STAFF
REPORTS

NO. 736
AUGUST 2015

REVISED
JUNE 2023

The Term Structure of the Price of Variance Risk

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

FEDERAL RESERVE BANK of NEW YORK

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 June 2023
JEL classification: G12, G13

Abstract

We empirically investigate the term structure of variance risk pricing and how it varies over time. Estimating the price of variance risk in a stochastic-volatility option pricing model separately for options of different maturities, we find a price of variance risk that decreases in absolute value with maturity but remains significantly different from zero up to the nine-month horizon. We show that the term structure is consistently 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.

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 Koijen (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 time-series variation 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. Using proprietary data on variance swaps, Dew-Becker et al. (2017) find that investors are willing to pay for insurance against variance risk at the one-month maturity but not at longer maturities, suggesting that the compensation for variance risk also has an important term structure component. In this paper, we focus on the term structure of variance risk pricing and its variation over time. We study two questions. First, is volatility risk never priced at longer horizons? Second, is there important temporal variation underlying the average term structure of volatility risk pricing? To answer these questions, we use standard index options data and the simple one-factor stochastic variance model of Christoffersen, Heston, and Jacobs (2013) which yields a single parameter capturing the per-unit price of variance risk.[2] 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 price of variance risk separately for options in distinct maturity buckets. This allows us to give clear answers

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

[2] Note that a positive price-of-variance-risk parameter implies that investors require compensation for exposure to variance risk. The literature defines the variance risk premium as the difference between expected variance under the physical and risk neutral measures. This is the payoff to being long variance and tends to be negative, consistent with investors requiring compensation for bearing variance risk (Carr and Wu, 2009). To avoid confusion, we tend to think in terms of positive numbers throughout, i.e. the absolute value of the price of variance risk.

to our first-order questions whether long horizon variance risk is priced and whether the relative pricing at different maturities (per unit of risk) varies systematically over time.

For our unconditional estimates, we find a significant price of variance risk for maturities up to nine months (options with 125 to 190 trading days to expiry). This finding is in contrast to Dew-Becker et al. (2017), who find a significant variance risk premium only at the shortest, one-month horizon. Our finding is robust to estimating the physical process of returns separately from options prices, and to allowing the expected process of returns to vary by maturity. The differences between our results and those of Dew-Becker et al. (2017) may be in part due to the more liquid nature of the market for index options we use to estimate the price of variance risk relative to their results which are found in the significantly less liquid market for variance swaps. In addition, we find that the term structure of the price of variance risk (in absolute value) is downward sloping, with short-maturity options exhibiting significantly higher risk premia than longer maturity options across maturities of up to a year.

Finally, 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 and low and high price-dividend ratios. These differing market conditions shift the level of the term structure of the compensation for variance risk, but the slope remains downward sloping in our estimates. That the slope of the term-structure of variance risk premia always remains downward sloping in the time series starkly contrasts to similar analysis on the term-structure of the equity risk premia in Gormsen (2021), Bansal et al. (2021), and Giglio, Kelly, and Kozak (2023).

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, have important implications for asset pricing models. Our first result, the strictly non-zero price of variance risk at short and medium horizons, implies that not just the immediate volatility shocks but also the shocks to expected volatility are priced; a fundamental difference to the results of Dew-Becker et al. (2017). This 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 result, that the average term-structure of variance risk premia is downward sloping, challenges the standard models. While the long-run-risk model of Bansal and Yaron (2004) 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). Finally, models that capture the variation in the slope of equity risk premia via variation in risks, such as Gormsen (2021) and Bansal et al. (2021), may not be able to match the downward sloping term-structures we obtain under all market conditions, our third result. Capturing both the upward/downward variations in the term-structure of equity premia and the constantly downward sloping term-structure of variance risk premia that we document constitutes a challenge and invites new theories, e.g. Andries, Eisenbach, and Schmalz (2023) who rationalize the evidence on equity risk premia via variations in market liquidity.

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 or "model-free" results that are inconsistent with a constant price of variance risk, but consistent with a horizon-depend price of risk, as derived in 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 sharply different and complementary to 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. Barras and Malkhozov (2016) find that institutional factors help explain differences in estimates of variance risk premia in the equity and option markets.

The paper proceeds as follows. Section 2 presents our data sources and parametric estimation procedure. Section 3 gives the empirical results and Section 4 provides robustness checks. Section 5 concludes.

2 Data and empirical results

2.1 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 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

Table 1: 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

the put. We convert all put prices to call prices by put-call parity, ignoring dividends.

Table 1 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.

2.2 Parametric procedure

We follow Christoffersen, Heston, and Jacobs (2013, hereafter CHJ) and Heston and Nandi (2000) to describe the dynamics of stock return and variance. Specifically we assume that the stock price S_t follows a GARCH-in-means process and the one-period excess return has variance h_t, as follows:

$$\log S_t = \log S_{t-1} + r_t + \left(\mu - \tfrac{1}{2}\right) h_t + \sqrt{h_t}\, z_t$$
$$h_t = \omega + \beta h_{t-1} + \alpha \left(z_{t-1} - \gamma \sqrt{h_{t-1}}\right)^2,$$

with $z_t \sim \mathcal{N}(0,1)$, where μ governs the equity premium, ω governs the unconditional mean of the shock process, β governs its persistence, α captures the kurtosis of the distribution, i.e. how fat the tails of the variance process are, and γ makes the distribution asymmetric and captures the correlation of the variance with the stock return. The stochastic discount factor is given by

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

where δ and η capture time preferences, and ϕ captures aversion to equity risk. Our focus is on the parameter ξ which multiplies variance in the stochastic discount factor and therefore captures aversion to variance risk "per unit of variance." We therefore refer to ξ as the price of variance risk (PVR) parameter.

Given the physical GARCH parameters $\Theta = \{\omega, \beta, \alpha, \mu, \gamma\}$ and the PVR parameter ξ, the equity risk aversion ϕ is pinned down as $\phi = -\left(\mu - \frac{1}{2} + \gamma\right)(1 - 2\alpha\xi) + \gamma - \frac{1}{2}$ (see Heston and Nandi (2000) and CHJ for additional details). CHJ show that the processes can be risk-neutralized as

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

with

$$h_t^* = \frac{1}{1 - 2\alpha\xi} h_t, \quad \omega^* = \frac{1}{1 - 2\alpha\xi} \omega,$$
$$\alpha^* = \frac{1}{1 - 2\alpha\xi} \alpha, \quad \gamma^* = \gamma - \phi,$$

and $z_t^* \sim \mathcal{N}(0,1)$. To compensate for variance risk, the risk neutral variance process has a higher long-run mean and higher persistence for $\xi > 0$. The PVR parameter ξ therefore directly translates into the variance risk premium $E[h_t] - E[h_t^*]$, entering multiplicatively and therefore capturing the per-unit pricing of variance risk.

CHJ estimate the GARCH parameters and a common PVR parameter ξ jointly with a likelihood that incorporates both returns and option prices. We follow their approach, except that 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. Other than that, we adapt their joint estimation procedure for the full sample but, in contrast to them, allow ξ to vary by maturity category.

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 PVR parameter ξ for the maturity category to which option i belongs. We use four maturity categories and assign different PVR parameters $\xi_1, \ldots, \xi_4$ for options with maturities of 20 to 60 days, 60 to 125 days, 125 to 190 days, and 190 to 252 days, respectively. Define $\Xi = \{\xi_1, \ldots, \xi_4\}$. 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 Θ, Ξ, and σ^2:

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

Maximum likelihood can then be used to estimate both Θ and Ξ,

$$\left\{\widehat{\Theta}, \widehat{\Xi}, \widehat{\sigma}^2\right\} = \underset{\{\Theta, \Xi, \sigma^2\}}{\text{argmax}}\left(\ell_{\text{ret}} + \ell_{\text{opt}}\right).$$

2.3 Sample splits

Our interest is in whether estimated PVR parameters vary 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 par-

Table 2: 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							
$\leq$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

ticular, 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 2 reports summary statistics over these subsamples. Times of high volatility, 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 analysis below is designed to determine how much of these differences are due to the underlying index return process versus the price of variance risk.

3 Estimation results

3.1 Term structure of the price of variance risk

Figure 1 shows the term structure of the PVR parameter ξ estimated on the full sample. The paramater is significantly positive ($\xi > 1$) at all maturities except the longest maturity (190–252 days), indicating that investors require compensation for bearing variance risk and considerably more so (per unit of variance) at shorter horizons. Table 3, column (1)

Table 3: Estimation on full sample and subsamples. The table shows 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, 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)
	Full sample	VIX split		PD ratio split		GDP growth split	
		High	Low	Low	High	Low	High
β	0.735	0.747		0.725		0.733	
	(0.004)	(0.003)		(0.004)		(0.003)	
α	1.3×10^{-6}	1.1×10^{-6}		1.3×10^{-6}		1.3×10^{-6}	
	(0.02×10^{-6})	(0.02×10^{-6})		(0.03×10^{-6})		(0.03×10^{-6})	
γ	451.24	463.62		459.38		445.72	
	(5.72)	(4.35)		(6.55)		(6.14)	
ξ_{20-60}	1.212	1.896	1.065	1.377	1.166	1.321	1.169
	(0.018)	(0.020)	(0.019)	(0.026)	(0.020)	(0.026)	(0.020)
ξ_{60-125}	1.108	1.621	1.020	1.205	1.082	1.198	1.065
	(0.013)	(0.015)	(0.014)	(0.018)	(0.015)	(0.015)	(0.014)
$\xi_{125-190}$	1.025	1.414	1.000	1.116	1.007	1.105	1.000
	(0.010)	(0.011)	(0.010)	(0.014)	(0.011)	(0.013)	(0.013)
$\xi_{190-252}$	1.000	1.320	1.000	1.056	1.000	1.048	1.000
	(0.009)	(0.010)	(0.009)	(0.012)	(0.009)	(0.011)	(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	

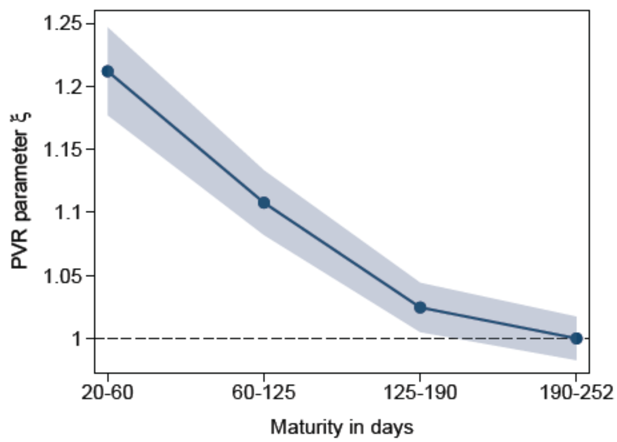

Figure 1: Term structure of variance risk pricing. The figure shows estimates of the PVR parameter ξ by maturity on the full sample (Table 3, column 1). Shaded areas indicates 95% confidence intervals.

shows the full estimation results, including the GARCH parameters α, β and γ. 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. Estimates of the parameters of the GARCH process are similar to those in CHJ, who in their joint estimation find $\widehat{\beta} = 0.756$, $\widehat{\alpha} = 1.41 \times 10^{-6}$ and $\widehat{\gamma} = 515.57$. Our estimated GARCH parameters are lower than theirs, implying 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 volatlity period from 2013 to 2018 as well as spikes in volatility in 2010 and 2011.

3.2 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 the PVR parameter ξ separately on subsamples of the data. Figure 2 shows estimates of the term structure of ξ when splitting the sample into high/low VIX, price-dividend ratio, and GDP growth. Table 3, columns (2) to (7) show the full 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 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 PVR parameter is significantly greater than 1 out to the 60–125 day bin, showing that the result that

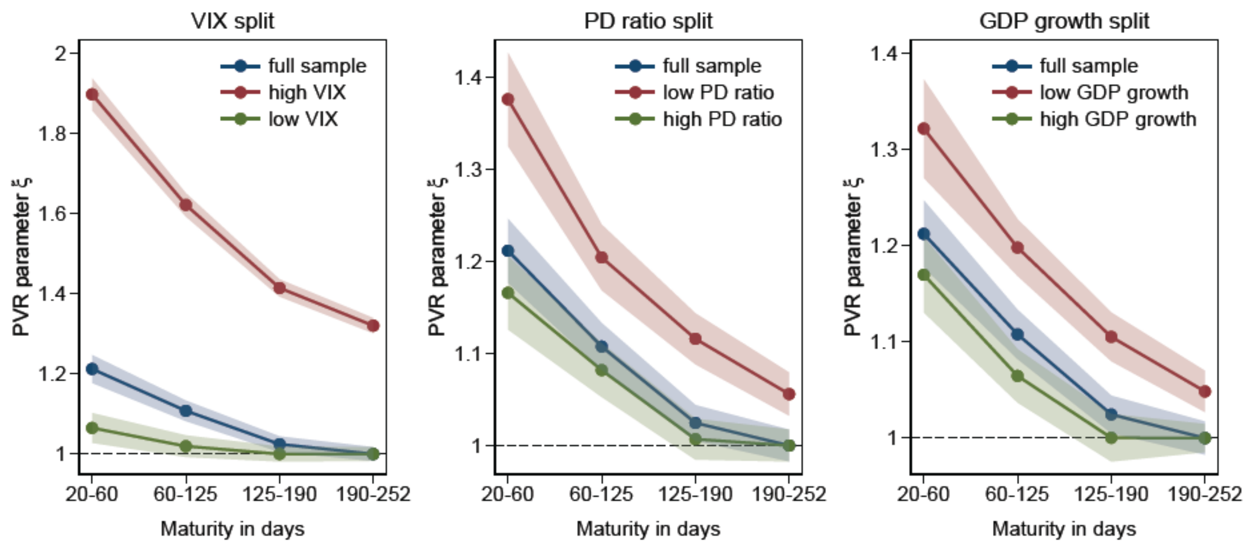

Figure 2: Term structure in different subsamples. The figure shows estimates of ξ on subsamples and compares them to the full-sample estimates (Table 3, columns 2–7). Shaded areas indicates 95% confidence intervals.

longer-maturity options display a significant price of variance risk persists even in these subsamples.

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 subsample split, again emphasizing the close relationship between the VIX and the price of variance risk (bottom of Table 3).

4 Robustness

4.1 Joint vs. sequential estimation

Our baseline specification follows CHJ in estimating the PVR parameter ξ jointly with the GARCH parameters α, β, and γ. Figure 3 shows the coefficient estimates for ξ if we do the

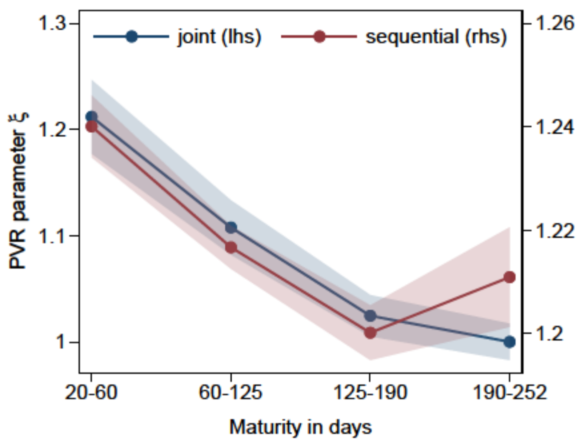

Figure 3: Estimating parameters jointly or sequentially. The figure shows estimates of ξ by maturity on the full sample, comparing the joint estimation of GARCH parameters and PVR parameter (Table 4, column 1) to a sequential estimation, first estimating the GARCH process using returns and then estimating ξ using option prices (Table 4, column 2). Shaded areas indicates 95% confidence intervals.

estimation in two sequential steps, first estimating the GARCH parameters without making use of the options data and then estimating the PVR parameter 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.

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.

4.2 Allowing the GARCH parameters to vary with horizon

Our benchmark specification only allows the PVR parameter ξ to vary with the horizon but maintains a single set of GARCH parameters α, β, and γ. Figure 4 shows the coefficient estimates if we allow the GARCH parameters to vary with the horizon. Table 4, columns

Table 4: Sequential estimation and allowing the GARCH parameters to vary with horizon. The table shows 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, 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)
	Joint	Sequential	Joint, by maturity			
β	0.735	0.819	0.714	0.705	0.775	0.787
	(0.004)	(0.013)	(0.006)	(0.007)	(0.009)	(0.006)
α	1.3×10^{-6}	3.9×10^{-6}	1.5×10^{-6}	1.2×10^{-6}	1.2×10^{-6}	1.1×10^{-6}
	(0.02×10^{-6})	(0.0×10^{-6})	(0.03×10^{-6})	(0.04×10^{-6})	(0.05×10^{-6})	(0.04×10^{-6})
γ	451.24	193.63	429.69	488.67	417.86	436.44
	(5.72)	(1.94)	(5.14)	(10.19)	(4.42)	(11.50)
ξ_{20-60}	1.212	1.240	1.227			
	(0.018)	(0.003)	(0.015)			
ξ_{60-125}	1.108	1.217		1.141		
	(0.013)	(0.002)		(0.019)		
$\xi_{125-190}$	1.025	1.200			1.000	
	(0.010)	(0.003)			(0.020)	
$\xi_{190-252}$	1.000	1.211				1.000
	(0.009)	(0.005)				(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

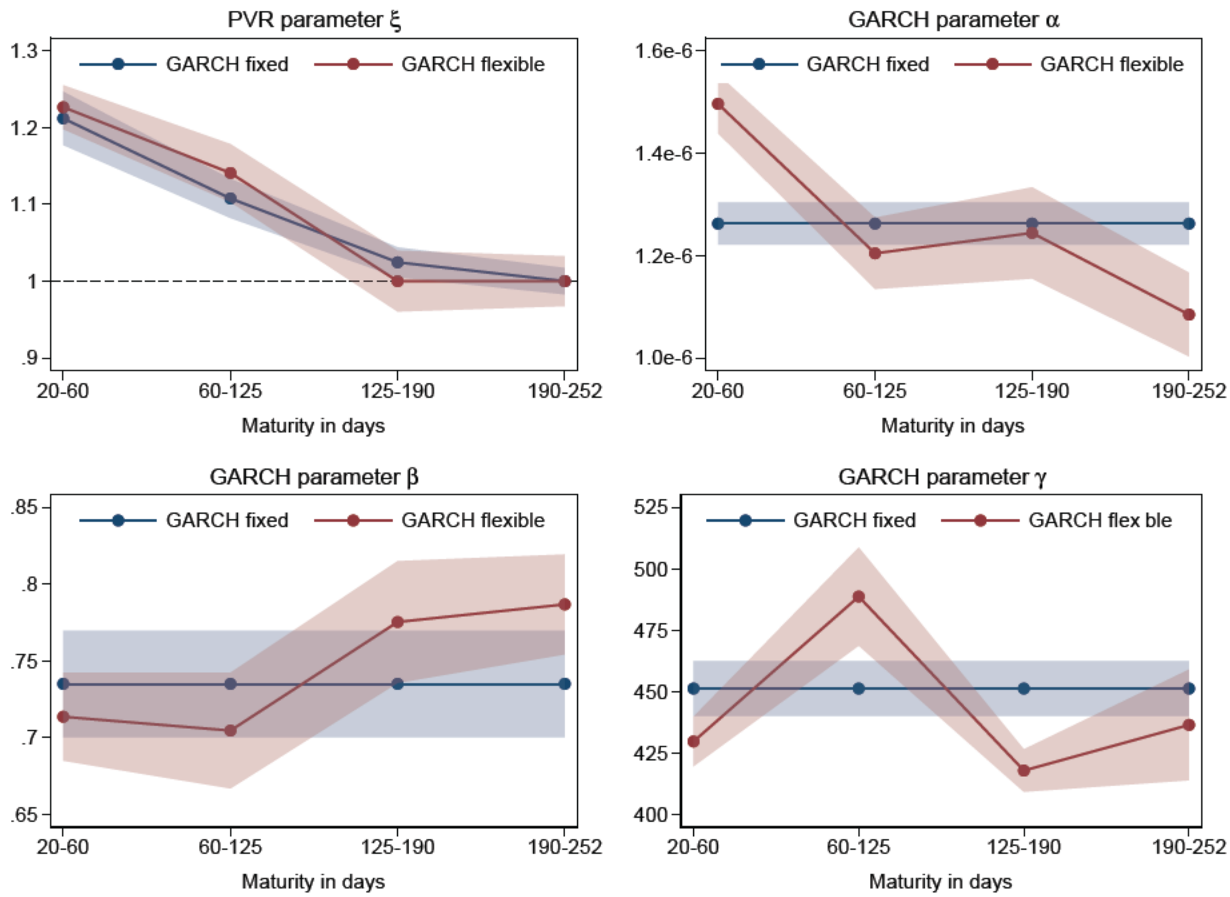

Figure 4: Allowing the GARCH parameters to vary with horizon. The figure shows estimates of ξ and the GARCH parameters, comparing the estimation where GARCH parameters are fixed across maturities (Table 4, column 1) to the estimation which allows both GARCH parameters and ξ to vary by maturity (Table 4, columns 3–6). Shaded areas indicates 95% confidence intervals.

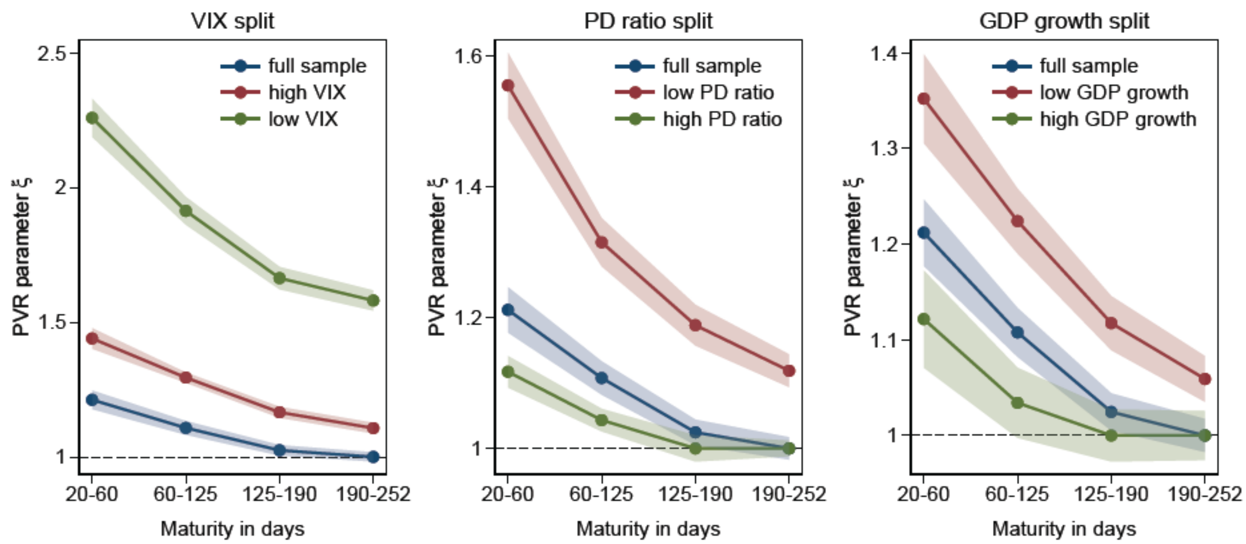

Figure 5: Allowing GARCH parameters to vary with sample splits. The figure shows estimates of ξ on subsamples, allowing for the GARCH parameters to vary with the sample splits (Table 5). Shaded areas indicates 95% confidence intervals.

(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 is negligible (top-left panel of Figure 4). 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.

4.3 Allowing GARCH parameters to vary with sample splits

Our benchmark specification estimates one set of GARCH parameters and only estimates the PVR parameter ξ separately on different subsamples. Due to the auto-regressive nature of the process, estimating separate GARCH parameters for different 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 5 shows the coefficient estimates of 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 is negligible if we split the sample on PD ratio and on GDP growth.

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

However, in case of the VIX split, the ordering of the term structures changes as now the PVR parameter 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 distaste for 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 the PVR parameter 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 VRP 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.

5 Conclusion

We provide 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) model, based on the empirical approach developed by Christoffersen, Heston, and Jacobs (2013). We find that the price of insurance against increases in volatility 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. 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. A comparative advantage of our approach to the existing literature is a focus on the price per unit 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, Eisenbach, and Schmalz (2023).

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

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. 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.