columns (2)–(7) repeat the sample splits of Table 3 but allow for the GARCH parameters to vary with the sample splits. The top panel shows the estimates for GARCH parameters, the middle panel shows estimates for the variance risk preference parameter, ξ , and the bottom panel shows the likelihood from the fit to returns, Φ_R , the fit to options prices, Φ_O , and the joint likelihood which is the sum of the returns likelihood and the options prices likelihood.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	VIX split		PD ratio split		GDP growth split		
Joint	High	Low	High	Low	High	Low	High
β	0.735 (0.004)	0.758 (0.005)	0.781 (0.004)	0.636 (0.006)	0.756 (0.005)	0.707 (0.004)	0.765 (0.027)
α	1.3×10^{-6} (0.02×10^{-6})	1.7×10^{-6} (0.02×10^{-6})	2.9×10^{-7} (0.06×10^{-7})	9.3×10^{-7} (0.18×10^{-7})	1.6×10^{-6} (0.02×10^{-6})	1.2×10^{-6} (0.03×10^{-6})	1.6×10^{-6} (0.05×10^{-6})
γ	451.24 (5.72)	365.79 (5.87)	856.22 (11.01)	620.34 (9.26)	383.09 (5.37)	485.78 (7.44)	376.43 (8.34)
Ξ_{20-60}	1.212 (0.018)	1.440 (0.020)	2.258 (0.036)	1.555 (0.026)	1.117 (0.012)	1.352 (0.024)	1.122 (0.026)
Ξ_{60-125}	1.108 (0.013)	1.294 (0.010)	1.913 (0.026)	1.315 (0.019)	1.043 (0.009)	1.224 (0.017)	1.034 (0.019)
$\Xi_{125-190}$	1.025 (0.010)	1.166 (0.010)	1.663 (0.021)	1.189 (0.016)	1.000 (0.010)	1.117 (0.014)	1.000 (0.014)
$\Xi_{190-252}$	1.000 (0.009)	1.107 (0.010)	1.580 (0.019)	1.119 (0.013)	1.000 (0.006)	1.059 (0.012)	1.000 (0.013)
Φ_R	29,678	29,725	29,396	29,586	29,739	29,675	29,749
Φ_O	70,951	44,501	32,659	25,991	45,892	30,733	40,491
$\Phi_R + \Phi_O$	100,630	74,226	62,055	55,576	75,631	60,409	70,239

6 Conclusion

We provide parametric and non-parametric estimates of the term structure of variance risk pricing and how it varies over time by estimating the price of variance risk in a [Heston \(1993\)](#) stochastic-volatility model. We find that the price of insurance against volatility shocks varies with the horizon of the risk insured: short-term insurance is significantly more expensive than long-term insurance, and this effect is more pronounced in times of higher volatility. The price is significant across short- and longer-maturity options and the term structure is consistently downward sloping (in absolute value) across normal times and periods of stress.

These results extend the accumulating evidence for non-trivial term structures of risk prices to the market for variance risk, and suggest the need for a new generation of option pricing models that allow for horizon-dependent risk aversion and/or more general physical dynamics.

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. Combined with the time series variations in our variance price term-structure estimates, which contrast to similar analysis on the term-structure of the equity risk pricing, our finding may help inspire future generations of asset pricing models and econometricians' interpretation of economic forecasts.

References

- Aït-Sahalia, Y., M. Karaman, and L. Mancini (2020). The term structure of equity and variance risk premia. *Journal of Econometrics* 219(2), 204–230.
- Amengual, D. (2008). The term structure of variance risk premia. Working Paper.
- Andries, M., T. M. Eisenbach, and M. C. Schmalz (2023). Horizon-dependent risk aversion and the timing and pricing of uncertainty. Working Paper.
- Bakshi, G., C. Cao, and Z. Chen (1997). Empirical performance of alternative option pricing models. *Journal of Finance* 52(5), 2003–2049.
- Bansal, R., D. Kiku, I. Shaliastovich, and A. Yaron (2014). Volatility, the macroeconomy, and asset prices. *Journal of Finance* 69(6), 2471–2511.
- Bansal, R., D. Kiku, and A. Yaron (2012). An empirical evaluation of the long-run risks model for asset prices. *Critical Finance Review* 1, 183–221.
- Bansal, R., S. Miller, D. Song, and A. Yaron (2021). The term structure of equity risk premia. *Journal of Financial Economics* 142(3), 1209–1228.
- 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.
- Bardgett, C., E. Gourier, and M. Leippold (2019). Inferring volatility dynamics and risk premia from the S&P 500 and VIX markets. *Journal of Financial Economics* 131(3), 593–618.
- 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.
- 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. H. and R. S. Koijen (2016). The term structure of returns: Facts and theory. *Journal of Financial Economics* (forthcoming).
- 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. (2018). The VIX premium. *Review of Financial Studies* 32(1), 180–227.
- Choi, H., P. Mueller, and A. Vedolin (2017). Bond variance risk premiums. *Review of Finance* 21(3), 987–1022.
- 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.
- Cochrane, J. H. (2017). Macro-Finance. *Review of Finance* 21(3), 945–985.
- Coval, J. D. and T. Shumway (2001). Expected option returns. *Journal of Finance* 56(3), 983–1009.
- Dew-Becker, I., S. Giglio, A. Le, and M. Rodriguez (2017). The price of variance risk. *Journal of Financial Economics* 123(2), 225–250.
- Giglio, S., B. T. Kelly, and S. Kozak (2023). Equity term structures without dividend strips data. *Journal of Finance* (forthcoming).
- Gormsen, N. J. (2021). Time variation of the equity term structure. *Journal of Finance* 76(4), 1959–1999.
- Gruber, P. H., C. Tebaldi, and F. Trojani (2021). The price of the smile and variance risk premia. *Management Science* 67(7), 4056–4074.
- 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., K. Jacobs, and H. J. Kim (2022). Exploring risk premia, pricing kernels, and no-arbitrage restrictions in option pricing models. Working Paper.
- Heston, S. and S. Nandi (2000). A closed-form GARCH option valuation model. *Review of Financial Studies* 13(3), 585–625.
- Hülsbusch, H. and A. Kraftschik (2018). Consistency between S&P500 and VIX derivatives: Insights from model-free VIX futures pricing. *Journal of Futures Markets* 38(8), 977–995.
- Oh, D. H. and Y.-H. Park (2023). GARCH option pricing with volatility derivatives. *Journal of Banking & Finance* 146, 106718.

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

Rubinstein, M. (1976). The valuation of uncertain income streams and the pricing of options. *Bell Journal of Economics* 7(2), 407–425.

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

Appendix

Estimation procedure

Define daily index returns $R_t = \log(S_t/S_{t-1})$ and the risk-free rate r_t . The return log likelihood is only a function of the GARCH parameters $\Theta = \{\omega, \beta, \alpha, \mu, \gamma\}$

$$\ell_{\text{ret}}(\Theta) = -\frac{1}{2} \sum_{t=1}^T \left[\log h_t + \frac{1}{h_t} \left(R_t - r_t - \left(\mu - \frac{1}{2} \right) h_t \right)^2 \right],$$

where

$$h_1 = \frac{\omega + \alpha}{1 - \beta - \alpha\gamma^2}.$$

Define Black-Scholes vega weighted pricing errors as

$$\varepsilon_i = \frac{P_i^{\text{mkt}} - P_i^{\text{mod}}}{\text{BSV}_i^{\text{mkt}}},$$

where P_i^{mkt} is the market price of option i , $\text{BSV}_i^{\text{mkt}}$ is the market Black-Scholes vega of option i , and P_i^{mod} is the model price for option i . Note that P_i^{mod} depends on both the GARCH parameters Θ as well as the variance risk aversion factor Ξ for the maturity category to which option i belongs. We use four maturity categories and assign different variance risk aversion factors Ξ_1, \dots, Ξ_4 for options with maturities of 20 to 60 days, 60 to 125 days, 125 to 190 days, and 190 to 252 days, respectively. Assume that the Black-Scholes vega weighted pricing errors are i.i.d. normal with mean zero and variance σ^2 . The option likelihood is then a function of $\Theta, \Xi_1, \dots, \Xi_4$, and σ^2 :

$$\ell_{\text{opt}}(\Theta, \Xi_1, \dots, \Xi_4, \sigma^2) = -\frac{1}{2} \sum_{i=1}^N \left(\log \sigma^2 + \frac{\varepsilon_i^2}{\sigma^2} \right)$$

Maximum likelihood can then be used to estimate Θ and Ξ_1, \dots, Ξ_4 ,

$$\left\{ \hat{\Theta}, \hat{\Xi}_1, \dots, \hat{\Xi}_4, \hat{\sigma}^2 \right\} = \underset{\{\Theta, \Xi_1, \dots, \Xi_4, \sigma^2\}}{\text{argmax}} \quad (\ell_{\text{ret}} + \ell_{\text{opt}}).$$

We also follow CHJ in setting $\omega = \omega^* = 0$ in the estimation and μ to its maximum likelihood estimate from just the returns data.

Additional figures




Figure A1: Non-parametric term structure by sample period. The figure shows Sharpe ratios of delta-neutral straddles by maturity and sample period (pre/post Dec 2007). Shaded areas indicate 95% confidence intervals.

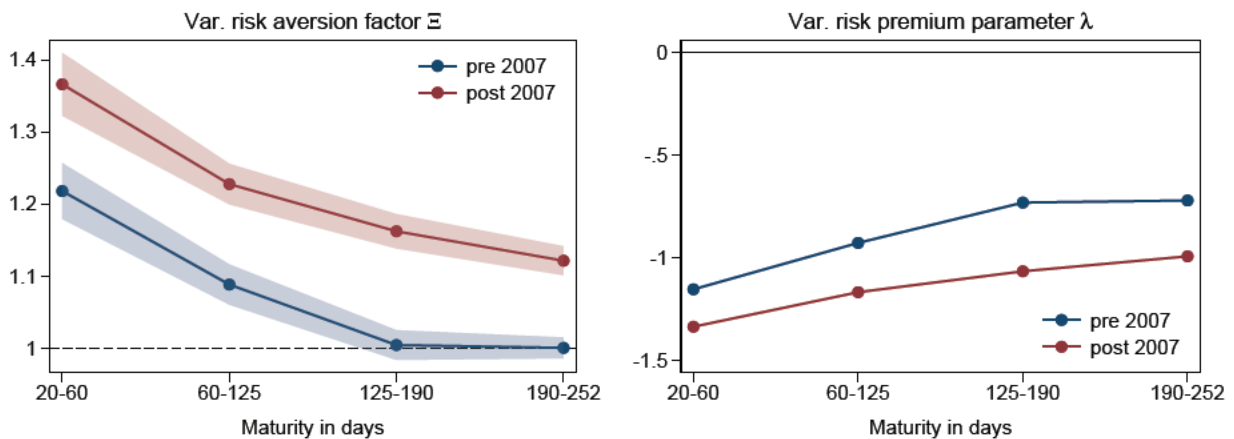


Figure A2: Parametric term structure of variance risk pricing. The figure shows estimates of the variance risk aversion factor Ξ (left panel) and resulting calibrations of the variance risk premium parameter λ (right panel) by maturity and sample period (pre/post Dec 2007). Shaded areas indicate 95% confidence intervals.