

Systematic Variance Risk Everywhere in Equity Option Markets

March 16, 2026

ABSTRACT

We propose systematic variance risk as the primary factor for pricing the cross-section of equity option returns. Using a parsimonious factor model with time-varying risk premiums, we find pervasive negative variance risk premiums in index, stock, and ETF options, consistent with asset pricing theory. The magnitudes of the risk premiums are plausible and similar across markets and the cross-section of options. This highlights the central and unifying role of systematic variance risk in option markets. We emphasize aspects of model specification and empirical implementation which play a critical role in establishing this stylized fact and explain differences with existing studies.

JEL Classification: G10; G12; G13.

Keywords: Option Returns; Factor Models; Systematic Variance Risk; Option-Implied Factor Risk Premiums; Time-Varying Exposures; Machine Learning.

1. Introduction

Option prices are highly sensitive to fluctuations in the volatility of the underlying security. Systematic variance risk—which we proxy by the change in aggregate stock market variance—therefore emerges as the pre-eminent candidate risk factor for explaining the cross-section of equity option returns. Indeed, determining whether systematic variance risk is priced in equity option returns can be regarded as the option market analogue of assessing the stock market factor’s role in pricing the cross-section of equity returns.¹ In the absence of frictions, theory predicts that systematic variance risk should carry a similar negative premium across markets;² however, existing evidence that it is priced is largely confined to index options, and the extent of its role in broader equity option markets remains an open question.

This paper conclusively demonstrates that systematic variance risk is pervasively priced in equity option returns across index, stock, and ETF markets.³ Consistent with theoretical predictions, we find a negative systematic variance risk premium in all three markets. Moreover, the magnitude of the systematic variance risk premium estimate is plausible and very similar in the three option markets we investigate, as well as for separate samples of call and put options, and across option maturities. Our findings therefore also provide important new evidence that these three equity option markets are highly integrated.

The existing literature documents a negative price of variance risk in index option markets, but reports mixed evidence based on stock options, and lacks evidence on ETF options. Our finding of a negative systematic variance risk premium in all three markets therefore complements earlier evidence by documenting compensation for variance risk that is consistent across option markets. This is due to three features of our empirical framework and implementation that distinguish our contribution from the existing literature: 1) The scope of our analysis is much wider; 2) We study the pricing of systematic variance risk, as

¹See Giglio and Xiu (2021) and Bryzgalova, Huang, and Julliard (2023) for recent evidence on the performance of the market factor in the cross-section of equity returns.

²The existence of a systematic variance risk factor is supported by option pricing models as well as numerous asset pricing models—including the intertemporal CAPM, recursive-preference consumption models, and long-run risk frameworks—which predict that the variance risk premium is linearly and negatively related to the level of aggregate variance.

³Henceforth we refer to options on common stocks as stock options. We use the terminology “equity option” to refer collectively to index, stock, and ETF options.

opposed to total variance risk;⁴ 3) Our estimates rely on a factor model that accounts for the state-dependent nature of option returns. We now discuss these three important aspects in more detail.

First, the scope of our empirical analysis is extensive. We estimate the compensation for systematic variance risk across three option markets: a sample with options on the S&P 500 index, a sample with options on 422 common stocks, and a sample of options on 96 ETFs. With more than 12 million observations across these three samples, spanning a wide range of maturity and moneyness over identical sample periods, our study therefore provides the most comprehensive analysis to date of the role of systematic variance risk in equity option markets. We apply a unified empirical framework to these three samples, ensuring that the resulting estimates are directly comparable.

The second important feature of our contribution is that we study the pricing of *systematic* variance risk and we compare its sign and magnitude across the three option markets. This focus stems not only from the theoretical motivation discussed above, but also from a desire to identify and compare the same construct across different markets. This is critical because of the stylized facts mentioned above: while index option studies consistently find a negative price of variance risk, the evidence based on stock options is mixed. However, most existing studies of stock options estimate compensation for *total* variance risk, as opposed to *systematic* variance risk. While theory suggests that idiosyncratic variance risk should bear no compensation, some empirical studies have argued that this may not be the case. Our estimates are not affected by assumptions regarding idiosyncratic variance risk and therefore facilitate comparisons across option markets.

Third, we analyze systematic variance risk in these three option markets using a factor model tailored to accommodate specific features of option returns and the dynamics of the variance risk premium. Specifically, we rely on the conditional factor model for option returns in Fournier, Jacobs, and Orłowski (2024). This approach captures the nonlinear and time-varying nature of option exposures in a way that is consistent with option pricing theory, without imposing restrictive distributional assumptions. We model factor loadings as

⁴The total variance risk premium refers to the difference between a security's total variance under the historical and risk-neutral measures. The systematic variance risk premium is the risk premium associated with systematic variance risk, which as mentioned above we proxy by the change in aggregate stock market variance.

nonparametric functions of option characteristics, which allows exposures to vary flexibly and nonlinearly across contracts and over time. We use this approach to study models with a systematic variance risk factor. Our baseline model is a parsimonious one. After delta-hedging, this model contains a systematic variance risk factor, measured by changes in the squared VIX, and the underlying-specific squared return.

We show that unbiased and reliable estimation of systematic variance risk premiums with this setup requires several key model ingredients: (i) delevered option returns, which address measurement errors in option prices and returns (Duarte, Jones, and Wang, 2024); (ii) joint, nonparametric estimation of option deltas and other exposures, which capture nonlinear relations between option characteristics, economic conditions, and returns with greater accuracy; and (iii) a conditional variance risk premium, which varies with the level of aggregate variance, consistent with theory.

Our empirical investigation yields important novel insights regarding the role of systematic variance risk in the cross-section of equity option returns. First, systematic variance risk contributes meaningfully to explaining the variation in option returns in all three markets, and a specification search across additional observable and latent factors confirms that the baseline model performs very well. Second, we find that systematic variance risk is priced in the three option markets we investigate and that the price of variance risk is consistently negative across index, stock, and ETF options. This result also holds across option types, maturities, and industries. Identifying factors that command premiums in option returns—particularly factors that are directly observable and economically meaningful—has proven difficult, which makes our finding especially consequential. Third, the average conditional risk premiums of our baseline specification are similar across markets, with values of -20.3% for index options, -18.5% for ETFs, and -18.8% for stock options. These estimates are statistically significant within each market and do not differ significantly across markets, suggesting that these three equity option markets are highly integrated. Our fourth finding addresses the time variation in the price of systematic variance risk. In each of the three option markets, the estimated premium loads negatively on the aggregate level of variance as measured by the square of the level of the VIX index. That is, the conditional variance risk premium always becomes larger (more negative) when economic uncertainty increases, in line with modern asset pricing theories.

In summary, we robustly identify the existence of an economically meaningful systematic variance risk premium “everywhere” in equity option markets. We conclude that systematic variance risk plays a central role across option markets, and should be part of any factor model for options. This central role of systematic variance risk is consistent with many theoretical models, and therefore analogous to the role of the market factor in factor models for stocks and corporate bonds.

The remainder of the paper is organized as follows. Section 2 provides a literature review. Sections 3 and 4 describe the empirical methodology, data and variable construction. Section 5 presents the baseline empirical results, including model explanatory power and estimates of the variance risk premiums. Section 6 highlights key modeling choices required to obtain reliable estimates of the systematic variance premium, and Section 7 presents robustness exercises and additional analysis. Section 8 concludes.

2. Literature Review

We now provide a more detailed overview of the existing literature on the VRP. We first discuss the evidence on the sign and magnitude of the VRP in different option markets. We then highlight the problems in comparing existing estimates. Subsequently we discuss the existing literature on factor models of option returns, and we conclude by summarizing the literature on variance risk in the cross-section of stock returns.

The Sign and Magnitude of the VRP

The lion’s share of the literature on the VRP investigates its presence in index markets. While the debate about the exact size of the premium, its dynamics, and economic drivers is ongoing, there is consensus on the existence of a large negative VRP in index options.⁵

The literature measuring either the systematic or the total VRP in the cross-section of stock options is much less voluminous and contains conflicting findings. In early work, Bakshi and Kapadia (2003) conclude that the total VRP is negative on average, but positive for many firms. Similarly, Carr and Wu (2009) document substantial cross-sectional variation in total VRP estimates that are on average negative. Using S&P 100 stock options, Driessen,

⁵See for instance Pan (2002), Todorov (2010), and Bollerslev, Tauchen, and Zhou (2009) among many others.

Maenhout, and Vilkov (2009) find no evidence of a negative risk premium. In contrast, in recent work Duarte, Jones, and Wang (2024) forcefully argue for the presence of a negative VRP in stock options.

Comparing VRP Estimates Across and Within Option Markets

One of the central questions in the VRP literature is its relative importance and magnitude in different markets. However, comparing existing estimates of variance risk premiums *across* and even *within* given segments of the equity options market is far from straightforward. Studies that analyze multiple option markets using the same empirical methodology and sample period are naturally better-suited for comparisons, but the problems with comparing existing estimates run deeper.

To illustrate the differences between existing VRP studies, first consider the literature on the VRP estimated from index option markets. An important part of this literature estimates the VRP based on index returns and options to calibrate parametric option pricing models.⁶ Another strand of this literature also computes the VRP as the difference between the historical and implied volatility, but relies on a more nonparametric approach (Carr and Wu, 2009; Bollerslev, Tauchen, and Zhou, 2009). Yet another approach to extract the index VRP uses variance swap data instead of index options (Aït-Sahalia, Karaman, and Mancini, 2020). Moreover, the VRP is sometimes modeled and measured based on the level of variance, while other estimates are based on variance innovations.

Comparisons between results from the literature on the index VRP and studies that go beyond index markets are further complicated because most of the latter studies focus on the analysis of the *total* VRP.⁷ Moreover, studies in this literature exhibit an even wider variation in their empirical strategies, the types of derivative securities considered, and the proxies used for variance risk. Some studies use variance swap methodologies (see Carr and Madan (1998); Britten-Jones and Neuberger (2000) and Carr and Wu (2009)) while others analyze average delta-hedged stock option returns (Bakshi and Kapadia, 2003). More recent work

⁶See, among others Bates (2000), Pan (2002), Jones (2003), Eraker (2004), and Andersen, Fusari, and Todorov (2017). See Hurn, Lindsay, and McClelland (2015) for a comprehensive review.

⁷Christoffersen, Fournier, and Jacobs (2018) use CAPM-style dynamic models with systematic variance risk to analyze stock options. While this approach enables a clear identification of systematic variance risk, the systematic VRP in this type of models is inferred from index data and then imposed on stock options, which limits the ability to compare VRP estimates in the two markets.

uses the cross-section of stock option returns primarily relying on nonparametric portfolio sorts and cross-sectional Fama-MacBeth (FMB) regressions to study the VRP as in Driessen, Maenhout, and Vilkov (2009) or Duarte, Jones, and Wang (2024).⁸ Duarte, Jones, and Wang (2024) use delevered Black-Scholes vegas to measure the total variance risk exposure of stock options, and estimate the VRP with FMB regressions. Goyal and Saretto (2009) identify a premium for trading volatility by forming a long-short portfolio of straddle and delta-hedged call option returns from a decile sort on the difference between historical and implied volatility.

While these empirical methodologies are all valid in their own right, the heterogeneity in design and implementation further complicates direct comparisons of risk premium estimates across different sample periods, and market segments. We contribute to this literature by providing a unified approach to estimate the systematic VRP across option markets with the same empirical methodology, same sample period, and the same proxy for variance risk.

Factor Models of Option Returns

Our empirical implementation relies on a factor model of option returns. This paper is therefore part of a rapidly expanding area of research that explicitly models option returns rather than option prices.⁹ While the analysis of option returns faces measurement problems because option returns are, from a microstructure perspective, very noisy (Duarte, Jones, and Wang, 2024), it provides a more natural connection to the literature that analyzes the underlying securities, which focuses almost exclusively on returns. Whilst the majority of prior research on stock option returns focuses on characteristic pricing,¹⁰ this paper fits into a smaller segment of the literature that analyzes the cross-section of returns through the lens of observable, economically motivated factors (Jones, 2006). Moreover, our focus on systematic variance risk differs from prior research, which predominantly examines the underlying’s total

⁸Additionally, Carr and Wu (2009) and Driessen, Maenhout, and Vilkov (2009) compute the VRP as the difference between the physical and risk-neutral total stock variance measures. If idiosyncratic risk is priced (Herskovic, Kelly, Lustig, and Van Nieuwerburgh, 2016; Bégin, Dorion, and Gauthier, 2020), the resulting estimate differs from the systematic variance risk premium. The difference can be compounded by cross-sectional differences in market betas.

⁹See, among others, Jones (2006); Broadie, Chernov, and Johannes (2009); Horenstein, Vasquez, and Xiao (2025), Heston, Jones, Khorram, Li, and Mo (2023) and Duarte, Jones, and Wang (2024).

¹⁰This literature typically relies on nonparametric portfolio sorts and FMB cross-sectional linear regressions. See for instance Goyal and Saretto (2009), Christoffersen, Goyenko, Jacobs, and Karoui (2018), Cao and Han (2013), and Cao, Han, Tong, and Zhan (2022).

variance or option exposures to total variance (see, e.g., Driessen, Maenhout, and Vilkov, 2009; Duarte, Jones, and Wang, 2024, among others). Instead of computing the difference between the conditional expectations of the total variance *level* under the physical and risk-neutral measures, we focus on the pricing of systematic variance *increments*, a specification choice consistent with the literature on factor models.

Our empirical methodology also sets our work apart from most of the literature on factor models of option returns, and allows us to study a broad array of option contracts spanning varying degrees of maturity and moneyness. Some empirical studies of the cross-section of option returns rely on principal components analysis (PCA) (Israelov and Kelly, 2017; Horenstein, Vasquez, and Xiao, 2025) or instrumented principal components analysis (IPCA) (Büchner and Kelly, 2022; Goyal and Saretto, 2024).¹¹ We instead use the approach in Fournier, Jacobs, and Orłowski (2024), a factor model with time-varying exposures and risk premiums, which builds on Gagliardini, Ossola, and Scaillet (2016). This approach allows the analysis of large panels of option returns while instrumenting option exposures, but the factors are economic variables instead of principal components. The use of observed economic factors enables a consistent analysis of risk pricing across multiple markets—an objective that would be more difficult to achieve using principal components. Our objectives are instead more closely aligned with the early work of Jones (2006) on index option returns.

Finally, we document the performance of pricing factors across three different option markets. Jones (2006), Büchner and Kelly (2022), and Fournier, Jacobs, and Orłowski (2024) study factor models of index option returns. Goyal and Saretto (2024) and Horenstein, Vasquez, and Xiao (2025) analyze the cross-section of stock option returns. Our focus on multiple markets and segments within these markets (e.g., specific industries, maturities, option types) provides unique insights into the similarities and differences in the pricing of systematic variance risk. Moreover, it enables us to comment on the extent to which these markets are integrated.

Variance Risk in the Cross-Section of Stock Returns

Because we study delta-hedged option returns, our empirical findings are by construction not directly informative about the cross-section of the underlying returns on common stocks.

¹¹Bali, Beckmeyer, Mörke, and Weigert (2023) use machine learning techniques to predict the cross-section of stock returns.

Nevertheless, it is important to mention the existing literature on the pricing of systematic variance risk in the cross-section of stock returns, which uses a variance factor very similar to ours. Ang, Hodrick, Xing, and Zhang (2006) show that systematic variance risk is priced in the cross-section of equity returns. Detzel, Duarte, Kamara, Siegel, and Sun (2023) document a negative systematic VRP in equity returns in a wide variety of samples.¹² In a different strand of literature, Barras and Malkhozov (2016) document discrepancies between the magnitudes of the systematic VRPs in the stock and index option markets.

3. Model and Estimation

This section presents a conditional factor model for delevered excess equity option returns. We implement the framework used by Fournier, Jacobs, and Orłowski (2024) to model index options. We specify a conditional factor structure in equity option returns with minimal parametric assumptions beyond the choice of observable factors and characteristics for the exposures. For brevity and expositional simplicity, our discussion in this section focuses on the baseline model. Internet Appendix B provides a discussion of the more general framework.

By relying on observable rather than latent factors, our framework is particularly well-suited for analyzing and comparing cross-sections of options across various underlying securities and for studying the pricing of systematic risks across markets.

3.1. A Theory of Delevered Excess Equity Option Returns

Market, ETF, and stock dynamics. Consider an economy in which the dynamics of the market index $S_{M,t}$ and its stochastic variance $\sigma_{M,t}^2$ are characterized by the following equations under the physical measure P :

$$\frac{dS_{M,t}}{S_{M,t}} = \mu_t dt + \sigma_{M,t} dW_t^M \quad (1)$$

$$d\sigma_{M,t}^2 = \kappa_M(\theta_M - \sigma_{M,t}^2) dt + \sigma_{M,t} \delta_M dW_t^{\sigma_M}, \quad (2)$$

¹²These studies are related to but distinct from a voluminous literature that studies (idiosyncratic) volatility as a cross-sectional predictor of returns, as well as a literature that studies factors based on idiosyncratic volatility (Ang et al., 2006; Herskovic et al., 2016).

where μ_t is the expected market return. The systematic market variance $\sigma_{M,t}^2$ follows a square-root process, with parameters κ_M , θ_M , and δ_M representing the mean-reversion speed, long-run mean, and volatility of variance, respectively. The Brownian motions dW_t^M and $dW_t^{\sigma_M}$ may be correlated, though we do not impose additional restrictions on their correlation structure at this stage.

The economy features a cross-section of stocks and ETFs, denoted by $S_{j,t}$ for $l \neq M$, which follow a one-factor structure:

$$\frac{dS_{j,t}}{S_{j,t}} = (r^f - q_l + \alpha_l) dt + \beta_l \left(\frac{dS_{M,t}}{S_{M,t}} - r^f dt \right) + \sigma_{j,t} dW_t^l, \quad (3)$$

$$d\sigma_{j,t}^2 = \kappa_l(\theta_l - \sigma_{j,t}^2) dt + \sigma_{j,t} \delta_l dW_t^{\sigma_l}, \quad (4)$$

where r^f is the constant risk-free rate, q_l denotes the dividend yield, α_l captures any pricing error or alpha, β_l is the beta exposure, dW_t^l is a standard Brownian motion representing idiosyncratic price risk, and $\sigma_{j,t}$ is the stochastic volatility of idiosyncratic returns. The idiosyncratic variance follows a square-root process with analogous parameter interpretations. The Brownian motions dW_t^l and $dW_t^{\sigma_l}$ may also be correlated. We discuss below the potential dependence of $dW_t^{\sigma_l}$ on systematic risk factors and the possibility of cross-sectional correlation in idiosyncratic variance shocks across securities.

Option returns. Consider O_t^{ij} , the value of equity option i written on underlying security $S_{j,t}$ with $j \in \{M, l\}$.¹³ Under no arbitrage, there exists a risk-neutral measure Q such that the option earns the risk-free rate in expectation: $E_t^Q[dO_t^{ij}] = r^f O_t^{ij} dt$. Combining this restriction with Itô's lemma yields the option excess return dynamics:

$$\begin{aligned} \left(\frac{dO_t^{ij}}{O_t^{ij}} - r^f dt \right) &= \frac{1}{O_t^{ij}} \frac{\partial O^{ij}}{\partial \sigma_M^2} \left(d\sigma_{M,t}^2 - E_t^Q[d\sigma_{M,t}^2] \right) + \frac{S_{j,t}}{O_t^{ij}} \frac{\partial O^{ij}}{\partial S_j} \left(\frac{dS_{j,t}}{S_{j,t}} - r^f dt \right) \\ &\quad + \mathbf{1}_{\{j=l\}} \frac{1}{O_t^{ij}} \frac{\partial O^{ij}}{\partial \sigma_l^2} \left(d\sigma_{j,t}^2 - E_t^Q[d\sigma_{j,t}^2] \right), \end{aligned} \quad (5)$$

where $\mathbf{1}_{\{j=l\}}$ equals one for stock and ETF options and zero otherwise. Option excess returns exhibit a conditional factor structure with time-varying exposures levered by $1/O_t^{ij}$.

¹³Technically, $O_t^{ij} = O^{ij}(S_{j,t}, \sigma_{M,t}^2)$ for $j = M$ (index options) and $O_t^{ij} = O^{ij}(S_{j,t}, \sigma_{M,t}^2, \sigma_{j,t}^2)$ for $j = l$ (stock or ETF options).

Delevered option returns. Following Büchner and Kelly (2022) and Fournier, Jacobs, and Orłowski (2024), our empirical analysis focuses on the modeling of delevered returns. After deleveraging by multiplying equation (5) by $\frac{O_t^{ij}}{S_{j,t}}$, we have the following representation for delevered option returns:

$$\frac{O_t^{ij}}{S_{j,t}} \cdot \left(\frac{dO_{t+1}^{ij}}{O_t^{ij}} - r_t^f \right) \quad (6)$$

$$\begin{aligned} &= \frac{1}{S_{j,t}} \frac{\partial O^{ij}}{\partial \sigma_M^2} \left(d\sigma_{M,t}^2 - E_t^Q[d\sigma_{M,t}^2] \right) + \frac{\partial O^{ij}}{\partial S_j} \underbrace{\left(\frac{dS_{j,t}}{S_{j,t}} - r^f dt \right)}_{drs_{j,t}} + \underbrace{\mathbf{1}_{\{j=l\}} \frac{1}{S_{j,t}} \frac{\partial O^{ij}}{\partial \sigma_l^2} \left(d\sigma_{j,t}^2 - E_t^Q[d\sigma_{j,t}^2] \right)}_{d\varepsilon_{ij,t}}, \\ &= \gamma_{ij,t} dt + \frac{1}{S_{j,t}} \frac{\partial O^{ij}}{\partial \sigma_M^2} d\sigma_{M,t}^2 + \frac{\partial O^{ij}}{\partial S_j} drs_{j,t} + d\varepsilon_{ij,t}, \end{aligned} \quad (7)$$

where the (instantaneous) intercept $\gamma_{ij,t} dt$ is equal to $\frac{1}{S_{j,t}} \frac{\partial O^{ij}}{\partial \sigma_M^2} E_t^Q[d\sigma_{M,t}^2]$, $drs_{j,t}$ is the excess return on security j , and $d\varepsilon_{ij,t}$ denotes the residual. When idiosyncratic variance innovations are not priced, $E_t^P[d\varepsilon_{ij,t}] = 0$ as $E_t^P[d\sigma_{j,t}^2] = E_t^Q[d\sigma_{j,t}^2]$ where $E^P[\cdot]$ is the conditional expectation under the physical probability measure. Note that both stock and ETF option return dynamics are impacted by the underlying-specific idiosyncratic variance factor.

After delevering, the delta exposure becomes $\frac{\partial O^{ij}}{\partial S_j}$ (the standard option delta), while the variance exposures are scaled by the underlying price rather than the option price. Most affine option pricing models are homogeneous of degree one in the stock-strike pair.¹⁴ Specifically, $\frac{\partial O^{ij}}{\partial S_j}$ depends only on the moneyness, maturity, and variance states but not on the underlying price level. A similar argument applies to the vega term $\frac{1}{S_{j,t}} \frac{\partial O^{ij}}{\partial \sigma_M^2}$.

While both the theoretical and empirical frameworks can be adapted to alternative return definitions, using delevered returns addresses three key limitations of simple returns. First, it mitigates the option return measurement bias documented by Duarte, Jones, and Wang (2024). Second, delevering improves the stability and tractability of exposure estimation through regularization. Third, simple returns mechanically lever exposures by $1/O_t^{ij}$, which can overweight deep out-of-the-money options in estimation and in turn distort exposure and risk-premium estimates.

¹⁴Standard affine stochastic volatility models, including Heston (1993) and its extensions, satisfy this property: if both the stock price and strike are scaled by a constant, the option price scales by the same factor.

Extension with systematic jumps. The baseline model can be extended to incorporate systematic jump risk, where the market return follows

$$\frac{dS_{M,t}}{S_{M,t}} = \mu_t dt + \sigma_{M,t} dW_t^M + J (dN_t - \lambda^P dt), \quad (8)$$

with dN_t a Poisson process of intensity λ^P under P , constant jump size J , and $dN_t - \lambda^P dt$ the corresponding compensated martingale. Each security l inherits jump exposure through its market beta β_l . When jump risk is priced, $\lambda^Q \neq \lambda^P$. When a jump occurs, the underlying moves from $S_{j,t}$ to $S_{j,t}(1 + \beta_j J)$ with $\beta_j = 1$ for $j = M$, and the resulting contribution to the delevered option excess return in equation (7) is

$$\frac{1}{S_{j,t}} [O^{ij}(S_{j,t}(1 + \beta_j J)) - O^{ij}(S_{j,t})] (dN_t - \lambda^Q dt). \quad (9)$$

The loading on the Q -compensated innovation $dN_t - \lambda^Q dt$ is the option jump revaluation scaled by $1/S_{j,t}$. The jump risk premium follows from taking the P -expectation: $E_t^P[dN_t - \lambda^Q dt] = (\lambda^P - \lambda^Q) dt$, which captures the wedge between physical and risk-neutral arrival rates.

3.2. A Conditional Factor Model of Delevered Excess Equity Option Returns

The representations above are stated in continuous time, while the empirical analysis uses weekly returns. The discrete delevered excess return on O^{ij} – the value of equity option i written on the S&P 500 index, the stock price of a firm, or the price of an ETF, with a given exercise price K_{ij} , time to maturity τ_{ij} , and payoff function (i.e., call or put) – is given by:

$$r_{ij,t+1} \equiv \frac{O_t^{ij}}{S_{j,t}} \cdot \left(\frac{\tilde{\Delta} O_{t+1}^{ij}}{O_t^{ij}} - r_t^f \right), \quad (10)$$

where $\tilde{\Delta}$ is the discrete increment operator over the time interval $[t, t+1]$ such that $\tilde{\Delta} O_{t+1}^{ij} = O_{t+1}^{ij} - O_t^{ij}$, and r_t^f is the risk-free rate of return per period. Henceforth, we refer to returns calculated in this manner simply as “option returns”.

A second-order Taylor expansion of the option price combined with the no-arbitrage condition (derived in Appendix A.5) yields the following factor structure of discrete option

returns:

$$r_{ij,t+1} = \gamma_{ij,t} + \underbrace{\mathcal{V}_{ij,t} \cdot f_{v,t+1}}_{\tilde{\Delta}\sigma_{M,t+1}^2} + \underbrace{\Delta_{ij,t} \cdot rs_{j,t+1}}_{\frac{\tilde{\Delta}S_{j,t+1}}{S_{j,t}} - r_t^f} + \Gamma_{ij,t} \cdot rs_{j,t+1}^2 + \varepsilon_{ij,t+1} \quad (11)$$

where $\gamma_{ij,t}$ is the conditional intercept, $f_{v,t+1}$ denotes the realization of the systematic variance factor between t and $t+1$, $rs_{j,t+1}$ is the underlying security j 's excess return over the same interval, and $rs_{j,t+1}^2$ is the squared excess returns of the underlying security j . $f_{v,t+1}$ is not Q -demeaned, so its conditional Q -expectation is absorbed into $\gamma_{ij,t}$ as we discuss below. Note that both $rs_{j,t+1}$ and $rs_{j,t+1}^2$ are specific to each underlying security, while the systematic variance factor is common to all options. $\varepsilon_{ij,t+1}$ is the model residual.¹⁵ Moreover, $\mathcal{V}_{ij,t}$, $\Delta_{ij,t}$, and $\Gamma_{ij,t}$ denote the option's systematic vega, its delta, and its gamma, respectively. Comparing the continuous with the discrete factor model, three important remarks are in order.

First, the squared excess return $rs_{j,t+1}^2$ in the discrete model plays a dual role. On the one hand, it captures the option's convexity (gamma) exposure to large price movements over the discrete interval. In continuous time this term is absorbed by the pricing equation because quadratic variation is locally deterministic; in discrete time it becomes genuinely stochastic. When jumps are present, a second-order expansion of the exact jump revaluation (9) shows that the gamma representation of systematic jump price impact is $\Gamma_{ij,t} \cdot (\beta_j J)^2$, which maps directly into large realizations of $rs_{j,t+1}^2$ (see Appendix A.5 for details). On the other hand, since $rs_{j,t+1}^2$ reflects the total realized variance of security j —both systematic and idiosyncratic—it controls for the potentiality of underlying-specific variance innovations being priced. In such case, omitting $rs_{j,t+1}^2$ would contaminate inference on the systematic variance risk premium $f_{v,t+1}$. Incorporating the squared excess returns of the underlying security j when estimating discrete exposures enhances the precision of risk premium estimates of systematic factors by reducing the incidence of omitted variable bias.

Second, the residual $\varepsilon_{ij,t+1}$ in equation (11) captures deviations from the factor structure and our baseline model assumes $E_t^P[\varepsilon_{ij,t+1}] = 0$. Relative to the theory, $\varepsilon_{ij,t+1}$ can be interpreted as the idiosyncratic variance innovation, together with higher-order terms (see

¹⁵We do not explicitly model heteroskedasticity and the correlation of residuals across time, options, and stocks, but we account for these patterns when conducting inference on the parameter estimates.

equation (IA.36)):

$$\varepsilon_{ij,t+1} \stackrel{\text{theor.}}{\underset{\text{equiv.}}{\approx}} \frac{1}{S_{j,t}} \left[\frac{\tilde{\Delta} O^{ij}}{\tilde{\Delta} \sigma_l^2} (\tilde{\Delta} \sigma_{j,t+1}^2 - E_t^Q[\tilde{\Delta} \sigma_{j,t+1}^2]) + (\text{H.O.T.} - E_t^Q[\text{H.O.T.}]) \right]. \quad (12)$$

When idiosyncratic variance innovations are not priced and higher order terms beyond gamma exposure are quantitatively small, $\varepsilon_{ij,t+1}$ would satisfy $E_t^P[\varepsilon_{ij,t+1}] = 0$.¹⁶ This condition holds because, when idiosyncratic variance is not priced, $E_t^P[\tilde{\Delta} \sigma_{j,t+1}^2] = E_t^Q[\tilde{\Delta} \sigma_{j,t+1}^2]$, so the idiosyncratic innovation inside $\varepsilon_{ij,t+1}$ has zero conditional mean under P as well.

When $E_t^P[\varepsilon_{ij,t+1}] = 0$, the no-arbitrage intercept is then

$$\gamma_{ij,t} = -\mathcal{V}_{ij,t} \cdot E_t^Q[f_{v,t+1}] - \Gamma_{ij,t} \cdot E_t^Q[rs_{j,t+1}^2]. \quad (13)$$

Thus, $\gamma_{ij,t}$ depends solely on the product of exposures with the conditional risk-neutral expectation of the systematic variance factor and the squared excess return, both of which are treated as nontradable. It is straightforward to extend the model to allow $\gamma_{ij,t}$ to be augmented by an option-level pricing error $\alpha_{ij,t}$.

Third, while the economic interpretation of the conditional exposures in the discrete model is similar to their continuous counterpart,

$$\mathcal{V}_{ij,t} \stackrel{\text{theor.}}{\underset{\text{equiv.}}{\approx}} \frac{1}{S_{j,t}} \frac{\tilde{\Delta} O^{ij}}{\tilde{\Delta} \sigma_M^2}, \quad \Delta_{ij,t} \stackrel{\text{theor.}}{\underset{\text{equiv.}}{\approx}} \frac{\tilde{\Delta} O^{ij}}{\tilde{\Delta} S_j}, \quad \Gamma_{ij,t} \stackrel{\text{theor.}}{\underset{\text{equiv.}}{\approx}} \frac{1}{2} S_{j,t} \frac{\tilde{\Delta}^2 O^{ij}}{\tilde{\Delta} S_j^2}, \quad (14)$$

they differ in an important way. In continuous time, exposures are partial derivatives that do not account for the potential dependence and covariance structure of the factors. In the discrete time model, the conditional exposures are estimated jointly using realized option returns and factor realizations, and as a result account for any comovement between factors.¹⁷

Our main focus is on studying systematic variance risk premiums in delta-hedged delevered excess equity option returns: $r_{ij,t+1}^\Delta \equiv r_{ij,t+1} - \Delta_{ij,t} rs_{j,t+1}$. The two-stage estimation procedure

¹⁶If a systematic component were driving the pricing of idiosyncratic variance innovations or H.O.T., an additional priced factor would be required. Such a factor would, however, be partially spanned by $rs_{j,t+1}^2$, since total realized variance reflects both systematic and idiosyncratic components.

¹⁷While $\Delta_{ij,t}$ is directly equivalent to the Black-Scholes delta ($\Delta_{ij,t}^{BS}$) and $\Gamma_{ij,t}$ to $\frac{1}{2} S_{j,t} \Gamma_{ij,t}^{BS}$, $\mathcal{V}_{ij,t}$ corresponds to $\mathcal{V}_{ij,t}^{BS} \left(\partial \sigma_M / \partial \sqrt{\beta_j^2 \sigma_M^2 + \sigma_j^2} \right)$. The vega mapping arises because Black-Scholes vega captures exposure to the total volatility of the underlying security, whereas our approach isolates exposure to systematic variance.

exploits the factor structure in equation (11) to identify these premiums from the cross-section of option returns.

3.2.1. Least-squares estimation of option exposures

Our baseline specification is supported by an analysis of biases present in least-squares estimation of option sensitivities. Biases arise through the omitted variables mechanism. If the regression model ignores terms implied by the theoretical specification that are correlated with the included terms, they are present even in the continuous-time observation limit. Estimation on discretized data introduces additional sources of bias which scale with the length of the discretization interval Δt . We rewrite the delevered return equation (7), this time explicitly including the term arising from idiosyncratic volatility:

$$\frac{dO_{j,t}}{S_{j,t}} - \frac{S_{j,t}}{O_{j,t}} r_t^f = \gamma_{ij,t} dt + \frac{1}{S_{j,t}} \frac{\partial O^{ij}}{\partial \sigma_M^2} d\sigma_{M,t}^2 + \frac{\partial O^{ij}}{\partial S_j} drs_{j,t} + \frac{1}{S_{j,t}} \frac{\partial O^{ij}}{\partial \sigma_j^2} d\sigma_{j,t}^2 \quad (15)$$

and consider OLS estimation of the conditional intercept and sensitivities. To simplify the exposition, we examine the regression for a single option and keep the state variables $\sigma_{M,t}$ and $\sigma_{j,t}$ constant, and drive the model one-period forward via the realizations of the four Brownian shocks. In this setting, the conditional intercept and option sensitivities are constant. We analyze the following regression specification:

$$\frac{dO_{j,t}}{S_{j,t}} - \frac{S_{j,t}}{O_{j,t}} r_t^f = \gamma + \beta_{\sigma_M} \Delta \sigma_{M,t+1}^2 + \beta_S \Delta rs_{t+1} + \beta_\Gamma (\Delta rs_{t+1})^2 + u_{t+1}. \quad (16)$$

This specification explicitly ignores the idiosyncratic variance term $\Delta \sigma_{j,t+1}^2$ but includes the “gamma” term $(\Delta rs_{t+1})^2$ which is absent from the Itô expansion but arises in discrete time. Further such terms are absorbed, together with the idiosyncratic variance term, into u_{t+1} :

$$u_{t+1} = \beta_{\sigma_j} \Delta \sigma_{j,t+1}^2 + \beta_{S\sigma_M} \Delta rs_{t+1} \Delta \sigma_{M,t+1}^2 + \beta_{\sigma_M \sigma_S} \Delta \sigma_{M,t+1}^2 \Delta \sigma_{S,t+1}^2 + \tilde{u}_{t+1}. \quad (17)$$

Analysis of the bias. The bias on the estimated coefficients is determined by the covariance between the omitted terms in (17) and the included regressors. In the remainder of this section we characterize the bias in terms of the model’s parameters. To do so, denote by $\epsilon_1, \dots, \epsilon_4$ the standardized Brownian increments driving W^M, Z^M, W^S, Z^S , with $\text{Corr}(\epsilon_1, \epsilon_2) = \rho$,

$\text{Corr}(\epsilon_2, \epsilon_4) = \rho_v$, and all other pairs independent. Define

$$\sigma_{j,\text{Tot.}}^2 \equiv \beta^2 \sigma_M^2 + \sigma_j^2, \quad \sigma_{j,\perp}^2 \equiv \sigma_j^2 + \beta^2 \sigma_M^2 (1 - \rho^2),$$

the total instantaneous stock variance and its component orthogonal to market variance innovations, respectively.

The three omitted terms in (17) produce biases at two distinct orders. The key observation is that the discretized regressors $\Delta\sigma_{M,t+1}^2$ and Δrs_{t+1} are linear in the Brownian increments at leading order, as is the omitted $\Delta\sigma_{j,t+1}^2 \propto \epsilon_4 \sqrt{\Delta t}$. In contrast, the cross-variation $\Delta rs_{t+1} \Delta\sigma_{M,t+1}^2$ and the cross-volga $\Delta\sigma_{M,t+1}^2 \Delta\sigma_{j,t+1}^2$ are quadratic in the shocks. Key for the analysis is the observation that, for jointly Gaussian variables, centered products are orthogonal to linear terms:¹⁸

$$E[\epsilon_k \cdot (\epsilon_l \epsilon_m - E[\epsilon_l \epsilon_m])] = 0 \quad \text{for all } k, l, m. \quad (18)$$

This result zeroes out the leading-order covariance between the omitted terms $\Delta rs_{t+1} \Delta\sigma_{M,t+1}^2$ and $\Delta\sigma_{M,t+1}^2 \Delta\sigma_{S,t+1}^2$, and the linear regressors Δrs_{t+1} and $\Delta\sigma_{M,t+1}^2$. Correlation arises only through the interaction of drift and diffusion in the discretization, producing an $O(\Delta t)$ bias that vanishes in the continuous-time limit. No such orthogonality shields the $\Delta\sigma_j^2$ channel: since $\Delta\sigma_{j,t+1}^2 \propto \epsilon_4$ and $\Delta\sigma_{M,t+1}^2 \propto \epsilon_2$ are both linear, then if $\rho_v \neq 0$, their covariance is the same order as the regressor variances, yielding $O(1)$ bias.

An exception arises for $\hat{\beta}_\Gamma$. The regressor $(\Delta rs_{t+1})^2$ is itself quadratic in the shocks, so (18) does not apply between it and the omitted cross-variation. Their covariance is $O((\Delta t)^2)$ —the same order as $\text{Var}((\Delta rs_{t+1})^2)$ —producing an additional $O(1)$ contamination.

¹⁸This result generalized to elliptically symmetric distributions.

Working through the omitted variable bias formula, the biases are:

$$\text{bias}(\hat{\beta}_{\sigma_M}) = \underbrace{\frac{\sigma_{j,\text{Tot.}}^2 \sigma_j \delta_j \rho_v}{\sigma_M \delta_M \sigma_{j,\perp}^2} \beta_{\sigma_j}}_{\text{from } \Delta\sigma_j^2, O(1)} + \underbrace{\mu_j \beta_{S\sigma_M} \Delta t}_{\text{from cross-var, } O(\Delta t)} \quad (19)$$

$$\text{bias}(\hat{\beta}_S) = \underbrace{\frac{-\beta \sigma_j \delta_j \rho \rho_v}{\sigma_{j,\perp}^2} \beta_{\sigma_j}}_{\text{from } \Delta\sigma_j^2, O(1)} + \underbrace{\left[\kappa_M - \frac{2\mu_j \beta \sigma_M^2 \delta_M \rho}{\sigma_{j,\text{Tot.}}^2} \right] \beta_{S\sigma_M} \Delta t}_{\text{from cross-var, } O(\Delta t)} \quad (20)$$

$$\text{bias}(\hat{\beta}_\Gamma) = \underbrace{\frac{\beta \sigma_M^2 \delta_M \rho}{\sigma_{j,\text{Tot.}}^2} \beta_{S\sigma_M}}_{\text{from cross-var, } O(1)} \quad (21)$$

where $\mu_j \equiv (r^f - q_j) + \beta(\mu_M - r^f)$ is the stock's conditional drift, $\kappa_M \equiv K_m(\theta_M - \sigma_M^2)$ the mean-reversion rate of market variance, and δ_M, δ_j the vol-of-vol parameters in the variance dynamics.

The $O(1)$ bias on $\hat{\beta}_{\sigma_M}$ reflects direct confounding: when $\rho_v > 0$, positive shocks to market volatility coincide with positive shocks to idiosyncratic volatility, and the regression misattributes the option's sensitivity to $\Delta\sigma_j^2$ to $\Delta\sigma_M^2$. The bias scales with ρ_v ; the factor $\sigma_{j,\text{Tot.}}^2/\sigma_{j,\perp}^2$ reflects the partial-regression adjustment after controlling for Δrs_{t+1} . The $O(1)$ bias on $\hat{\beta}_S$ is a suppressor effect. Although $\text{Cov}(\Delta rs_{t+1}, \Delta\sigma_{j,t+1}^2) = 0$, the Frisch–Waugh residual of Δrs_{t+1} on $\Delta\sigma_{M,t+1}^2$ acquires a spurious correlation with $\Delta\sigma_{j,t+1}^2$ running through the leverage channel (ρ). Both biases scale linearly in ρ_v .

The $O(1)$ bias on $\hat{\beta}_\Gamma$ in (21) arises because the squared return and the cross-variation are both quadratic in the shocks. When $\rho < 0$ (leverage), large squared returns coincide with negative cross-variation, and the regression misattributes part of the cross-variation effect to gamma. The magnitude is governed by $\beta^2 \sigma_M^2 / b$, the market-driven fraction of total stock variance, which is typically modest.

The $O(\Delta t)$ terms in (19)–(20) arise because the cross-variation correlates with $\Delta\sigma_{M,t+1}^2$ only through the stock's drift (μ_j), and with Δrs_{t+1} only through the mean-reversion drift of market variance (κ_M). The cross-volga contributes only at $O(\Delta t)$ as well—it is deterministic in continuous time—and is omitted from (19)–(21). At daily observation frequency, both

$O(\Delta t)$ channels are negligible.

Including $(\Delta rs_{t+1})^2$ as a regressor eliminates an additional $O(\Delta t)$ bias of $2\mu_j \beta_\Gamma \Delta t$ on $\hat{\beta}_S$ that would appear if the regression omitted the gamma term, at the cost of introducing the $O(1)$ contamination (21).

Quantification of the bias. To gauge magnitudes, we evaluate (19)–(21) at representative parameter values: $\sigma_M = 0.16$, $\delta_M = 0.50$, $\rho = -0.70$, $\beta = 1$, $\sigma_j = 0.25$, $\delta_j = 0.40$, and $\Delta t = 5/252$ (weekly), for an at-the-money one-month call (we obtain approximate $\beta_{\sigma_M} = \beta_{\sigma_j} \approx 0.19$, $\beta_S \approx 0.53$, $\beta_\Gamma \approx 2.3$, $\beta_{S\sigma_M} \approx 0.03$ from Black-Scholes greeks). Based on our stock option dataset, we calibrate $\rho_v = 0.04$. At this value the $O(1)$ channel inflates $\hat{\beta}_{\sigma_M}$ by approximately 6% and $\hat{\beta}_S$ by approximately 1.4%. The contamination of $\hat{\beta}_\Gamma$ is below 0.2%. All $O(\Delta t)$ terms are of order 10^{-5} —negligible at weekly frequency.

3.3. Modeling Conditional Exposures

The specifications in equation (11) above is fairly general and require additional structure on the conditional exposures and the risk premiums. Building on the option pricing literature, we closely follow Fournier, Jacobs, and Orłowski (2024) and model exposures nonparametrically using basis expansion theory. Consider, for instance, the modeling of the systematic vega exposure of option i written on security j to the aggregate variance risk in equation (11). We assume that $\mathcal{V}_{ij,t}$ satisfies $\mathcal{V}_{ij,t} = \mathcal{V}^j(\mathbf{c}_{ij,t})$ where $\mathbf{c}_{ij,t}$ is a vector composed of option $\{ij\}$ ’s maturity, moneyness, and a measure of implied volatility.¹⁹ $\mathcal{V}^j(\cdot)$ represents a nonlinear transformation of $\mathbf{c}_{ij,t}$ capturing the exposures to systematic variance risk applied to all options on underlying security j .²⁰ To provide more intuition, note that the function $\mathcal{V}^j(\cdot)$, which maps the characteristics of any option written on security j to its exposure to systematic variance, is related to the vega $\mathcal{V}(\cdot)$ function in the Black-Scholes model. However, while the Black-Scholes vega captures an option’s exposure to the total volatility of its underlying asset, $\mathcal{V}^j(\cdot)$ isolates the exposure to systematic variance risk. In

¹⁹Unlike SPX options, ETF and stock options exhibit greater contract-level idiosyncratic variation not fully captured by security-level implied volatility. For these markets, using option-level implied volatilities to instrument exposures improves the model’s ability to account for variation across contracts.

²⁰Note that $\mathbf{c}_{ij,t}$ can also include factor-specific variables but we omit this dependence to simplify the notation.

addition, the functional form of the Black-Scholes vega does not depend on the underlying asset, while ours is security specific.

To account for nonlinearities in exposures, we employ Thin Plate Regression Spline (TPRS) basis functions (see Wood, 2003) to map the vector of characteristics $\mathbf{c}_{ij,t}$ into a higher-dimensional basis vector, $\boldsymbol{\phi}_{ij,t}$, consisting of nonlinear transformations and interactions of the characteristics. Thus, $\boldsymbol{\phi}_{ij,t}$ is fixed given our choice of basis expansion and option characteristics, and can be treated as a vector of explanatory variables. We then assume that option exposures are linear functions of the basis vector so that, $\mathcal{V}_{ij,t} = \mathcal{V}^j(\mathbf{c}_{ij,t}) = \boldsymbol{\phi}_{ij,t}' \mathbf{b}_{v,j}$. The modeling of $\Delta_{ij,t}$ and $\Gamma_{ij,t}$ in equation (11) follows the same approach.

3.4. Modeling Conditional Risk Premiums

We are now left with the specification of conditional expectations and risk premiums. Parameterizing the risk-neutral conditional expectations of $f_{v,t+1}$ and $rs_{j,t+1}^2$ is required to estimate $\gamma_{ij,t}$ in equation (12), while specifying their risk premiums is essential for conducting inference on risk compensation. Although the model can accommodate nonlinear specifications, our empirical analysis focuses on linear functions. For squared return risk, we adopt a parsimonious constant expectation/premium specification common to all underlying securities, such that $E_t^Q[rs_{j,t+1}^2] = \mu_\Gamma^Q$ and $(E_t^P - E_t^Q)[rs_{j,t+1}^2] = \lambda_\Gamma$.²¹ We model the conditional expectations of f_v under the P and Q measures as an affine function of selected state variables, which implies an affine specification of its risk premium. Our baseline specification assumes that the conditional expectation of the systematic variance risk under the risk-neutral measure is given by $E_t^Q[f_{v,t+1}] = \mu_v^Q \cdot VIX_t^2$ where the Q superscript denotes the risk-neutral parameters, and the squared VIX is used as the state variable.²² Consequently, the factor risk premium satisfies $(E_t^P - E_t^Q)[f_{v,t+1}] = \lambda_v \cdot VIX_t^2$, where λ_v is the risk premium parameter.

To support this model feature, we now demonstrate that the expected return on a portfolio formed with SPX options that tracks the $f_{v,t+1}$ factor, which we proxy using weekly changes

²¹This assumption does not affect the empirical results, because $rs_{j,t+1}^2$ does not have a significant economic impact on expected option returns (see also Fournier, Jacobs, and Orłowski, 2024). It is important to include it in the model because it has significant explanatory power for realized option returns.

²²The literature contains alternative predictors of risk premiums. Bollerslev, Tauchen, and Zhou (2009) and Bollerslev, Todorov, and Xu (2015) find that the market variance risk premium predicts market equity returns. Bollerslev and Todorov (2011) and Bollerslev, Todorov, and Xu (2015) emphasize that compensation for rare events constitutes a substantial part of the market equity risk premium.

in the squared VIX index, is strongly negatively related to the level of systematic variance, proxied by VIX_t^2 . Each week we use the 50 options with the highest open interest to construct the portfolio. We calculate portfolio weights using the estimated exposures to the S&P 500 index return, its square, and the $f_{v,t+1}$ factor.²³

The excess return of this mimicking portfolio averages -14% per year. The scatterplot in Figure 1 documents the relation between the excess returns on this portfolio and VIX_t^2 . The green line is the regression line of the portfolio return on VIX_t^2 ; its coefficients are reported in column (1) of Table A.1 in the Appendix. The slope coefficient of -0.0774 is statistically significant at the 10% level and close to our main estimates in Table 5. The orange line is obtained from a nonparametric regression of portfolio returns on VIX_t^2 . While it fluctuates due to local return variation, it consistently maintains a negative slope. Both lines thus indicate that the return on the VAR-mimicking portfolio is strongly negatively related to the level of VIX_t^2 . Regression estimates in column (2) of Table A.1, obtained from a sample which discards VIX_t^2 values above their 90th percentile (i.e., all points to the right of the rightmost vertical grey line in Figure 1), demonstrate that the result is not driven by extreme outcomes.

Finally, note that the nonparametric estimate of the expected return (represented by the orange line) is consistently negative and very close to zero when VIX_t^2 is low. This evidence supports our factor risk premium specification: a risk premium proportional to VIX_t^2 (no intercept) provides a good approximation of the nonparametric estimate over most of the VIX_t^2 distribution. In Section 6, we show the importance of modeling this time variation. If we instead assume a constant expected variance risk premium specification, the resulting estimates of the variance risk premium are biased and often switch sign. There we also discuss the benefits of constraining the intercept to zero in the factor risk premium specification. In the absence of this assumption, the estimated average risk premiums remain negative, but they vary more across markets and option types.

²³Specifically, at each time t we set $\hat{\mathbf{B}}_t = [\hat{\Delta}_t \hat{\Gamma}_t \hat{\mathbf{V}}_t]$ (of size 50×3) and find the weights $\mathbf{w}_{\text{VAR},t}$ that satisfy $\hat{\mathbf{B}}_t' \mathbf{w}_{\text{VAR},t} = [0 \ 0 \ 1]'$. Such $\mathbf{w}_{\text{VAR},t}$ is not unique, hence we pick the one with the lowest sum of squared weights, using the Moore-Penrose generalized inverse to solve this equation.

3.5. Estimation

Together, our choice of basis expansions to model option exposures and of the functional form of conditional expectations/risk premiums imply that the conditional factor model in (11) can be written as a linear (unconditional) asset pricing model with additional explanatory variables (see Nagel and Singleton, 2011; Gagliardini, Ossola, and Scaillet, 2016). This, in turn, allows for straightforward estimation of the model parameters using two-stage OLS.

For instance, consider $\mathcal{V}_{ij,t} \cdot f_{v,t+1}$ in equation (11). Given the definition of exposures, $\mathcal{V}_{ij,t} \cdot f_{v,t+1}$ is equivalent to $(\phi'_{ij,t} \mathbf{b}_{v,j}) \cdot f_{v,t+1} = (f_{v,t+1} \phi'_{ij,t}) \cdot \mathbf{b}_{v,j}$ where $f_{v,t+1} \phi'_{ij,t}$ can be treated as a vector of explanatory variables and $\mathbf{b}_{v,j}$ is the vector of parameters to be estimated using full sample OLS. Similarly, $\gamma_{ij,t}$ in equation (12) can also be expressed as the product of the explanatory variables and appropriately redefined model parameters.

In the first step, we estimate the exposures $\mathcal{V}_{ij,t}$, $\Delta_{ij,t}$, and $\Gamma_{ij,t}$, as well as the conditional intercept $\gamma_{ij,t}$ separately for each underlying security j using weekly panels of delevered return on individual options. With the appropriate mapping, this step is analogous to the standard first-stage regression in unconditional factor models with constant exposures and intercept.

After delta-hedging using first-stage estimates of $\Delta_{ij,t}$, the baseline model reduces to nontraded systematic variance risk factor, $f_{v,t+1}$, and $rs_{j,t+1}^2$. To estimate the associated risk premiums, we first project delevered delta-hedged option returns, $r_{ij,t+1} - \hat{\Delta}_{ij,t} rs_{j,t+1}$, onto the model intercept and the interactions of the remaining factors with their respective exposures. This projection amounts to $\hat{\gamma}_{ij,t} + \hat{\mathcal{V}}_{ij,t} \cdot f_{v,t+1} + \hat{\Gamma}_{ij,t} \cdot rs_{j,t+1}^2$. We then conduct a second-stage regression of the projected delta-hedged returns on $\hat{\mathcal{V}}_{ij,t} \cdot VIX_t^2$ and $\hat{\Gamma}_{ij,t}$ to estimate the risk premium parameters. This amounts to regressing the projected component of returns on the time-varying product of exposures and risk premiums, similarly to Fama-MacBeth regressions, but on the whole sample at once. This procedure is the empirical equivalent of taking the (physical) conditional expectations of the projection of delta-hedged returns after replacing $\gamma_{ij,t}$ with the expression in equation (12), which yields:

$$\begin{aligned} \mathcal{V}_{ij,t} \left(E_t^P - E_t^Q \right) [f_{v,t+1}] + \Gamma_{ij,t} \left(E_t^P - E_t^Q \right) [rs_{j,t+1}^2] \\ = (\mathcal{V}_{ij,t} VIX_t^2) \cdot \lambda_v + \Gamma_{ij,t} \cdot \lambda_\Gamma \end{aligned} \quad (22)$$

Our inference scheme accounts for measurement errors in the exposures from the first stage,

and their correlation with the measurement errors in the projection.

For a more detailed discussion of our modeling approach, estimation strategy, and asymptotic inference in a general framework with multiple traded and nontraded systematic factors, refer to Internet Appendix [B](#).

4. Data

We describe the samples for the three option markets used in the empirical analysis: index options, stock options, and ETF options. We provide descriptive statistics of the option characteristics and delta-hedged delevered option returns for these three samples.

4.1. Sample Composition

Our empirical analysis reports on three equity option markets. To investigate the existence and magnitude of a priced systematic variance risk premium, we consider options written on (i) the S&P 500 index (SPX); (ii) common stocks; and (iii) ETFs. The stock option sample consists of options written on stocks that were S&P 500 index constituents at any time from 1996-01 to 2023-02 and that had stock options listed on the Cboe during this period. The ETF sample includes all funds with listed equity options on the Cboe that do not track the main stock indices.²⁴ For all options with quotes on the business day closest to Wednesday, we calculate weekly delevered excess equity option returns for OptionMetrics data from 1996-01 to 2023-02. We use weekly delevered excess equity option returns as defined in equation (10), computed using end-of-day mid-quotes.

We retain options with maturities ranging from one to nine months at the beginning of each week. We discard option returns associated with stocks that have an ex-dividend date or earnings announcement in a given week. It is important for our estimation strategy to ensure the availability of a sufficient number of observations for each underlying security. We therefore restrict the sample to stocks and ETFs with at least 8000 weekly option return observations. For all underlying securities, we also discard weekly option returns if either of the options used to compute the return has a zero bid price, or if at the start of the period

²⁴Because the sample of ETF options is the least liquid of the three, we also include ETFs composed of securities that are not in the S&P 500 in this sample.

the ask price is five times larger than the bid price. For index and stock options, we discard returns if at the start of the return measurement period, the option has zero open interest or volume. We do not apply these filters to ETF options because the much lower trading activity would significantly lower the sample size.

Following Büchner and Kelly (2022) and Fournier, Jacobs, and Orłowski (2024), we do not rely on in-the-money (ITM) options for illiquidity reasons and we limit the sample to at-the-money (ATM) and out-of-the-money (OTM) options. Accordingly, we restrict the dataset to put options with moneyness ($K_{ij}/S_{j,t}$) below 1.025 and call options with moneyness above 0.975. We merge the equity option return data with the weekly returns for the corresponding underlying security, computed using the CRSP Daily File. We use weekly underlying excess returns and their squared values as proxies for $rs_{j,t+1}$ and $rs_{j,t+1}^2$, to estimate deltas and gammas, and to compute delta-hedged delevered equity option returns. Finally, we obtain VIX data from the Cboe website. As mentioned above, we proxy systematic variance innovations ($f_{v,t+1}$) using weekly changes in the squared VIX index.²⁵ For notational convenience, henceforth we refer to this proxy for the $f_{v,t+1}$ factor as VAR. In total, our sample includes more than 12 million observations on weekly equity option returns on the SPX index, 422 stocks, and 96 ETFs.

4.2. Descriptive Statistics: Option Characteristics

Throughout the empirical analysis, we report results for all options and calls and puts separately. On occasion we discuss maturity and moneyness effects. To provide insight into term structure effects, we distinguish between medium (1–2 months) and long-term (2–9 months) maturities. Because of the large variation in volatility and price levels across underlying securities, we categorize at-the-money (ATM) and out-of-the-money (OTM) options based on their Black-Scholes deltas, $\Delta_{ij,t}^{BS}$. To this end, we first compute the average absolute Black-Scholes delta based on the entire sample for all options written on a given underlying category (index, stocks, or ETFs), and of a given type, call or put, with $K_{ij}/S_{j,t} \in [0.975; 1.025]$. We define calls and puts to be ATM if their Black-Scholes deltas lie within ± 0.1 from this respective average. Any option that does not meet this criterion is categorized as OTM.

²⁵Note that we treat the VIX index in decimal form: we use 0.20 for a reported VIX value of 20, and an increase in squared VIX from 0.20^2 to 0.25^2 corresponds to a change of 0.0225.

Table 1 presents descriptive statistics of option characteristics for puts and calls written on the SPX (Panel A), stocks (Panel B), and ETFs (Panel C). The first column of Table 1 indicates that on average, SPX options have the shortest maturity, followed by stock options. ETF options have on average the longest maturity. Within the three asset classes, the average maturity of puts is similar to that of calls. Furthermore, for each asset class the average moneyness, measured using $K_{ij}/S_{j,t}$, deviates almost symmetrically from 1 for both puts and calls. The sample covers a large cross-section of option contracts, as evidenced by the 5th percentile of the put moneyness distribution and the 95th percentile of the call moneyness distribution. For the SPX options in Panel A, these percentiles correspond to 0.78 for puts and 1.17 for calls. For the stock options in Panel B, they are 0.66 and 1.41, and for the ETF options in Panel C they are 0.71 and 1.26.

Stock options have the highest average implied volatilities, followed by ETF options, and then SPX options. Put implied volatilities on average exceed call implied volatilities, indicating the presence of an implied volatility skew (OTM calls can be viewed as ITM puts based on put-call parity). Consistent with the existing literature, the wedge between put and call implied volatilities is most pronounced for SPX options. Comparing the average implied volatility of a given option type and underlying security across maturity buckets, we see that puts and calls have relatively flat term structures of implied volatilities. These empirical stylized facts about the implied volatility skew and term-structure are consistent with previous findings, see e.g., Christoffersen, Fournier, and Jacobs (2018).

The last column of Table 1 indicates that the sample is somewhat unbalanced, in the sense that the total number of observations for puts and calls differs within all three asset classes. The SPX and ETF sample contains more puts, whereas the reverse is true for stock options. Finally, note that compared to the existing literature, we study a very extensive cross-section of equity option contracts, consisting of more than 12 million return observations in total.

4.3. Descriptive Statistics: Option Returns

Table 2 reports various summary statistics for delta-hedged delevered option returns. We delta-hedge one long position in each delevered option using the estimated sensitivity to the underlying security return, on a weekly basis. We then calculate summary statistics for each underlying security and report the cross-sectional averages of those statistics, except

of course for the case of SPX options, where there is a single underlying. The average return and standard deviation are annualized assuming 52 trading weeks in a year. We focus on delta-hedged delevered option returns for two reasons. First, delevering is helpful to address biases when estimating using option returns. Duarte, Jones, and Wang (2024) highlight the importance of this problem. Second, delta hedging is a very standard procedure when studying option returns, because it removes the impact of the underlying security, and therefore accounts for a very large number of cross-sectional determinants of the underlying returns, the so-called factor zoo.

Table 2 highlights several interesting stylized facts. Panel A indicates that the average weekly returns of SPX calls and puts are negative, of comparable magnitude, and statistically significant from zero across the board. For ETF options (Panel C), average returns for both puts and calls are negative, but compared to index option returns we observe more pronounced differences across maturities and option types. The negative average ETF option returns are statistically significant only 21% of the time for puts (33% for calls). For stock options (Panel B), returns are statistically significant for only 16% of stocks in the case of puts (19% in the case of calls). The average returns on stock options with maturities exceeding two months are positive in most cases but are not statistically significantly different from zero. Returns on short-maturity calls and puts are negative (more so for calls), but smaller in absolute value compared to index or ETF options. For all markets and option moneyness, both calls and puts have an upward sloping term structure of average returns, which is more pronounced for calls than for puts.

Overall, the most important finding is that negative delta-hedged returns are less pronounced in stock and ETF options compared to index options. This suggests that systematic variance risk, which is most likely the primary driver of negative returns in index options, is relatively less important for stock and ETF options, and that other factors might be significant drivers of the returns on those options.

5. Systematic Variance Risk in Equity Option Markets

This section documents the pricing of systematic variance risk in three equity option markets. We first assess the explanatory power of systematic variance risk for SPX, stock, and ETF option returns, and we discuss the patterns in the systematic variance exposures.

We provide estimates of the systematic variance risk premium in the three markets.

5.1. The Explanatory Power of Systematic Variance Risk

We use the total (pooled) R_p^2 as metric of explanatory power for option returns. Recall that $r_{ij,t}^\Delta$ denotes the delta-hedged delevered excess return on option i written on underlying security j observed in week t . The total R_p^2 is then defined as

$$R_p^2 \equiv 1 - \frac{\sum_{i,j,t} (r_{ij,t}^\Delta - \hat{r}_{ij,t}^\Delta)^2}{\sum_{i,j,t} (r_{ij,t}^\Delta)^2}. \quad (23)$$

where $\hat{r}_{ij,t}^\Delta$ denotes the model value of the delta-hedged delevered excess return. Table 3 provides an overview of the model's explanatory power, along with a decomposition of this explanatory power into the contribution of VAR and $rs_{j,t+1}^2$. The three panels report results for the three option markets. Each panel also reports results for different moneyness and maturity buckets.

The column labeled *Tot.* reports the total R_p^2 . We obtain the highest explanatory power (80.0%) for longer-maturity OTM SPX puts, and the lowest (30.2%) for short-maturity OTM ETF puts. The total R_p^2 s for SPX calls and stock calls are similar in magnitude (60.4% and 60.7%), but the total R_p^2 s for SPX puts (72.2%) far exceeds that of stock puts (53.0%). For ETF options, the model's explanatory power is lower than for SPX and stock options, with a more balanced fit between puts and calls.

The R_p^2 decomposition indicates that the relative importance of the VAR factor in explaining option returns within this model depends on the underlying security and option type. For SPX options, VAR accounts for most of the return variation in both puts and calls, explaining 55.0% of SPX put returns (compared to 17.2% for $rs_{j,t+1}^2$) and 42.0% of SPX call returns (18.4% for $rs_{j,t+1}^2$). In contrast, for stock options, $rs_{j,t+1}^2$ plays a more significant role than VAR. On average, $rs_{j,t+1}^2$ explains 38.5% of the return variation in stock puts and 50.4% in stock calls, whereas VAR accounts for 14.5% and 10.4%, respectively. For ETF options, $rs_{j,t+1}^2$ is more important for explaining put returns, while VAR is the primary driver of variation in call returns. Overall, we conclude that the VAR factor has significant explanatory power across asset classes. However, the underlying-specific $rs_{j,t+1}^2$ also explains a substantial amount of return variation. It is therefore critically important to account for it and avoid

biases when estimating the VAR premium in equity and ETF options.

5.2. Patterns in Option Exposures

We next examine the unconditional exposures of delta-hedged option returns through the lens of the model. Table 4 presents the results from the first-stage estimation. To obtain the unconditional exposures for a given underlying asset class, we first compute the weekly exposures for each bucket and average them across all options written on the underlying securities, yielding one observation per week. We then take the average of these weekly exposures over the full sample. Panel A reports the sample average (unconditional) exposures of SPX options for puts and calls separately, as well as by moneyness and maturity buckets. Panels B and C present the results for stock and ETF options.

All three panels in Table 4 indicate that the unconditional exposures align with no-arbitrage option pricing models and are consistent with economic intuition for both calls and puts. Regardless of the underlying security or option type, vegas and gammas are positive on average. Moreover, the relative magnitudes of the estimated exposures align with economic intuition. Gammas tend to be larger for short-maturity and ATM options. In contrast, vegas generally increase with maturity, except for ETF options.

5.3. Estimates of the Systematic Variance Risk Premium

Table 5 presents the estimates of the systematic variance risk premium obtained from SPX options, stock options, and ETF options. Panel A presents estimates obtained using both puts and calls, while Panels B and C separately report on put and call options respectively. All nine sets of estimates report the estimated $\hat{\lambda}_v$ risk premium parameter based on the risk premium specification $(\mathbb{E}_t^P - \mathbb{E}_t^Q)[\text{VAR}_{t+1}] = \lambda_v \cdot \text{VIX}_t^2$ along with the corresponding standard error (obtained from the two-stage procedure described in Section 3.5), and descriptive statistics on the premium.

The most important observation is that the average VAR premiums (see the rows labeled *Avg. VRP*), which are the time-series averages of the product of $\hat{\lambda}_v$ and VIX^2 , are negative in all segments of the equity option markets and across option types. Moreover, all nine estimates of the $\hat{\lambda}_v$ risk premium parameter are statistically significant, and the average variance risk premiums in Panel A are very similar across the three equity option markets:

-20.3% for SPX options, -18.8% for stock options, and -18.5% for ETF options, respectively. The estimates of the average risk premiums for puts and calls in Panels B and C are nearly identical to the full-sample estimates in Panel A, indicating that the pricing of variance risk is integrated across index, stock, and ETF options. This consistency also addresses a potential concern of circularity in the specification of our model for index option returns. Because the variance factor is defined as a linear function of VIX^2 , and VIX itself is constructed from *SPX* option prices, one might interpret evidence of a variance-risk premium in index options as mechanical. Yet the presence of an equally pronounced systematic variance risk premium in stock and ETF options, two markets not used as inputs in the construction of VIX , eliminates this concern and reinforces the conclusion that systematic variance risk is broadly and consistently priced across option markets. Another potential concern is that VIX dynamics reflect both physical variance fluctuations and compensation for systematic variance risk. However, this issue is largely mitigated by our use of weekly VIX^2 increments to construct the factor. Because weekly changes in VIX^2 exhibit a very low signal-to-noise ratio, they primarily capture unexpected innovations rather than shifts in conditional expectations. Consequently, our factor is mainly driven by variance shocks, rather than by fluctuations in the risk premium component.

Are the magnitudes of the estimated VAR risk premiums plausible? We look for guidance in the existing literature. Specifically, because we estimate the price of systematic variance risk, we refer to existing estimates of the index variance risk premium. However, it is critical to realize that the object we estimate differs from the well-known variance risk premium defined through a variance swap (see for example Carr and Wu, 2009; Bollerslev, Tauchen, and Zhou, 2009). Instead, it can be shown that under the null hypothesis of a standard one-factor stochastic volatility model (Heston, 1993), our estimated VAR risk premium equals the spot VRP in that model. For instance, in the seminal study of Eraker (2004), this estimate is -12.3% . Estimates from models with multiple volatility factors are typically larger (in absolute value). The estimates in Table 5 therefore seem plausible.

6. Implementing Factor Models for Option Returns

In this section, we examine the importance of some of the modeling choices we implement, which yield statistically significant negative systematic VAR premium estimates that are very

similar across the three different equity option markets. We show that the most important model feature is to allow for time variation in the VAR premium. We also establish the benefits from imposing a zero intercept in the linear relation between the VAR premium and the variance level. The implications of this assumption are less critical, but we address it first because it is helpful to understand the mechanics underlying the problems emanating from the constant VAR premium assumption.

Next, we compare the results obtained using nonparametrically estimated deltas with results based on Black-Scholes deltas. The use of Black-Scholes deltas yields estimates of the VAR premium that are less consistent across markets due to an omitted variable argument. Finally, we explore the implications of using delevered rather than simple returns. We find that using simple returns leads to large discrepancies between the estimates from different option markets that cannot entirely be resolved by weighting the returns as recommended by Duarte, Jones, and Wang (2024).

6.1. Allowing for an Intercept in the VAR Premium

Based on our analysis of the return on the VAR factor’s replicating portfolio (see Figure 1 and the discussion in Section 3.4), our baseline implementation in Table 5 imposes a zero intercept in the predictive linear relation between the VAR premium and the variance level. We now generalize this specification to allow for a nonzero intercept: $(\mathbb{E}_t^P - \mathbb{E}_t^Q)[\text{VAR}_{t+1}] = \lambda_{v,0} + \lambda_{v,1} \cdot \text{VIX}_t^2$. Table 6 reports the results. As in Table 5, we present estimates of the systematic variance risk premium obtained from SPX options, stock options, and ETF options, and we report separate estimates based on put and call options in addition to full-sample estimates. We report the estimated risk premium parameters for VAR along with their corresponding standard errors, the F -statistics assessing the joint significance of both VAR premium coefficients, and descriptive statistics on the premium.

The most important finding is that while all nine estimates of the average VAR premium in Table 5 remain negative, we now find some notable differences across markets and option types (i.e., call versus put). The average variance risk premiums are -9.1% , -5.4% , and -12.7% for SPX, stock and ETF options respectively. Interestingly, the differences between the estimates for puts and calls are generally larger than the differences between the three equity option markets. The F -statistics indicate statistical significance of the VAR premium

for SPX and stock options, but not for ETF options.

Upon closer inspection, we find that these differences may be due to specification problems. Specifically, the differences between markets and option types are mostly due to differences in the $\hat{\lambda}_{v,0}$ intercept estimates. While all positive, these estimates are statistically significant in only five out of nine cases. In contrast, the estimates of $\hat{\lambda}_{v,1}$, which defines the premium's dependence on VIX_t^2 , are statistically significant in all cases and remarkably consistent, ranging from -0.1057 for SPX puts to -0.0781 for ETF calls and puts combined. None of the $\hat{\lambda}_{v,1}$ coefficient estimates are significantly different from one another at the 5% level in pairwise comparisons.

We conclude that while this specification with an intercept confirms our main result that variance risk premiums are negative for calls and puts in all three markets, imposing a zero intercept in the specification of the variance risk premium results in estimates that are much more consistent across markets. This motivates our choice of a baseline specification with zero intercept, which can be further supported in several ways. First, as discussed in Section 3.4, restricting the intercept to zero eliminates spurious positive expected returns on the variance factor portfolio. The constraint can therefore be motivated by the principle of parsimony to avoid implications that contradict economic theory due to inference from noisy data. Second, the specification of the conditional variance risk premium is equivalent to the specification of the variance-dependent component of the pricing kernel. An extensive literature on term structure models and stochastic volatility option pricing models has documented that specifications of the price of risk that allow for nonzero intercepts result in improved in-sample fit, but this sometimes come at the cost of counterintuitive economic implications.²⁶ In response, the literature on affine option pricing with stochastic volatility usually imposes a zero intercept restriction to address these problems, similar to our baseline specification.²⁷

²⁶See for instance Singleton (2006, p. 392), Heston, Loewenstein, and Willard (2006), Duarte (2003), Cheridito, Filipović, and Kimmel (2007), and Heston, Jacobs, and Kim (2025) on the empirical implications of the specification of the pricing kernel in these models.

²⁷See for instance Bates (1996), Pan (2002), Broadie, Chernov, and Johannes (2007), Hurn, Lindsay, and McClelland (2015), and Bardgett, Gourier, and Leippold (2019).

6.2. Imposing a Constant VAR Premium

We now turn to the most important aspect of the modeling of the variance risk premium: It is critical to model the systematic variance risk premium as time-varying. We demonstrate this by examining the consequences of imposing a constant premium. Table 7 presents the results, which are completely different from the baseline results in Table 5. The three estimates obtained using all options in a given market, reported in Panel A, are now positive (ranging from 0.4% per year in ETF options to 1.9% per year in stock options) and not statistically significant. When options are split into sub-samples by contract type (puts and calls), we observe substantial differences between the estimates. Whereas for SPX options the estimate is negative for puts (−0.6% per year) and positive for calls (5.8% per year), this pattern is reversed for stock and ETF options. For stock options, the difference between the estimates is 12.9 percentage points per year (9.5% in puts vs −3.4% in calls). Overall, the results are much less consistent across markets and option types compared to Table 5 and even Table 6.

To better understand these estimates, we provide some analytical guidance by investigating the consequences of omitting the time-variation in the second-stage regression in a simplified setting that only contains the VAR factor. It is instructive to consider a null hypothesis that includes both a time-varying and constant component. Consider therefore two simplified specifications of the factor model for delta-hedged delevered option returns, with only the VAR factor. In the first, true specification, the VAR premium is affine in VIX_t^2 , while in the second, misspecified model, the VAR premium is constant. These specifications imply the following expressions for expected delta-hedged option returns:

$$E_t[\widehat{r}_{ij,t+1}^\Delta] = \lambda_{v,0} \cdot \mathcal{V}_t^{ij} + \lambda_{v,1} \cdot \mathcal{V}_t^{ij} \cdot VIX_t^2 \quad (24)$$

$$E_t[\widehat{r}_{ij,t+1}^\Delta] = \lambda_{v,c} \cdot \mathcal{V}_t^{ij}. \quad (25)$$

Consider estimating the second specification in equation (25). The resulting OLS estimate's population value is

$$\text{plim } \widehat{\lambda}_{v,c} = \lambda_{v,0} + \lambda_{v,1} \cdot \underbrace{\frac{\text{Cov}(\mathcal{V}_t^{ij}, \mathcal{V}_t^{ij} \cdot VIX_t^2)}{\text{Var}(\mathcal{V}_t^{ij})}}_{\rho}. \quad (26)$$

The main implication of equation (26) is that the constant premium estimate will not necessarily be close to the average of the time-varying premium estimate, but rather depends on ρ , the regression coefficient of the product of exposures and risk premium drivers on exposures. Combining the values of ρ reported in Table 7 with estimates of $\lambda_{v,0}$ and $\lambda_{v,1}$ reported in Table 6 indeed approximates well the constant premium estimates reported in Table 7.

Equation (26) is also helpful to understand the differences between estimates based on put and call options in Panels B and C of Table 7. For stock and ETF options, put-based estimates are higher because the estimate of ρ for puts is smaller than for calls. In these two markets, the estimates of the intercept $\lambda_{v,0}$ in Table 6 for puts and calls are very similar and thus the difference between the intercept estimates are mostly due to the respective values of ρ . For SPX options on the other hand, the estimate of $\lambda_{v,0}$ for puts in Panel B of Table 6 is not statistically significant and much lower than the corresponding estimate for calls in Panel C.

We conclude that it is critical to allow for time-variation in the VAR premium when conducting inference on the compensation for aggregate variance risk. Failing to do so can lead to potentially misleading results and statistically insignificant estimates. The resulting estimate of the constant risk premium is not guaranteed to be close to the average of the time-varying estimates because of co-movement between option exposures and risk premiums.

6.3. Nonparametric versus Black-Scholes Deltas

Much of the empirical literature on equity options relies on Black-Scholes deltas to hedge option returns and draw inferences. This widespread practice stands in contrast to an extensive literature documenting that the Black-Scholes model is severely misspecified and advocating for the use of more refined models. We therefore investigate the role of the nonparametric deltas in the estimation of the VAR premium.

One likely consequence of using Black-Scholes deltas is that the delta-hedged returns may still be exposed to the underlying security. If the underlying security return is not included in the factor model, the resulting exposure estimates suffer from omitted variable bias, which then impacts the second-stage estimates of the VAR premium. The resulting bias is determined by the following quantities: the co-movement of the underlying security return

and the VAR factor, the magnitudes of the VAR and underlying return premiums, and the co-variation between the options' exposures to the omitted factor and their exposures to the other factors estimated in the misspecified model. We elucidate this point further in Internet Appendix C.

Table 8 reports on VAR premium estimates when option returns are first delta-hedged using Black-Scholes deltas before estimating the exposures to the VAR factor, while keeping the second-stage methodology unchanged. The resulting VAR premiums vary widely across underlying securities. The λ_v estimates are much less consistent across underlying securities compared to the baseline estimates in Table 5. Moreover, the resulting risk premiums are not always statistically significant.

We confirmed that the magnitude of the discrepancy between the estimates based on Black-Scholes deltas and those based on our nonparametric deltas is related to the explanatory power of the underlying security return for Black-Scholes delta-hedged returns. We estimated additional specifications of the model which include the underlying security return as an additional factor and we calculate the Shapley-Owen value of the total R^2 for the underlying security return for the three option categories. This explanatory power varies across underlying security categories and option types. For SPX options, the index return explains 3% of Black-Scholes delta-hedged return variation (4% for puts and 1% for calls). For stock options, underlying stock returns explain 6% of return variation (10% for calls and 1% for puts). In contrast, for ETF options, underlying fund returns have virtually zero explanatory power for Black-Scholes delta-hedged options.

We conclude that the use of Black-Scholes deltas can have a detrimental impact on inference when studying the variance risk premium.

6.4. Delevered versus Simple Returns

Finally, we investigate the benefits of using delevered delta-hedged option returns when estimating the model, as opposed to using simple delta-hedged returns. Table 9 presents estimates of the VAR premium obtained using simple delta-hedged returns that are not delevered. These results can once again be compared to the baseline estimates in Table 5. We present two sets of results: the column labeled "Standard" reports estimates from applying our baseline method to simple returns; the column labeled "Weighted" addresses Duarte,

Jones, and Wang’s concern that the use of simple returns may lead to biased estimates due to the presence of noise in the measured option returns, and weights each option observation by the option’s lagged gross return.

Comparing the “Standard” estimates with the baseline results in Table 5, the estimates are also negative in all three segments of the equity option markets. However, while the estimates of λ_v and the resulting average variance risk premiums in Table 5 are very similar across the three markets, the estimates in Table 9 are more variable. The estimates for SPX and ETF options are somewhat smaller (in absolute value), while the estimates for stock options in Panel B are much larger. The “Weighted” estimates indicate that weighting by the lagged gross return, as suggested by Duarte, Jones, and Wang (2024), brings the estimates in all three panels slightly closer to zero, but differences between markets persist.

We conclude that we obtain more consistent estimates of the VAR premium across underlying securities when using delevered returns.

7. Robustness

This section presents a series of robustness tests. First, we examine if alternative observable factors provide incremental value for explaining equity option returns relative to our baseline model, and we discuss the relation between our findings and latent factor models. Next, we examine how consistent the estimates of the price of systematic variance risk are across the maturity spectrum, and we compare estimates from stock options across industries.

7.1. Alternative Observable Factors

After delta-hedging, our baseline model consists of the systematic variance risk (VAR) factor and rs_j^2 , which is underlying-specific and serves as a control variable. It is straightforward to accommodate additional factors with our framework. We focus our empirical analysis on the pricing of the systematic variance risk factor across multiple equity option markets not only because of its economic content, but also because it performs well. We now present evidence on the performance of alternative factors suggested by the literature to further motivate our choice of baseline model.

We present evidence on the following four additional factors: VFTS, the weekly change

in the slope of the term structure of VIX futures prices (6M minus 1M) (Johnson, 2017); HKM, the weekly realization of the intermediary risk factor (He, Kelly, and Manela, 2017); VVIX, the weekly change in the square of the VVIX index; and SKEW, the weekly change in the Cboe SKEW index. Table 10 reports the incremental R_p^2 contribution of each factor, as measured by their Shapley-Owen value, along with the increase in R_p^2 when adding a given factor to the baseline (BL) model. The term structure slope VFTS provides the highest incremental value for SPX, ETF and stock options, regardless of the measure considered, and performs very well for SPX puts. However, note that in comparison to the Shapley-Owen value of VAR and rs_j^2 in Table 3, the explanatory power of VFTS is noticeably smaller in most cases, except for SPX puts. The other three factors do not meaningfully improve on the fit of the baseline model. For these factors, the highest Shapley-Owen value (3.3%) is obtained for the HKM factor in the case of ETF puts. The largest increment in R_p^2 (5.4%) is obtained for the VVIX factor in the case of stock calls.

Overall, the incremental benefit of including additional factors in the baseline model is modest. The most promising alternative factor (VFTS) captures another dimension of variance risk. The three other factors provide limited additional explanatory power and may result in additional noise when estimating exposures and risk premiums. We conclude that, consistent with the literature, it is not straightforward to identify observable factors that can explain the cross-section of option returns, and that the search for additional priced factors should constitute an important component of the future research agenda in this area.

7.2. The Performance of Latent Factors

Building on the work of Kelly, Pruitt, and Su (2018), a growing body of literature employs instrumented principal component analysis (IPCA) to study the cross-section of returns in derivative markets.²⁸ We now document the latent factor structure implied by IPCA in our samples of delta-hedged delevered SPX, stock, and ETF option returns. These IPCA estimates provide a useful benchmark to discuss our model’s explanatory power, but factor models with observable and latent factors are obviously very different, and our goal is not to run a horse race. Our main questions of interest is instead if the latent IPCA factors are

²⁸Büchner and Kelly (2022) find that IPCA models effectively capture the cross-section of monthly S&P 500 index option returns. For stock options, Goyal and Saretto (2024) demonstrate that the returns on volatility strategies can be explained by IPCA-implied factors.

related to the VAR factor, which we investigate below.

Table A.2 in the Appendix and Section D of the Internet Appendix describe the empirical procedure and the set of underlying characteristics we use to estimate the IPCA models. Panels A and B of Table 11 present the total R_p^2 from IPCA models fitted to delta-hedged delevered SPX, stock, and ETF option returns. We report on models with 1 to 5 latent factors. Panel A (B) uses the nonparametric (BS) deltas. Panels C and D report additional R_p^2 s for the stock and ETF options. These panels additionally include the stock/ETF-level characteristics. We can compare the R_p^2 s in Table 11 to the explanatory power for our baseline specification comprising the stock-level squared return (rs_j^2) and the aggregate variance factor (VAR) in Table 3.²⁹ For a given market, the R_p^2 s in Table 11 do not greatly differ across panels. For stock and ETF options, we therefore use the results in Panel C for this comparison. Compared to the two-factor IPCA specification, the R_p^2 s for our baseline model are lower for SPX options. The two-factor IPCA model produces an R_p^2 of 93.46% (91.19%) for puts (calls) whereas our model yields values of 72.2% (60.4%). For stock options, our model with observable rs_j^2 and VAR yields an R_p^2 of 53.00% (60.70%) for put (call) options, higher than that achieved by the two-factor IPCA: 31.45% (24.39%). For ETF options, the explanatory power of IPCA and our model are more similar. The two-factor IPCA model produces an R_p^2 of 53.76% (53.75%) for puts (calls) whereas our model yields values of 44.7% (44.9%). While we emphasize again that these comparisons of model fit need to be interpreted cautiously because the models are very different, one interesting conclusion is that our model with observable factors becomes more competitive in terms of fit when the cross-section displays more variation.

7.3. The Economic Content of Latent Factors

An important question is what the economic content of the IPCA latent factors might be, and if these factors are related to the VAR factor in our observable factor model. We now show that this is indeed the case: there is a very strong (linear) relationship between the IPCA risk factor in a single factor model and the VAR factor. Figure A.1 scatter plots the VAR factor against the IPCA factor for SPX, stock, and ETF equity options, as well as the regression line. The correlation between the VAR factor and the first IPCA factor

²⁹One important caveat is that rs_j^2 is underlying-specific, not an aggregate factor common to all stocks. In the IPCA estimation framework on the other hand, all of the latent factors are aggregate factors.

ranges from 0.69 for stock options to 0.86 for SPX options. Replacing VAR with its traded counterpart, the weekly return on VIX futures (starting in August 2008), we observe similar results with correlations of 0.72 and 0.80 for stock and SPX options respectively. Our results are consistent with the findings of Büchner and Kelly (2022) in the SPX option market and Goyal and Saretto (2024) for stock options, who find that the first IPCA factor is related to changes in the Cboe VIX index, as well as other proxies for systematic risk.

Our analysis reveals that the first latent factor identified by the IPCA model (i.e., the factor that explains the most return variation) across all option types exhibits a strong correlation with the systematic VAR factor. IPCA seeks to form statistically optimal factors. From this perspective, the strong relation between the IPCA risk factor and systematic aggregate variance risk is not surprising; indeed, it lends additional support to the importance of the VAR factor in explaining variation in option returns.

7.4. Systematic Variance Risk Premiums by Maturity

We compare VAR premiums estimated using options with different maturities. For this exercise, we jointly consider puts and calls, and we use either a sample of options with maturities between one and two months, or a sample of options with maturities between two and nine months. Panel A of Table 12 reports the estimates of the second-stage risk premium parameters. The magnitudes of the average VAR premiums are similar for short and long maturities, but the average VAR premium associated with longer-term options slightly exceeds the one associated with shorter-term options for stock and ETF options, while the opposite applies to SPX options. The estimates of $\hat{\lambda}_v$ are statistically significant for all subsamples, except for short-maturity ETF options.

We conclude that systematic variance risk is pervasively priced across the term structure of equity options.

7.5. Systematic Stock Variance Risk Premiums by Industry

So far, our empirical analysis has focused on the pricing of systematic variance risk in index, stock, and ETF option markets. We now investigate patterns in the pricing of systematic variance risk across industries, where we assign each stock in our sample to one of

the five Fama-French industries.³⁰ Panel B of Table 12 presents the resulting estimates. The average variance risk premium is negative in all five industries. There is some variation in the premium across industries, ranging from -24.3% for manufacturing to -9.5% for consumer goods. The estimates of $\hat{\lambda}_v$ are statistically significant for three of the five industries.

Overall, these findings suggest that systematic variance risk in stock options is robustly negatively priced across industries.

8. Conclusion

This paper provides a comprehensive analysis of the pricing of systematic variance risk in the cross-section of index, stock, and ETF option markets. Our empirical investigation analyzes more than 12 million weekly option returns from three distinct markets: options written on the S&P 500 index, on 422 common stocks, and on 96 exchange-traded funds (ETFs).

We make four key contributions to the literature on option pricing and the pricing of variance risk. First, systematic variance risk explains a significant percentage of the variation in delta-hedged option returns. This is an important finding, because the literature has struggled to identify such observable and economically motivated factors. Second, we find that aggregate variance risk is priced in all three equity option markets we investigate. Consistent with theory and economic intuition, the price of systematic variance risk is negative. Third, the magnitude of the systematic variance risk premium estimate is plausible and very similar for the three option markets we investigate, as well as for separate samples of call and put options and across maturities. This finding underscores the pervasive presence of a negative variance risk premium across markets, highlighting the importance of systematic variance risk as a first-order pricing factor that links these integrated markets.

Our fourth contribution concerns the specification, implementation, and estimation of conditional factor models for option returns. Our finding of the pervasive pricing of systematic variance risk in option markets contrasts with mixed evidence in the existing literature. This may be partly due to the scope of our empirical analysis, but we also demonstrate the critical importance of several modeling choices. Specifically, we show that it is important to use nonparametric deltas that generalize Black-Scholes deltas, to allow for time-variation in the

³⁰We use five industries to ensure sufficiently large samples.

price of risk, and that delevering option returns helps address well-known problems due to noise in measured option returns. In future work we plan to use our approach to identify additional factors that explain the cross-section of option returns.




Figure 1. VIX^2 and Returns on the f_v (VAR) Factor Replicating Portfolio. We plot the returns on the VAR factor replicating portfolio (defined in Section 3.4) against the lagged squared value of the Cboe VIX index. The green line shows the fitted values from the predictive regression $r_{v,t+1} = \lambda_0 + \lambda_1 VIX_t^2 + \varepsilon_{t+1}$ where $r_{v,t+1}$ denotes the factor replicating portfolio return. The orange line displays the corresponding nonparametric LOESS fit. Vertical dashed lines mark the 10th through 90th percentiles of the empirical distribution of VIX^2 . The bottom panel repeats the same information using a narrower range on the vertical axis to highlight the variation near the mean.

Table 1
Descriptive Statistics for Option Samples

This table describes our SPX, stock, and ETF option samples. The SPX option sample contains 671,090 observations of weekly option returns. The stock option sample contains 8,822,613 observations of returns for 422 firms. The ETF option sample contains 3,324,607 observations of returns for 96 funds. The table presents summary information on option maturities, relative strikes (K/S), Black-Scholes implied volatilities (in %), and the number of available observations across option types. ATM options are defined via $|\Delta_{ijt}^{BS}| \in [|\Delta_{ATM}^{BS}| \pm 0.1]$ for option j on stock i at time t , by option type, where Δ_{ATM}^{BS} is defined as the median Black-Scholes delta of all options of a given type in the whole sample, whose $K/S \in [0.975, 1.025]$. The remaining options are considered OTM. The sample is weekly and spans the period 1996-01 to 2023-02 for all option types.

τ		Maturity (months)			Moneyness (K/S)			Implied volatility (%)			N
		Avg.	Q05	Q95	Avg.	Q05	Q95	Avg.	Q05	Q95	
<i>Panel A: SPX options</i>											
Puts	All	2.5	1.0	6.5	0.91	0.78	1.00	24.1	12.8	37.9	417,693
	[1M, 2M]	1.4	1.0	1.9	0.91	0.79	1.00	24.2	12.1	39.6	232,892
	(2M, 9M]	3.8	2.1	7.7	0.91	0.77	1.01	23.8	13.7	35.5	184,801
Calls	All	2.6	1.0	6.8	1.06	0.99	1.17	16.1	8.6	26.2	253,397
	[1M, 2M]	1.4	1.0	1.9	1.05	0.99	1.14	15.6	8.1	26.8	132,379
	(2M, 9M]	3.9	2.1	7.9	1.07	0.99	1.20	16.6	9.3	25.7	121,018
<i>Panel B: Stock options</i>											
Puts	All	3.5	1.0	7.5	0.89	0.66	1.02	40.6	19.9	77.3	4,182,878
	[1M, 2M]	1.3	1.0	1.9	0.92	0.76	1.01	40.7	19.1	79.6	1,614,943
	(2M, 9M]	4.8	2.4	7.9	0.87	0.62	1.03	40.5	20.5	76.0	2,567,935
Calls	All	3.6	1.0	7.7	1.12	0.98	1.41	36.5	17.6	69.6	4,639,735
	[1M, 2M]	1.3	1.0	1.9	1.07	0.98	1.24	37.2	17.2	72.9	1,662,698
	(2M, 9M]	4.9	2.4	7.9	1.15	0.98	1.47	36.1	17.8	67.8	2,977,037
<i>Panel C: ETF options</i>											
Puts	All	4.2	1.0	7.7	0.91	0.71	1.02	28.0	14.0	51.0	1,978,547
	[1M, 2M]	1.4	1.0	1.9	0.94	0.83	1.01	28.6	13.1	54.4	516,581
	(2M, 9M]	5.2	2.4	7.9	0.90	0.69	1.02	27.8	14.4	49.7	1,461,966
Calls	All	4.2	1.0	7.9	1.07	0.98	1.26	23.8	11.2	45.6	1,346,060
	[1M, 2M]	1.4	1.0	1.9	1.04	0.99	1.14	25.3	10.8	50.2	371,216
	(2M, 9M]	5.3	2.4	7.9	1.08	0.98	1.29	23.2	11.4	43.7	974,844

Table 2
Summary Statistics: Delta-Hedged Delevered Option Returns

The table reports summary statistics of delta-hedged delevered returns for SPX, stock, and ETF options. For stock and ETF options, we report averages of summary statistics calculated for each underlying. ER denotes the average return, SD denotes standard deviation, Sk denotes skewness, and K denotes kurtosis. For stock and ETF options, “% sig.” denotes the fraction of average return estimates that are significantly different from zero at the 0.05 confidence level. For SPX options, a ✓ denotes statistical significance at the 0.05 level and ✗ a lack thereof. ER and SD are annualized using 52 weeks per year. ER, SD, and “% sig.” are reported in %. We use Driscoll and Kraay (1998) standard errors to determine statistical significance.

Mon.	τ	Put options					Call options				
		ER	% sig.	SD	Sk	K	ER	% sig.	SD	Sk	K
Panel A: SPX options											
All	All	−2.19	✓	2.73	3.5	55.4	−2.06	✓	2.80	2.3	25.5
	[1M, 2M]	−2.38	✓	2.64	3.9	55.9	−2.46	✓	2.92	2.7	26.6
	(2M, 9M]	−2.05	✓	2.89	2.7	51.5	−1.57	✓	2.75	2.2	25.3
ATM	All	−2.61	✓	3.09	3.2	36.1	−2.85	✓	2.98	2.9	28.8
	[1M, 2M]	−3.21	✓	3.25	3.5	35.1	−3.43	✓	3.20	3.2	30.6
	(2M, 9M]	−2.14	✓	3.00	2.7	36.0	−2.36	✓	2.86	2.6	27.0
OTM	All	−2.04	✓	2.63	3.6	66.5	−1.65	✓	2.71	2.0	23.3
	[1M, 2M]	−2.12	✓	2.43	3.9	67.0	−2.04	✓	2.80	2.4	24.0
	(2M, 9M]	−2.02	✓	2.88	2.8	59.7	−1.15	✓	2.69	1.9	24.3
Panel B: Stock options											
All	All	−0.50	16	5.81	4.3	59.0	−0.45	19	7.35	4.5	79.2
	[1M, 2M]	−1.20	24	6.51	4.0	48.5	−1.69	35	7.87	3.5	41.4
	(2M, 9M]	0.01	9	5.42	4.3	61.1	0.29	12	7.29	5.0	94.2
ATM	All	−0.36	13	6.49	4.0	47.5	−0.54	18	6.90	3.8	49.4
	[1M, 2M]	−1.22	19	7.24	3.7	38.8	−1.52	26	7.53	3.3	34.5
	(2M, 9M]	0.35	9	6.02	4.0	48.0	−0.08	12	6.72	3.7	47.9
OTM	All	−0.55	16	5.39	4.3	61.9	−0.45	18	7.44	4.6	83.3
	[1M, 2M]	−1.24	23	6.00	3.9	48.2	−1.92	38	7.93	3.1	36.4
	(2M, 9M]	−0.15	9	5.05	4.1	60.9	0.50	10	7.75	5.2	99.4
Panel C: ETF options											
All	All	−0.69	21	4.40	3.1	43.4	−0.98	33	4.70	2.4	27.8
	[1M, 2M]	−1.35	34	4.32	3.3	33.0	−2.78	63	4.74	2.2	18.8
	(2M, 9M]	−0.48	18	4.31	2.8	42.7	−0.50	16	4.63	2.3	29.1
ATM	All	−0.96	29	4.71	2.9	31.3	−1.29	28	4.71	2.1	21.4
	[1M, 2M]	−1.73	34	4.72	3.3	30.2	−2.48	57	4.63	2.2	19.4
	(2M, 9M]	−0.74	15	4.55	2.5	28.0	−1.07	18	4.66	1.9	21.3
OTM	All	−0.57	25	3.89	3.2	50.2	−0.97	32	4.42	2.5	27.9
	[1M, 2M]	−0.99	24	3.88	3.2	31.5	−4.05	70	4.66	1.9	13.7
	(2M, 9M]	−0.45	17	3.82	2.9	48.5	−0.19	16	4.24	2.6	29.0

Table 3
Model Explanatory Power

The table reports on the explanatory power of the baseline specification for returns on SPX, stock, and ETF call and put options (separately). We report the total (pooled) R_p^2 (in column labeled “Tot.”) and the Shapley-Owen value for the VAR factor and the rs_j^2 underlying asset-specific control. The total R_p^2 is defined as $R_p^2 = 1 - \sum_{ijt} (r_{ijt}^\Delta - \hat{r}_{ijt}^\Delta)^2 / \sum_{ijt} (r_{ijt}^\Delta)^2$, where r_{ijt}^Δ is the delta-hedged delevered return on option j written on underlying security i on day t , and $\hat{r}_{ijt}^\Delta$ is the fitted value from the model. All quantities are reported in percent.

		Put options			Call options		
Mon.	τ	VAR	rs_j^2	Tot.	VAR	rs_j^2	Tot.
<i>Panel A: SPX options</i>							
All	All	55.0	17.2	72.2	42.0	18.4	60.4
	[1M, 2M]	48.1	15.5	63.6	40.0	21.0	61.0
	(2M, 9M]	60.8	18.6	79.5	44.4	15.2	59.6
ATM	All	53.1	24.5	77.6	47.4	24.8	72.3
	[1M, 2M]	53.0	24.0	77.0	48.1	25.7	73.8
	(2M, 9M]	53.3	25.0	78.3	46.7	23.9	70.6
OTM	All	55.9	13.7	69.6	39.0	14.9	53.9
	[1M, 2M]	45.6	11.0	56.6	35.8	18.7	54.5
	(2M, 9M]	64.2	15.9	80.0	43.0	10.1	53.1
<i>Panel B: Stock options</i>							
All	All	14.5	38.5	53.0	10.4	50.4	60.7
	[1M, 2M]	8.6	42.7	51.3	8.9	53.3	62.2
	(2M, 9M]	19.6	34.9	54.5	11.6	48.0	59.6
ATM	All	15.9	45.5	61.4	11.6	47.7	59.2
	[1M, 2M]	11.5	50.7	62.2	9.4	51.1	60.5
	(2M, 9M]	20.2	40.4	60.6	13.0	45.3	58.4
OTM	All	13.5	33.4	46.9	9.7	51.8	61.6
	[1M, 2M]	6.2	35.9	42.0	8.7	54.3	62.9
	(2M, 9M]	19.2	31.4	50.6	10.7	49.7	60.3
<i>Panel C: ETF options</i>							
All	All	19.7	25.0	44.7	25.5	19.4	44.9
	[1M, 2M]	13.0	30.4	43.4	27.0	16.3	43.3
	(2M, 9M]	23.0	22.3	45.3	24.7	21.0	45.7
ATM	All	21.4	28.1	49.5	24.3	23.7	47.9
	[1M, 2M]	20.9	36.4	57.3	23.3	30.4	53.7
	(2M, 9M]	21.6	23.8	45.4	24.7	20.6	45.3
OTM	All	18.3	22.2	40.5	26.6	15.5	42.1
	[1M, 2M]	5.5	24.7	30.2	29.8	5.5	35.3
	(2M, 9M]	24.1	21.1	45.2	24.7	21.5	46.2

Table 4
Option Exposures

This table reports the average exposures of SPX, stock, and ETF options to the VAR factor and rs_j^2 by option type, maturity (across columns), and moneyness (across rows).

Mon. / τ	Put options						Call options					
	VAR			rs_j^2			VAR			rs_j^2		
	All	1-2M	2-9M	All	1-2M	2-9M	All	1-2M	2-9M	All	1-2M	2-9M
<i>Panel A: SPX options</i>												
All	0.14	0.14	0.15	1.47	1.52	1.41	0.09	0.08	0.09	1.29	1.40	1.18
ATM	0.14	0.13	0.15	1.48	1.60	1.38	0.13	0.13	0.14	1.46	1.58	1.35
OTM	0.14	0.14	0.15	1.47	1.50	1.42	0.07	0.06	0.07	1.23	1.34	1.10
<i>Panel B: Stock options</i>												
All	0.11	0.10	0.12	1.14	1.48	0.93	0.09	0.09	0.09	1.21	1.50	1.06
ATM	0.12	0.10	0.13	1.33	1.60	1.14	0.12	0.10	0.12	1.25	1.56	1.09
OTM	0.11	0.10	0.11	1.06	1.43	0.84	0.08	0.08	0.08	1.19	1.46	1.04
<i>Panel C: ETF options</i>												
All	0.11	0.12	0.11	1.30	1.76	1.14	0.10	0.11	0.09	1.47	1.88	1.32
ATM	0.11	0.12	0.11	1.51	1.94	1.35	0.11	0.12	0.11	1.52	1.95	1.36
OTM	0.11	0.12	0.11	1.20	1.66	1.05	0.08	0.11	0.07	1.44	1.83	1.29

Table 5
Estimates of the VAR Premium

This table reports estimates of the parameters characterizing the systematic variance risk premium and the time-series properties of the premium for our baseline specification. The premium is $(\mathbb{E}_t^P - \mathbb{E}_t^Q)[\text{VAR}_{t+1}] = \lambda_v \cdot \text{VIX}_t^2$. Standard errors are reported in parentheses. The following descriptive statistics of the VAR premium estimates are also reported: the average annualized premium (Avg. VRP) and its annualized standard deviation.

	SPX	Stock	ETF
<i>Panel A: All options</i>			
λ_v	−0.0804** (0.0331)	−0.0743* (0.0399)	−0.0737* (0.0408)
Avg. VRP	−20.3	−18.8	−18.5
Std. dev.	2.9	2.7	2.8
<i>Panel B: Put options</i>			
λ_v	−0.1043** (0.0425)	−0.0714* (0.0379)	−0.0708* (0.0417)
Avg. VRP	−26.4	−18.1	−17.8
Std. dev.	3.8	2.6	2.6
<i>Panel C: Call options</i>			
λ_v	−0.0677** (0.0292)	−0.0843** (0.0430)	−0.0700* (0.0410)
Avg. VRP	−17.1	−21.3	−17.6
Std. dev.	2.5	3.1	2.6

Table 6
Estimates of the Affine VAR Premium

This table reports estimates of the parameters characterizing the systematic variance risk premium and the time-series properties of the premium for the affine specification. The premium is $(\mathbb{E}_t^P - \mathbb{E}_t^Q)[\text{VAR}_{t+1}] = \lambda_{v,0} + \lambda_{v,1} \cdot \text{VIX}_t^2$. Standard errors are reported in parentheses. The following descriptive statistics of the VAR premium estimates are also reported: the average annualized premium (Avg. VRP) and its annualized standard deviation.

	SPX	Stock	ETF
<i>Panel A: All options</i>			
$\lambda_{v,0}$	0.0026** (0.0012)	0.0032** (0.0014)	0.0013 (0.0011)
$\lambda_{v,1}$	-0.0899** (0.0361)	-0.0881** (0.0444)	-0.0781* (0.0438)
<i>F</i> -stat.	6.72**	5.85*	3.30
<i>p</i> -val.	(0.03)	(0.05)	(0.19)
Avg. VRP	-9.1	-5.4	-12.7
Std. dev.	3.3	3.2	2.9
<i>Panel B: Put options</i>			
$\lambda_{v,0}$	0.0014 (0.0011)	0.0038*** (0.0012)	0.0015 (0.0011)
$\lambda_{v,1}$	-0.1057** (0.0426)	-0.0826** (0.0408)	-0.0732* (0.0429)
<i>F</i> -stat.	7.45**	10.95***	3.04
<i>p</i> -val.	(0.02)	(0.00)	(0.22)
Avg. VRP	-19.6	-1.2	-10.5
Std. dev.	3.8	3.0	2.7
<i>Panel C: Call options</i>			
$\lambda_{v,0}$	0.0040*** (0.0015)	0.0034** (0.0016)	0.0014 (0.0012)
$\lambda_{v,1}$	-0.0989** (0.0385)	-0.0996** (0.0475)	-0.0765* (0.0460)
<i>F</i> -stat.	7.88**	5.07*	3.07
<i>p</i> -val.	(0.02)	(0.08)	(0.22)
Avg. VRP	-4.3	-7.7	-12.2
Std. dev.	3.6	3.6	2.9

Table 7
Estimates of the Constant VAR Premium

This table reports the estimates of the VAR premium when it is constrained to be constant. Standard errors are reported in parentheses. Additionally reported are ρ , the coefficient from the regression of $\widehat{\mathcal{V}}_t^{ij} \cdot VIX_t^2$ on $\widehat{\mathcal{V}}_t^{ij}$, and the average annualized risk premium (Avg. VRP).

	SPX	Stock	ETF
<i>Panel A: All options</i>			
$\lambda_{v,c}$	0.0002 (0.0009)	0.0004 (0.0007)	0.0001 (0.0007)
ρ	0.0178	0.0332	0.0173
Avg. VRP	1.3	1.9	0.4
<i>Panel B: Put options</i>			
$\lambda_{v,c}$	-0.0001 (0.0014)	0.0018*** (0.0007)	0.0009 (0.0009)
ρ	-0.0044	0.0227	0.0101
Avg. VRP	-0.6	9.5	4.7
<i>Panel C: Call options</i>			
$\lambda_{v,c}$	0.0011 (0.0009)	-0.0007 (0.0012)	-0.0003 (0.0007)
ρ	0.0350	0.0421	0.0246
Avg. VRP	5.8	-3.4	-1.6

Table 8
Estimates of the VAR Premium. Black-Scholes Delta-Hedged Returns

This table reports estimates of the parameters characterizing the systematic variance risk premium, and the time-series properties of the premium. Before estimating the model with the VAR factor and rs_j^2 of underlying, we hedge each option using its Black-Scholes delta and the return on the underlying asset. The first column of each panel presents the estimates based on all options, whereas the second and third columns report estimates from the subsamples of put and call options, respectively. The premium is proportional to VIX_t^2 : $(\mathbb{E}_t^P - \mathbb{E}_t^Q)[VAR_{t+1}] = \lambda_v \cdot VIX_t^2$. Standard errors are reported in parentheses. The following descriptive statistics of the VAR premium estimates are also reported: the average annualized risk premium (Avg. VRP) and its annualized standard deviation (all in percent).

	SPX	Stock	ETF
<i>Panel A: All options</i>			
λ_v	−0.1033** (0.0455)	−0.0477 (0.0341)	−0.0540 (0.0489)
Avg. VRP	−26.1	−12.0	−13.6
Std. dev.	3.7	1.7	2.0
<i>Panel B: Put options</i>			
λ_v	−0.1078** (0.0483)	−0.0636* (0.0376)	−0.0544 (0.0485)
Avg. VRP	−27.2	−16.1	−13.7
Std. dev.	3.9	2.3	2.0
<i>Panel C: Call options</i>			
λ_v	−0.1260** (0.0563)	−0.0580* (0.0345)	−0.0612 (0.0519)
Avg. VRP	−31.9	−14.7	−15.4
Std. dev.	4.6	2.1	2.3

Table 9
Estimates of the VAR Premium. Simple Returns

This table reports estimates of the parameters characterizing the systematic variance risk premium obtained from simple option returns, and the time-series properties of the premium. The first column of each panel labeled “Standard” presents the estimates obtained from our standard procedure. The second column labeled “Weighted” contains the estimates obtained by additionally weighting option observations by lagged gross returns, as in Duarte, Jones, and Wang (2024). The premium is proportional to VIX_t^2 : $(\mathbb{E}_t^P - \mathbb{E}_t^Q)[VAR_{t+1}] = \lambda_v \cdot VIX_t^2$. Standard errors are reported in parentheses. The following descriptive statistics of the VAR premium estimates are also reported: the average annualized risk premium (Avg. VRP) and its annualized standard deviation (all in percent).

	Standard	Weighted
<i>Panel A: SPX options</i>		
λ_v	−0.0658* (0.0373)	−0.0293 (0.0375)
Avg. VRP	−16.6	−7.4
Std. dev.	2.4	1.1
<i>Panel B: Stock options</i>		
λ_v	−0.1773** (0.0815)	−0.1492** (0.0655)
Avg. VRP	−44.8	−37.7
Std. dev.	6.4	5.4
<i>Panel C: ETF options</i>		
λ_v	−0.1067*** (0.0224)	−0.0970*** (0.0210)
Avg. VRP	−26.8	−24.4
Std. dev.	4.0	3.6

Table 10
Explanatory Power of Alternative Factors

This table reports on the explanatory power of four alternative factors for the returns on SPX, stock, and ETF options. The factors are: VFTS, the change in the slope of the term structure of VIX futures; HKM, the intermediary factor of He, Kelly, and Manela (2017), VVIX, the change in the square of the VVIX index, and SKEW, the change in the Cboe SKEW index. For each factor, and for put and call options separately, we report the Shapley-Owen value of R^2 and the increment in R^2 vis-à-vis our baseline (BL) specification with VAR and rs_j^2 . The total R_p^2 is defined as $R_p^2 = 1 - \sum_{ijt} (r_{ijt}^\Delta - \hat{r}_{ijt}^\Delta)^2 / \sum_{ijt} (r_{ijt}^\Delta)^2$, where r_{ijt}^Δ is the delta-hedged return on option j written on underlying security i on day t , and $\hat{r}_{ijt}^\Delta$ is the fitted value from the model. All quantities are reported in percent.

	Type	VFTS		HKM		VVIX		SKEW	
		SO R^2	vs. BL	SO R^2	vs. BL	SO R^2	vs. BL	SO R^2	vs. BL
SPX	Calls	-2.1	6.1	0.4	3.1	-0.2	4.6	0.5	0.7
	Puts	23.4	8.1	2.9	6.1	3.7	2.8	1.3	0.8
Stock	Calls	7.3	8.1	1.2	0.9	0.9	5.4	0.7	0.8
	Puts	9.8	6.8	1.5	1.0	0.4	3.5	0.6	0.5
ETF	Calls	8.0	4.5	2.5	1.8	0.4	2.6	0.7	0.6
	Puts	8.5	6.9	3.3	4.3	1.3	2.1	0.7	0.7

Table 11
IPCA Factor Model Explanatory Power

This table reports on the explanatory power of models estimated via IPCA using returns on SPX, stock, and ETF equity options. We consider models with 1-5 factors. For put and call options separately, we report the total (pooled) R_p^2 for each latent factor model. The R_p^2 is defined as $R_p^2 = 1 - \sum_{ijt} (r_{ijt}^\Delta - \hat{r}_{ijt}^\Delta)^2 / \sum_{ijt} (r_{ijt}^\Delta)^2$, where r_{ijt}^Δ is the delta-hedged delevered return on option j written on underlying security i on day t , and $\hat{r}_{ijt}^\Delta$ is the fitted value from the model. All quantities are reported in %. All models are estimated pooled over puts and calls including a constant with all characteristics interacted with a dummy variable encoded as 1 for puts and zero for calls. Panels A (B) are estimated with the nonparametric (Black-Scholes) delevered delta-hedged option returns with 10 option-based characteristics (22 variables in total, including the constant and put dummy interactions). Panel C (D) follows the same procedure with an additional 16 equity-based characteristics (54 variables in total, including the constant and put dummy interactions).

	Put options					Call options				
	# of Latent IPCA Factors					# of Latent IPCA Factors				
	1	2	3	4	5	1	2	3	4	5
<i>Panel A: Nonparametrically estimated delta-hedged returns</i>										
SPX	76.64	93.46	96.18	97.43	97.89	75.78	91.19	94.98	96.57	97.19
Stock	29.81	30.03	30.95	32.24	33.35	17.91	23.30	24.78	25.75	26.40
ETF	44.88	50.41	53.32	54.88	56.01	45.91	50.72	54.19	55.96	57.13
<i>Panel B: Black-Scholes delta-hedged returns</i>										
SPX	92.50	96.21	97.65	98.91	98.90	86.58	92.26	95.00	95.76	97.81
Stock	32.73	33.68	35.08	36.07	36.79	19.20	21.43	22.44	23.01	23.50
ETF	45.32	50.54	53.15	54.60	55.80	42.99	47.59	50.84	52.53	53.47
<i>Panel C: Nonparametrically delta-hedged returns – Including additional underlying-based characteristics</i>										
Stock	30.53	31.45	32.65	34.39	35.42	18.53	24.39	26.88	28.07	29.10
ETF	47.73	53.76	57.02	58.96	60.58	47.64	53.75	57.30	59.47	61.66
<i>Panel D: Black-Scholes delta-hedged returns – Including additional underlying-based characteristics</i>										
Stock	33.75	34.85	37.21	38.40	39.47	19.58	23.37	24.92	26.17	27.08
ETF	48.39	53.41	57.34	59.72	61.09	45.66	50.80	54.71	57.49	59.72

Table 12
Estimates of the VAR Premium in Sub-samples

This table reports estimates of the parameters characterizing the systematic variance risk premium obtained from sub-samples of our dataset. Panel A contains estimates obtained from SPX, stock, and ETF options samples at a shorter (1–2 months, inclusive) and longer (2–9 months) maturity. Panel B contains estimates obtained from stock options divided into five Fama-French industry groups. The premium is proportional to VIX_t^2 : $(\mathbb{E}_t^P - \mathbb{E}_t^Q)[VAR_{t+1}] = \lambda_v \cdot VIX_t^2$. Standard errors are reported in parentheses. The following descriptive statistics of the VAR premium estimates are also reported: the average annualized risk premium (Avg. VRP) and its annualized standard deviation (all in percent).

<i>Panel A: by Maturity</i>						
	SPX options		Stock options		ETF options	
	1M–2M	2M–9M	1M–2M	2M–9M	1M–2M	2M–9M
λ_v	−0.085** (0.038)	−0.079** (0.032)	−0.082* (0.045)	−0.091** (0.043)	−0.064 (0.042)	−0.072* (0.043)
Avg. VRP	−21.4	−20.0	−20.7	−23.1	−16.1	−18.2
Std. dev.	3.1	2.9	3.0	3.3	2.4	2.7
<i>Panel B: in Stock Options by Industry</i>						
	Consumer	Health	Manuf.	Tech	Other	
λ_v	−0.037 (0.047)	−0.051 (0.051)	−0.096* (0.050)	−0.072** (0.035)	−0.094* (0.054)	
Avg. VRP	−9.5	−12.9	−24.3	−18.2	−23.7	
Std. dev.	1.4	1.9	3.5	2.6	3.4	

Appendix: Supplementary Tables and Figures

Table A.1
Predictive Regressions for the VAR Replicating Portfolio

This table reports estimates of predictive regressions of the return on the VAR-replicating portfolio on the squared VIX index, $r_{v,t+1} = \lambda_0 + \lambda_1 VIX_t^2 + \varepsilon_{t+1}$. The regressions in columns (1) and (3) are estimated on the full sample, with the constant in model (3) restricted to zero. The regression in column (2) is estimated on a sample consisting of VIX^2 values below the 90th percentile of its distribution.

	(1)	(2)	(3)
λ_0	0.0011 (0.0016)	0.0012 (0.0009)	
λ_1	-0.0774* (0.0437)	-0.0818*** (0.0312)	-0.0667** (0.0278)

Table A.2
List of Option and Equity Characteristics Used in the IPCA Estimation.

This table presents the list of option (Panel A) and equity (Panel B) characteristics used in the estimation of the latent IPCA factors. For each characteristic, we report its ID, a brief description, and the data source. All variables are sampled weekly and aligned in time with each option contract at week t .

#	ID	Description	Reference	Source
<i>Panel A: Option Characteristics</i>				
1	AbsSpread	Absolute bid-ask spread	–	OptionMetrics
2	Delta	Black-Scholes option delta	Buchner & Kelly (2022)	OptionMetrics
3	EmbedLev	$ \Delta \cdot S / \text{MidPrice} $	Buchner & Kelly (2022)	OptionMetrics
4	IV	Implied volatility	Buchner & Kelly (2022)	OptionMetrics
5	MidPrice	Midpoint of bid and ask (log)	Goyal & Saretto (2024)	OptionMetrics
6	Moneyess	Strike divided by underlying price	Buchner & Kelly (2022)	OptionMetrics
7	OpenInt	Open interest	Goyal & Saretto (2024)	OptionMetrics
8	RelSpread	Relative spread (AbsSpread / MidPrice)	Goyal & Saretto (2024)	OptionMetrics
9	TTM	Option time to maturity	Buchner & Kelly (2022)	OptionMetrics
10	Volume	Option volume traded	Goyal & Saretto (2024)	OptionMetrics
<i>Panel B: Equity Characteristics</i>				
11	Amihud	Amihud illiquidity ratio (5-day avg.)	Amihud (2002)	CRSP
12	AutoCorr	Rolling autocorr. of log returns (252-days)	Jeon et al. (2021)	CRSP
13	Beta	Rolling CAPM beta (252-day window)	Ang et al. (2006)	CRSP
14	BidAsk	Daily stock bid-ask spread	–	CRSP
15	IdiosynVol	Residual SD from 3-factor model (252-day window)	Ang et al. (2006)	CRSP
16	MarketCap	Log of market capitalization	Goyal & Saretto (2024)	CRSP
17	Max10	Rolling avg. of 10 highest daily returns (3-month window)	Bali et al. (2011)	CRSP
18	Mom12	12-month return (skip last month)	Goyal & Saretto (2024)	CRSP
19	Mom6	6-month return (skip last month)	–	CRSP
20	Price	Log of stock price	–	CRSP
21	Rkurt	Rolling return kurtosis (252-day window)	Goyal & Saretto (2024)	CRSP
22	Rskew	Rolling return skewness (252-day window)	Goyal & Saretto (2024)	CRSP
23	RV	Rolling realized volatility (252-day window)	Goyal & Saretto (2024)	CRSP
24	StockRetM	Monthly total return	–	CRSP
25	StockRetW	Weekly total return	Goyal & Saretto (2024)	CRSP
26	VolDev	$\log(\text{RV}) - \log(\text{IV})$	Goyal & Saretto (2024)	CRSP/OptionMetrics




Figure A.1. The IPCA Single-Factor Model and the Systematic VAR Factor. The figure shows the relation between the returns from the single-factor IPCA model (y -axis) and the VAR factor (x -axis) for the samples of SPX, stock, and ETF equity options. The single-factor IPCA model is estimated with the ten option-level characteristics documented in Table A.2 of the Appendix including a constant and put dummy interactions (22 characteristics in total).

REFERENCES

- Andersen, Torben G., Nicola Fusari, and Viktor Todorov, 2017, Short-Term Market Risks Implied by Weekly Options, *The Journal of Finance* 72, 1335–1386.
- Ang, Andrew, Robert J. Hodrick, Yuhang Xing, and Xiaoyan Zhang, 2006, The cross-section of volatility and expected returns, *The Journal of Finance* 61, 259–299.
- Aït-Sahalia, Yacine, Mustafa Karaman, and Lorian Mancini, 2020, The term structure of equity and variance risk premia, *Journal of Econometrics* 219, 204–230, Annals Issue: Econometric Estimation and Testing: Essays in Honour of Maxwell King.
- Bakshi, Gurdip, and Nikunj Kapadia, 2003, Volatility Risk Premiums Embedded in Individual Equity Options, *The Journal of Derivatives* 11, 45–54.
- Bali, Turan G, Heiner Beckmeyer, Mathis Mörke, and Florian Weigert, 2023, Option return predictability with machine learning and big data, *The Review of Financial Studies* 36, 3548–3602.
- Bardgett, Chris, Elise Gourier, and Markus Leippold, 2019, Inferring volatility dynamics and risk premia from the S&P 500 and VIX markets, *Journal of Financial Economics* 131, 593–618.
- Barras, Laurent, and Aytek Malkhozov, 2016, Does variance risk have two prices? Evidence from the equity and option markets, *Journal of Financial Economics* 121, 79–92.
- Bates, David S, 1996, Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutschemark Options, *Review of Financial Studies* 9, 69–107.
- Bates, David S., 2000, Post-'87 crash fears in the S&P 500 futures option market, *Journal of Econometrics* 94, 181–238.
- Bégin, Jean-François, Christian Dorion, and Geneviève Gauthier, 2020, Idiosyncratic Jump Risk Matters: Evidence from Equity Returns and Options, *The Review of Financial Studies* 33, 155–211.
- Bollerslev, T, George Tauchen, and H Zhou, 2009, Expected Stock Returns and Variance Risk Premia, *Review of Financial Studies* 22 (11), pp. 4463–4492.

- Bollerslev, Tim, and Viktor Todorov, 2011, Tails, fears, and risk premia, *The Journal of Finance* 66, 2165–2211.
- Bollerslev, Tim, Viktor Todorov, and Lai Xu, 2015, Tail risk premia and return predictability, *Journal of Financial Economics* 118, 113–134.
- Britten-Jones, Mark, and Anthony Neuberger, 2000, Option prices, implied price processes, and stochastic volatility, *The Journal of Finance* 55, 839–866.
- Broadie, Mark, Mikhail Chernov, and Michael Johannes, 2007, Model specification and risk premia: Evidence from futures options, *Journal of Finance* 62 (3), 1–62.
- Broadie, Mark, Mikhail Chernov, and Michael Johannes, 2009, Understanding index option returns, *The Review of Financial Studies* 22, 4493–4529.
- Bryzgalova, Svetlana, Jiantao Huang, and Christian Julliard, 2023, Bayesian solutions for the factor zoo: We just ran two quadrillion models, *The Journal of Finance* 78, 487–557.
- Büchner, Matthias, and Bryan Kelly, 2022, A factor model for option returns, *Journal of Financial Economics* 143, 1140–1161.
- Cao, Jie, and Bing Han, 2013, Cross-Section of Option Returns and Idiosyncratic Stock Volatility, *Journal of Financial Economics* 108, 231–249.
- Cao, Jie, Bing Han, Qing Tong, and Xintong Zhan, 2022, Option Return Predictability, *Review of Financial Studies* 35, 1394–1432.
- Carr, P, and L Wu, 2009, Variance Risk Premia, *Review of Financial Studies* 22, 1311–1341.
- Carr, Peter, and Dilip Madan, 1998, Towards a theory of volatility trading, *Volatility: New estimation techniques for pricing derivatives* 29, 417–427.
- Cheridito, Patrick, Damir Filipović, and Robert L Kimmel, 2007, Market price of risk specifications for affine models: Theory and evidence, *Journal of Financial Economics* 83, 123–170.
- Christoffersen, Peter, Mathieu Fournier, and Kris Jacobs, 2018, The Factor Structure in Equity Options, *The Review of Financial Studies* 31, 595–637.

- Christoffersen, Peter, Ruslan Goyenko, Kris Jacobs, and Mehdi Karoui, 2018, Illiquidity Premia in the Equity Options Market, *The Review of Financial Studies* 31, 811–851.
- Detzel, Andrew, Jefferson Duarte, Avraham Kamara, Stephan Siegel, and Celine Sun, 2023, The cross-section of volatility and expected returns: Then and now, *Critical Finance Review* 12, 9–56.
- Driessen, Joost, Pascal J. Maenhout, and Grigory Vilkov, 2009, The price of correlation risk: Evidence from equity options, *Journal of Finance* 64, 1377–1406.
- Driscoll, John C., and Aart C. Kraay, 1998, Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data, *Review of Economics and Statistics* 80, 549–560.
- Duarte, Jefferson, 2003, Evaluating an alternative risk preference in affine term structure models, *The Review of Financial Studies* 17, 379–404.
- Duarte, Jefferson, Christopher S. Jones, and Junbo L. Wang, 2024, Very noisy option prices and inference regarding the volatility risk premium, *The Journal of Finance* .
- Eraker, Bjørn, 2004, Do Stock Prices and Volatility Jump? Evidence From Spot and Option Prices, *Journal of Finance* 59, 1367–1403.
- Fournier, Mathieu, Kris Jacobs, and Piotr Orłowski, 2024, Modeling conditional factor risk premia implied by index option returns, *The Journal of Finance* 79, 2289–2338.
- Gagliardini, Patrick, Elisa Ossola, and Olivier Scaillet, 2016, Time-Varying Risk Premium in Large Cross-Sectional Equity Data Sets, *Econometrica* 84, 985–1046.
- Giglio, Stefano, and Dacheng Xiu, 2021, Asset pricing with omitted factors, *Journal of Political Economy* 129, 1947–1990.
- Goyal, Amit, and Alessio Saretto, 2009, Cross-section of option returns and volatility, *Journal of Financial Economics* 94, 310–326.
- Goyal, Amit, and Alessio Saretto, 2024, Can equity option returns be explained by a factor model? IPCA says yes, *The Review of Financial Studies* 38, 1783–1821.

- He, Zhiguo, Bryan Kelly, and Asaf Manela, 2017, Intermediary asset pricing: New evidence from many asset classes, *Journal of Financial Economics* 126, 1–35.
- Herskovic, Bernard, Bryan Kelly, Hanno Lustig, and Stijn Van Nieuwerburgh, 2016, The common factor in idiosyncratic volatility: Quantitative asset pricing implications, *Journal of Financial Economics* 119, 249–283.
- Heston, Steven, Kris Jacobs, and Hyung Joo Kim, 2025, Volatility risk and monotonic pricing kernels, Working paper, University of Maryland.
- Heston, Steven L., 1993, A closed-form solution for options with stochastic volatility with applications to bond and currency options, *The Review of Financial Studies* 6, 327–343.
- Heston, Steven L., Christopher S. Jones, Mehdi Khorram, Shuaiqi Li, and Haitao Mo, 2023, Option Momentum, *Journal of Finance* 78, 3141–3192.
- Heston, Steven L., Mark Loewenstein, and Gregory A. Willard, 2006, Options and bubbles, *Review of Financial Studies* 20, 359–390.
- Horenstein, Alex R., Aurelio Vasquez, and Xiao Xiao, 2025, Common Factors in Equity Option Returns, *Review of Financial Studies* (forthcoming).
- Hurn, A Stan, Kenneth A Lindsay, and Andrew J McClelland, 2015, Estimating the parameters of stochastic volatility models using option price data, *Journal of Business & Economic Statistics* 33, 579–594.
- Israelov, Roni, and Bryan T. Kelly, 2017, Forecasting the Distribution of Option Returns, *SSRN Electronic Journal* .
- Johnson, Travis L, 2017, Risk premia and the VIX term structure, *Journal of Financial and Quantitative Analysis* 52, 2461–2490.
- Jones, Christopher S., 2003, The Dynamics of Stochastic Volatility: Evidence from Underlying and Options Markets, *Journal of Econometrics* 116, 181–224.
- Jones, Christopher S., 2006, A Nonlinear Factor Analysis of S&P 500 Index Option Returns, *The Journal of Finance* 61, 2325–2363.

- Kelly, Bryan T., Seth Pruitt, and Yinan Su, 2018, Instrumented Principal Component Analysis, *SSRN Electronic Journal* .
- Nagel, Stefan, and Kenneth J. Singleton, 2011, Estimation and evaluation of conditional asset pricing models, *The Journal of Finance* 66, 873–909.
- Pan, Jun, 2002, The Jump-risk Premia Implicit in Options: Evidence from an Integrated Time-series Study, *Journal of Financial Economics* 63, 3–50.
- Singleton, Kenneth, 2006, *Empirical Dynamic Asset Pricing: Model Specification and Econometric Assessment* (Princeton University Press, Princeton, NJ).
- Todorov, Viktor, 2010, Variance Risk-Premium Dynamics: The Role of Jumps, *Review of Financial Studies* 23, 345–383.
- Wood, Simon N, 2003, Thin plate regression splines, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 65, 95–114.

Internet Appendix

to

Systematic Variance Risk Everywhere in Equity Option Markets

Alex Dickerson, Mathieu Fournier, Kris Jacobs, Piotr Orłowski

Appendix A. Derivation of Option Return Dynamics

This appendix provides a detailed derivation of the option return dynamics presented in the main text.

A.1. Model Primitives

Market dynamics. Consider an economy in which the market index $S_{M,t}$ and its stochastic variance $\sigma_{M,t}^2$ evolve under the physical measure P according to:

$$\frac{dS_{M,t}}{S_{M,t}} = \mu_t dt + \sigma_{M,t} dW_t^M \quad (\text{IA.1})$$

$$d\sigma_{M,t}^2 = \kappa_M(\theta_M - \sigma_{M,t}^2) dt + \sigma_{M,t} \delta_M dW_t^{\sigma_M}, \quad (\text{IA.2})$$

where μ_t is the expected market return, and $\sigma_{M,t}^2$ follows a square-root (CIR) process with mean-reversion speed κ_M , long-run mean θ_M , and volatility of variance δ_M . The Brownian motions dW_t^M and $dW_t^{\sigma_M}$ may be correlated.

ETF and stock dynamics. The economy features a cross-section of stocks and ETFs, denoted by $S_{l,t}$ for $l \neq M$, which follow a one-factor structure:

$$\frac{dS_{l,t}}{S_{l,t}} = (r^f - q_l + \alpha_l) dt + \beta_l \left(\frac{dS_{M,t}}{S_{M,t}} - r^f dt \right) + \sigma_{l,t} dW_t^l, \quad (\text{IA.3})$$

$$d\sigma_{l,t}^2 = \kappa_l(\theta_l - \sigma_{l,t}^2) dt + \sigma_{l,t} \delta_l dW_t^{\sigma_l}, \quad (\text{IA.4})$$

where r^f is the constant risk-free rate, q_l denotes the dividend yield, α_l captures any pricing error, dW_t^l represents idiosyncratic price risk, and $\sigma_{l,t}$ is the stochastic volatility of idiosyncratic returns. The idiosyncratic variance also follows a square-root process. The Brownian motions dW_t^l and $dW_t^{\sigma_l}$ may be correlated, and $dW_t^{\sigma_l}$ may depend on systematic risk factors or exhibit cross-sectional correlation.

A.2. Option Pricing via No-Arbitrage

Consider O_t^{ij} , the value of equity option i written on underlying security $S_{j,t}$ with $j \in \{M, l\}$. The option value depends on the state variables:

- For index options ($j = M$): $O_t^{ij} = O^{ij}(S_{M,t}, \sigma_{M,t}^2)$
- For stock/ETF options ($j = l$): $O_t^{ij} = O^{ij}(S_{l,t}, \sigma_{M,t}^2, \sigma_{l,t}^2)$

Itô's lemma. The change in the option price under P is given by Itô's lemma:

$$\begin{aligned}
dO_t^{ij} = & \frac{\partial O^{ij}}{\partial t} dt + \frac{\partial O^{ij}}{\partial S_j} dS_{j,t} + \frac{\partial O^{ij}}{\partial \sigma_M^2} d\sigma_{M,t}^2 + \mathbf{1}_{\{j=l\}} \frac{\partial O^{ij}}{\partial \sigma_l^2} d\sigma_{l,t}^2 \\
& + \frac{1}{2} \frac{\partial^2 O^{ij}}{\partial S_j^2} (dS_{j,t})^2 + \frac{1}{2} \frac{\partial^2 O^{ij}}{\partial (\sigma_M^2)^2} (d\sigma_{M,t}^2)^2 + \mathbf{1}_{\{j=l\}} \frac{1}{2} \frac{\partial^2 O^{ij}}{\partial (\sigma_l^2)^2} (d\sigma_{l,t}^2)^2 \\
& + \frac{\partial^2 O^{ij}}{\partial S_j \partial \sigma_M^2} dS_{j,t} d\sigma_{M,t}^2 + \mathbf{1}_{\{j=l\}} \frac{\partial^2 O^{ij}}{\partial S_j \partial \sigma_l^2} dS_{j,t} d\sigma_{l,t}^2 \\
& + \mathbf{1}_{\{j=l\}} \frac{\partial^2 O^{ij}}{\partial \sigma_M^2 \partial \sigma_l^2} d\sigma_{M,t}^2 d\sigma_{l,t}^2,
\end{aligned} \tag{IA.5}$$

where $\mathbf{1}_{\{j=l\}}$ is an indicator equal to one for stock/ETF options and zero for index options.

Risk-neutral valuation. Under no arbitrage, there exists a risk-neutral measure Q such that the option earns the risk-free rate in expectation:

$$E_t^Q[dO_t^{ij}] = r^f O_t^{ij} dt. \tag{IA.6}$$

Taking Q -conditional expectations of equation (IA.5), noting that all terms proportional to Brownian increments have zero Q -expectation, and applying condition (IA.6) yields:

$$\begin{aligned}
& \frac{\partial O^{ij}}{\partial t} dt + \frac{\partial O^{ij}}{\partial S_j} E_t^Q[dS_{j,t}] + \frac{\partial O^{ij}}{\partial \sigma_M^2} E_t^Q[d\sigma_{M,t}^2] + \mathbf{1}_{\{j=l\}} \frac{\partial O^{ij}}{\partial \sigma_l^2} E_t^Q[d\sigma_{l,t}^2] \\
& + \frac{1}{2} \frac{\partial^2 O^{ij}}{\partial S_j^2} E_t^Q[(dS_{j,t})^2] + \frac{1}{2} \frac{\partial^2 O^{ij}}{\partial (\sigma_M^2)^2} E_t^Q[(d\sigma_{M,t}^2)^2] + \mathbf{1}_{\{j=l\}} \frac{1}{2} \frac{\partial^2 O^{ij}}{\partial (\sigma_l^2)^2} E_t^Q[(d\sigma_{l,t}^2)^2] \\
& + \frac{\partial^2 O^{ij}}{\partial S_j \partial \sigma_M^2} E_t^Q[dS_{j,t} d\sigma_{M,t}^2] + \mathbf{1}_{\{j=l\}} \frac{\partial^2 O^{ij}}{\partial S_j \partial \sigma_l^2} E_t^Q[dS_{j,t} d\sigma_{l,t}^2] \\
& + \mathbf{1}_{\{j=l\}} \frac{\partial^2 O^{ij}}{\partial \sigma_M^2 \partial \sigma_l^2} E_t^Q[d\sigma_{M,t}^2 d\sigma_{l,t}^2] = r^f O_t^{ij} dt.
\end{aligned} \tag{IA.7}$$

Rearranging equation (IA.7) to isolate the drift term:

$$\begin{aligned}
\left(r^f O_t^{ij} - \frac{\partial O^{ij}}{\partial t} \right) dt &= \frac{\partial O^{ij}}{\partial S_j} E_t^Q[dS_{j,t}] + \frac{\partial O^{ij}}{\partial \sigma_M^2} E_t^Q[d\sigma_{M,t}^2] + \mathbf{1}_{\{j=l\}} \frac{\partial O^{ij}}{\partial \sigma_l^2} E_t^Q[d\sigma_{l,t}^2] \\
&+ \frac{1}{2} \frac{\partial^2 O^{ij}}{\partial S_j^2} E_t^Q[(dS_{j,t})^2] + \frac{1}{2} \frac{\partial^2 O^{ij}}{\partial (\sigma_M^2)^2} E_t^Q[(d\sigma_{M,t}^2)^2] \\
&+ \mathbf{1}_{\{j=l\}} \frac{1}{2} \frac{\partial^2 O^{ij}}{\partial (\sigma_l^2)^2} E_t^Q[(d\sigma_{l,t}^2)^2] \\
&+ \frac{\partial^2 O^{ij}}{\partial S_j \partial \sigma_M^2} E_t^Q[dS_{j,t} d\sigma_{M,t}^2] + \mathbf{1}_{\{j=l\}} \frac{\partial^2 O^{ij}}{\partial S_j \partial \sigma_l^2} E_t^Q[dS_{j,t} d\sigma_{l,t}^2] \\
&+ \mathbf{1}_{\{j=l\}} \frac{\partial^2 O^{ij}}{\partial \sigma_M^2 \partial \sigma_l^2} E_t^Q[d\sigma_{M,t}^2 d\sigma_{l,t}^2].
\end{aligned} \tag{IA.8}$$

This restriction implies that the drift term $\frac{\partial O^{ij}}{\partial t}$ and all second-order terms in equation (IA.5) can be re-expressed using risk-neutral expectations.

A.3. Option Return Dynamics

Substituting the restriction (IA.8) back into the Itô expansion (IA.5) eliminates the drift and second-order terms:

$$\begin{aligned}
dO_t^{ij} &= r^f O_t^{ij} dt + \frac{\partial O^{ij}}{\partial S_j} \left(dS_{j,t} - E_t^Q[dS_{j,t}] \right) \\
&+ \frac{\partial O^{ij}}{\partial \sigma_M^2} \left(d\sigma_{M,t}^2 - E_t^Q[d\sigma_{M,t}^2] \right) \\
&+ \mathbf{1}_{\{j=l\}} \frac{\partial O^{ij}}{\partial \sigma_l^2} \left(d\sigma_{l,t}^2 - E_t^Q[d\sigma_{l,t}^2] \right).
\end{aligned} \tag{IA.9}$$

Dividing both sides by O_t^{ij} yields the option return:

$$\begin{aligned}
\frac{dO_t^{ij}}{O_t^{ij}} &= r^f dt + \frac{1}{O_t^{ij}} \frac{\partial O^{ij}}{\partial S_j} \left(dS_{j,t} - E_t^Q[dS_{j,t}] \right) \\
&+ \frac{1}{O_t^{ij}} \frac{\partial O^{ij}}{\partial \sigma_M^2} \left(d\sigma_{M,t}^2 - E_t^Q[d\sigma_{M,t}^2] \right) \\
&+ \mathbf{1}_{\{j=l\}} \frac{1}{O_t^{ij}} \frac{\partial O^{ij}}{\partial \sigma_l^2} \left(d\sigma_{l,t}^2 - E_t^Q[d\sigma_{l,t}^2] \right).
\end{aligned} \tag{IA.10}$$

To express this in terms of the underlying return, note that:

$$dS_{j,t} - E_t^Q[dS_{j,t}] = S_{j,t} \left(\frac{dS_{j,t}}{S_{j,t}} - E_t^Q \left[\frac{dS_{j,t}}{S_{j,t}} \right] \right) = S_{j,t} \left(\frac{dS_{j,t}}{S_{j,t}} - r^f dt \right), \quad (\text{IA.11})$$

where the last equality uses that under Q , the expected return equals the risk-free rate.

Substituting this expression into equation (IA.10) yields:

$$\begin{aligned} \frac{dO_t^{ij}}{O_t^{ij}} &= r^f dt + \frac{S_{j,t}}{O_t^{ij}} \frac{\partial O^{ij}}{\partial S_j} \left(\frac{dS_{j,t}}{S_{j,t}} - r^f dt \right) \\ &\quad + \frac{1}{O_t^{ij}} \frac{\partial O^{ij}}{\partial \sigma_M^2} \left(d\sigma_{M,t}^2 - E_t^Q[d\sigma_{M,t}^2] \right) \\ &\quad + \mathbf{1}_{\{j=l\}} \frac{1}{O_t^{ij}} \frac{\partial O^{ij}}{\partial \sigma_l^2} \left(d\sigma_{l,t}^2 - E_t^Q[d\sigma_{l,t}^2] \right), \end{aligned} \quad (\text{IA.12})$$

where all terms are expressed as innovations relative to their Q -conditional expectations. This is the result presented in the main text.

A.4. Extension: Jump Risk in Option Returns

The baseline model can be extended to incorporate jump risk in the market portfolio. Suppose the market return follows:

$$\frac{dS_{M,t}}{S_{M,t}} = \mu_t dt + \sigma_{M,t} dW_t^M + J (dN_t - \lambda^P dt), \quad (\text{IA.13})$$

where dN_t is a Poisson process with intensity λ^P under P and J is the constant jump size. Individual securities inherit jump exposure through their market beta with $\beta_j = 1$ for $j = M$ and $\beta_j = \beta_l$ for $j = l$.

When jump risk is priced, the risk-neutral jump intensity differs from the physical intensity: $\lambda^Q = \lambda^P e^{-\xi_J}$, where ξ_J is the market price of jump risk. Let $\boldsymbol{\sigma}_t^2$ denote the vector of variance state variables, where $\boldsymbol{\sigma}_t^2 = (\sigma_{M,t}^2)$ for index options and $\boldsymbol{\sigma}_t^2 = (\sigma_{M,t}^2, \sigma_{l,t}^2)$ for stock/ETF options.

When a jump occurs, the underlying price changes instantaneously from $S_{j,t}$ to $S_{j,t}(1 + \beta_j J)$.

The resulting change in option value is given exactly by

$$O^{ij}(S_{j,t}(1 + \beta_j J), \boldsymbol{\sigma}_t^2) - O^{ij}(S_{j,t}, \boldsymbol{\sigma}_t^2). \quad (\text{IA.14})$$

Following the standard derivation and delevering, the delevered option excess return becomes:

$$\begin{aligned} \frac{O_t^{ij}}{S_{j,t}} \left(\frac{dO_t^{ij}}{O_t^{ij}} - r^f dt \right) &= \frac{\partial O^{ij}}{\partial S_j} \left(\frac{dS_{j,t}}{S_{j,t}} - r^f dt \right) + \frac{1}{S_{j,t}} \frac{\partial O^{ij}}{\partial \sigma_M^2} \left(d\sigma_{M,t}^2 - E_t^Q[d\sigma_{M,t}^2] \right) \\ &\quad + \mathbf{1}_{\{j=l\}} \frac{1}{S_{j,t}} \frac{\partial O^{ij}}{\partial \sigma_l^2} \left(d\sigma_{l,t}^2 - E_t^Q[d\sigma_{l,t}^2] \right) \\ &\quad + \frac{1}{S_{j,t}} [O^{ij}(S_{j,t}(1 + \beta_j J), \boldsymbol{\sigma}_t^2) - O^{ij}(S_{j,t}, \boldsymbol{\sigma}_t^2)] (dN_t - E_t^Q[dN_t]). \end{aligned} \quad (\text{IA.15})$$

where the last term correspond to the delevered jump loading $\frac{1}{S_{j,t}}[O^{ij}(S_{j,t}(1 + \beta_j J), \boldsymbol{\sigma}_t^2) - O^{ij}(S_{j,t}, \boldsymbol{\sigma}_t^2)]$. Under the physical measure P , the expected jump innovation is:

$$E_t^P[dN_t - E_t^Q[dN_t]] = (\lambda^P - \lambda^Q) dt = \lambda^P(1 - e^{-\xi_J}) dt, \quad (\text{IA.16})$$

where ξ_J is the market price of jump risk. This jump risk premium is positive when $\xi_J > 0$.

A.5. Derivation of the Discrete-Time Factor Model

This appendix provides a detailed derivation of the discrete-time factor representation in equation (11) of the main text, starting from the continuous-time delevered option return in equation (7).

State variables. The option value at time t depends on the underlying price $S_{j,t}$, the systematic variance $\sigma_{M,t}^2$, the idiosyncratic variance $\sigma_{l,t}^2$ (for stock/ETF options), and calendar time t :

$$O_t^{ij} = O^{ij}(S_{j,t}, \sigma_{M,t}^2, \sigma_{l,t}^2, t). \quad (\text{IA.17})$$

For index options ($j = M$), the dependence on $\sigma_{l,t}^2$ is absent.

Discrete delevered excess return. Define the discrete delevered excess return on option i written on underlying j over the interval $[t, t + 1]$ as

$$r_{ij,t+1} \equiv \frac{O_{t+1}^{ij} - O_t^{ij} (1 + r^f)}{S_{j,t}}. \quad (\text{IA.18})$$

Notation. We denote discrete changes by $\tilde{\Delta}X_{t+1} \equiv X_{t+1} - X_t$ and the underlying excess return by

$$r_{Sj,t+1} \equiv \frac{\tilde{\Delta}S_{j,t+1}}{S_{j,t}} - r^f = \frac{S_{j,t+1} - S_{j,t}}{S_{j,t}} - r^f. \quad (\text{IA.19})$$

A.5.1. Step 1: Second-Order Taylor Expansion

Expand $O_{t+1}^{ij} = O^{ij}(S_{j,t+1}, \sigma_{M,t+1}^2, \sigma_{l,t+1}^2, t + 1)$ around the time- t state $(S_{j,t}, \sigma_{M,t}^2, \sigma_{l,t}^2, t)$ to second order. All partial derivatives are evaluated at time t .

First-order terms:

$$O_{t+1}^{ij} \approx O_t^{ij} + \frac{\partial O^{ij}}{\partial t} \cdot 1 + \frac{\partial O^{ij}}{\partial S_j} \tilde{\Delta}S_{j,t+1} + \frac{\partial O^{ij}}{\partial \sigma_M^2} \tilde{\Delta}\sigma_{M,t+1}^2 + \frac{\partial O^{ij}}{\partial \sigma_l^2} \tilde{\Delta}\sigma_{l,t+1}^2 \quad (\text{IA.20})$$

Second-order terms:

$$\begin{aligned} & + \frac{1}{2} \frac{\partial^2 O^{ij}}{\partial S_j^2} (\tilde{\Delta}S_{j,t+1})^2 + \frac{1}{2} \frac{\partial^2 O^{ij}}{\partial (\sigma_M^2)^2} (\tilde{\Delta}\sigma_{M,t+1}^2)^2 + \frac{1}{2} \frac{\partial^2 O^{ij}}{\partial (\sigma_l^2)^2} (\tilde{\Delta}\sigma_{l,t+1}^2)^2 \\ & + \frac{\partial^2 O^{ij}}{\partial S_j \partial \sigma_M^2} \tilde{\Delta}S_{j,t+1} \tilde{\Delta}\sigma_{M,t+1}^2 + \frac{\partial^2 O^{ij}}{\partial S_j \partial \sigma_l^2} \tilde{\Delta}S_{j,t+1} \tilde{\Delta}\sigma_{l,t+1}^2 \\ & + \frac{\partial^2 O^{ij}}{\partial \sigma_M^2 \partial \sigma_l^2} \tilde{\Delta}\sigma_{M,t+1}^2 \tilde{\Delta}\sigma_{l,t+1}^2. \end{aligned} \quad (\text{IA.21})$$

Combining (IA.20) and (IA.21):

$$\begin{aligned} O_{t+1}^{ij} \approx O_t^{ij} & + \frac{\partial O^{ij}}{\partial t} + \frac{\partial O^{ij}}{\partial S_j} \tilde{\Delta}S_{j,t+1} + \frac{\partial O^{ij}}{\partial \sigma_M^2} \tilde{\Delta}\sigma_{M,t+1}^2 + \frac{\partial O^{ij}}{\partial \sigma_l^2} \tilde{\Delta}\sigma_{l,t+1}^2 \\ & + \frac{1}{2} \frac{\partial^2 O^{ij}}{\partial S_j^2} (\tilde{\Delta}S_{j,t+1})^2 + \text{H.O.T.}, \end{aligned} \quad (\text{IA.22})$$

where H.O.T. denotes higher-order terms (the remaining second-order cross-partials and variance-of-variance terms).³¹

A.5.2. Step 2: Form the Delevered Excess Return

Substituting (IA.22) into the definition (IA.18):

$$\begin{aligned}
r_{ij,t+1} &= \frac{O_{t+1}^{ij} - O_t^{ij}(1 + r^f)}{S_{j,t}} \\
&\approx \frac{1}{S_{j,t}} \left[\frac{\partial O^{ij}}{\partial t} + \frac{\partial O^{ij}}{\partial S_j} \tilde{\Delta} S_{j,t+1} + \frac{\partial O^{ij}}{\partial \sigma_M^2} \tilde{\Delta} \sigma_{M,t+1}^2 + \frac{\partial O^{ij}}{\partial \sigma_l^2} \tilde{\Delta} \sigma_{l,t+1}^2 \right. \\
&\quad \left. + \frac{1}{2} \frac{\partial^2 O^{ij}}{\partial S_j^2} (\tilde{\Delta} S_{j,t+1})^2 + \text{H.O.T.} - r^f O_t^{ij} \right]. \tag{IA.23}
\end{aligned}$$

A.5.3. Step 3: Decompose into Risk-Neutral Expectations and Innovations

The key step is to decompose each stochastic term into its Q -conditional expectation plus an innovation. Write:

$$\tilde{\Delta} S_{j,t+1} = E_t^Q[\tilde{\Delta} S_{j,t+1}] + \left(\tilde{\Delta} S_{j,t+1} - E_t^Q[\tilde{\Delta} S_{j,t+1}] \right) = r^f S_{j,t} + S_{j,t} r_{s_{j,t+1}}, \tag{IA.24}$$

$$\tilde{\Delta} \sigma_{M,t+1}^2 = E_t^Q[\tilde{\Delta} \sigma_{M,t+1}^2] + \left(\tilde{\Delta} \sigma_{M,t+1}^2 - E_t^Q[\tilde{\Delta} \sigma_{M,t+1}^2] \right), \tag{IA.25}$$

$$\tilde{\Delta} \sigma_{l,t+1}^2 = E_t^Q[\tilde{\Delta} \sigma_{l,t+1}^2] + \left(\tilde{\Delta} \sigma_{l,t+1}^2 - E_t^Q[\tilde{\Delta} \sigma_{l,t+1}^2] \right), \tag{IA.26}$$

$$(\tilde{\Delta} S_{j,t+1})^2 = E_t^Q[(\tilde{\Delta} S_{j,t+1})^2] + \left((\tilde{\Delta} S_{j,t+1})^2 - E_t^Q[(\tilde{\Delta} S_{j,t+1})^2] \right), \tag{IA.27}$$

$$\text{H.O.T.} = \text{H.O.T.}^Q + (\text{H.O.T.} - \text{H.O.T.}^Q), \tag{IA.28}$$

where H.O.T.^Q denotes the Q -expectations of the higher-order terms. Equation (IA.24) uses the fact that under Q , the expected return on the underlying equals the risk-free rate: $E_t^Q[\tilde{\Delta} S_{j,t+1}/S_{j,t}] = r^f$.

³¹The dominant term among these is $\frac{\partial^2 O^{ij}}{\partial S_j \partial \sigma_M^2} \tilde{\Delta} S_{j,t+1} \tilde{\Delta} \sigma_{M,t+1}^2$, whose unconditional mean is small relative to the leading-order terms. The terms involving $(\tilde{\Delta} \sigma_{M,t+1}^2)^2$ and $(\tilde{\Delta} \sigma_{l,t+1}^2)^2$ are of order $O((\tilde{\Delta} t)^2)$ in the diffusive setting and are negligible at weekly frequency.

A.5.4. Step 4: Substitute and Group Terms

Substituting the decompositions (IA.24)–(IA.28) into (IA.23) and grouping:

$$\begin{aligned}
r_{ij,t+1} &\approx \underbrace{\left[\frac{\partial O^{ij}}{\partial t} + \frac{\partial O^{ij}}{\partial S_j} r^f S_{j,t} + \frac{1}{2} \frac{\partial^2 O^{ij}}{\partial S_j^2} E_t^Q [(\tilde{\Delta} S_j)^2] + \frac{\partial O^{ij}}{\partial \sigma_M^2} E_t^Q [\tilde{\Delta} \sigma_M^2] + \frac{\partial O^{ij}}{\partial \sigma_l^2} E_t^Q [\tilde{\Delta} \sigma_l^2] + \text{H.O.T.}^Q - r^f O_t^{ij} \right]}_{S_{j,t}} \\
&\equiv A \\
&+ \frac{\partial O^{ij}}{\partial S_j} r s_{j,t+1} \\
&+ \frac{1}{S_{j,t}} \frac{\partial O^{ij}}{\partial \sigma_M^2} \left(\tilde{\Delta} \sigma_{M,t+1}^2 - E_t^Q [\tilde{\Delta} \sigma_{M,t+1}^2] \right) \\
&+ \frac{1}{2} S_{j,t} \frac{\partial^2 O^{ij}}{\partial S_j^2} \left(\left(\frac{\tilde{\Delta} S_{j,t+1}}{S_{j,t}} \right)^2 - E_t^Q \left[\left(\frac{\tilde{\Delta} S_{j,t+1}}{S_{j,t}} \right)^2 \right] \right) \\
&+ \frac{1}{S_{j,t}} \frac{\partial O^{ij}}{\partial \sigma_l^2} \left(\tilde{\Delta} \sigma_{l,t+1}^2 - E_t^Q [\tilde{\Delta} \sigma_{l,t+1}^2] \right) + \frac{1}{S_{j,t}} (\text{H.O.T.} - \text{H.O.T.}^Q), \tag{IA.29}
\end{aligned}$$

where in the gamma line we have divided and multiplied by $S_{j,t}^2$ to express the quadratic variation in return form.

A.6. Step 5: Apply the No-Arbitrage Condition

Risk-neutral pricing restriction. Under no arbitrage, the option price must satisfy

$$E_t^Q[O_{t+1}^{ij}] = O_t^{ij} (1 + r^f). \tag{IA.30}$$

Taking the Q -expectation of the Taylor expansion (IA.22):

$$\begin{aligned}
E_t^Q[O_{t+1}^{ij}] &\approx O_t^{ij} + \frac{\partial O^{ij}}{\partial t} + \frac{\partial O^{ij}}{\partial S_j} E_t^Q[\tilde{\Delta} S_{j,t+1}] + \frac{\partial O^{ij}}{\partial \sigma_M^2} E_t^Q[\tilde{\Delta} \sigma_{M,t+1}^2] + \frac{\partial O^{ij}}{\partial \sigma_l^2} E_t^Q[\tilde{\Delta} \sigma_{l,t+1}^2] \\
&+ \frac{1}{2} \frac{\partial^2 O^{ij}}{\partial S_j^2} E_t^Q[(\tilde{\Delta} S_{j,t+1})^2] + \text{H.O.T.}^Q. \tag{IA.31}
\end{aligned}$$

Combining (IA.30) and (IA.31):

$$\begin{aligned}
& \frac{\partial O^{ij}}{\partial t} + \frac{\partial O^{ij}}{\partial S_j} r^f S_{j,t} \\
& + \frac{1}{2} \frac{\partial^2 O^{ij}}{\partial S_j^2} E_t^Q [(\tilde{\Delta} S_{j,t+1})^2] + \frac{\partial O^{ij}}{\partial \sigma_M^2} E_t^Q [\tilde{\Delta} \sigma_{M,t+1}^2] \\
& + \frac{\partial O^{ij}}{\partial \sigma_l^2} E_t^Q [\tilde{\Delta} \sigma_{l,t+1}^2] + \text{H.O.T.}^Q - r^f O_t^{ij} = 0.
\end{aligned} \tag{IA.32}$$

This is the discrete-time analogue of the risk-neutral pricing PDE (or, equivalently, of the drift restriction). It implies that the bracket A in equation (IA.29) equals zero:

$$A = 0. \tag{IA.33}$$

A.6.1. Step 6: Express in Terms of Return Innovations

With $A = 0$, equation (IA.29) reduces to:

$$\begin{aligned}
r_{ij,t+1} & \approx \frac{\partial O^{ij}}{\partial S_j} r s_{j,t+1} + \frac{1}{S_{j,t}} \frac{\partial O^{ij}}{\partial \sigma_M^2} \left(\tilde{\Delta} \sigma_{M,t+1}^2 - E_t^Q [\tilde{\Delta} \sigma_{M,t+1}^2] \right) \\
& + \frac{1}{2} S_{j,t} \frac{\partial^2 O^{ij}}{\partial S_j^2} \left(r s_{j,t+1}^2 - E_t^Q [r s_{j,t+1}^2] \right) + \epsilon_{ij,t+1},
\end{aligned} \tag{IA.34}$$

where we have used the approximation

$$\left(\frac{\tilde{\Delta} S_{j,t+1}}{S_{j,t}} \right)^2 = (r s_{j,t+1} + r^f)^2 \approx r s_{j,t+1}^2, \tag{IA.35}$$

since r^f is small relative to $r s_{j,t+1}$ at weekly frequency (the cross-term $2 r^f r s_{j,t+1}$ and $(r^f)^2$ are negligible).³² The residual is

$$\epsilon_{ij,t+1} = \frac{1}{S_{j,t}} \frac{\partial O^{ij}}{\partial \sigma_l^2} \left(\tilde{\Delta} \sigma_{l,t+1}^2 - E_t^Q [\tilde{\Delta} \sigma_{l,t+1}^2] \right) + \frac{1}{S_{j,t}} (\text{H.O.T.} - \text{H.O.T.}^Q), \tag{IA.36}$$

³²For a weekly risk-free rate of approximately 0.1%, the approximation error is of order $10^{-3} r s_{j,t+1}$, which is negligible relative to the leading term $r s_{j,t+1}^2$.

which collects the idiosyncratic variance innovation and higher-order remainder terms.

On the idiosyncratic variance innovation. In the absence of systematic drivers of idiosyncratic variance (i.e., $dW_t^{\sigma_l}$ is purely idiosyncratic), the change of measure from P to Q does not affect the drift of $\sigma_{l,t}^2$, so that $E_t^Q[\tilde{\Delta}\sigma_{l,t+1}^2] = E_t^P[\tilde{\Delta}\sigma_{l,t+1}^2]$. In this case the idiosyncratic variance innovation has zero P -expectation and is unpriced.³³

A.6.2. Step 7: Absorb Risk-Neutral Expectations into the Intercept

Equation (IA.34) can be rewritten as:

$$\begin{aligned} r_{ij,t+1} = & -\mathcal{V}_{ij,t} \cdot E_t^Q[f_{v,t+1}] - \Gamma_{ij,t} \cdot E_t^Q[rs_{j,t+1}^2] \\ & + \Delta_{ij,t} \cdot rs_{j,t+1} + \mathcal{V}_{ij,t} \cdot f_{v,t+1} + \Gamma_{ij,t} \cdot rs_{j,t+1}^2 + \epsilon_{ij,t+1}. \end{aligned} \quad (\text{IA.37})$$

where the systematic variance factor is

$$f_{v,t+1} \equiv \Delta\sigma_{M,t+1}^2. \quad (\text{IA.38})$$

Derivation of (IA.37). Starting from (IA.34), expand the innovation terms:

$$\begin{aligned} r_{ij,t+1} = & \Delta_{ij,t} \cdot rs_{j,t+1} + \mathcal{V}_{ij,t} \left(f_{v,t+1} - E_t^Q[f_{v,t+1}] \right) + \Gamma_{ij,t} \left(rs_{j,t+1}^2 - E_t^Q[rs_{j,t+1}^2] \right) + \epsilon_{ij,t+1} \\ = & \Delta_{ij,t} \cdot rs_{j,t+1} + \mathcal{V}_{ij,t} \cdot f_{v,t+1} - \mathcal{V}_{ij,t} \cdot E_t^Q[f_{v,t+1}] + \Gamma_{ij,t} \cdot rs_{j,t+1}^2 - \Gamma_{ij,t} \cdot E_t^Q[rs_{j,t+1}^2] + \epsilon_{ij,t+1}. \end{aligned} \quad (\text{IA.39})$$

Defining the conditional intercept as

$$\gamma_{ij,t} \equiv -\mathcal{V}_{ij,t} \cdot E_t^Q[f_{v,t+1}] - \Gamma_{ij,t} \cdot E_t^Q[rs_{j,t+1}^2], \quad (\text{IA.40})$$

³³If a systematic component were present in $dW_t^{\sigma_l}$, the residual $\epsilon_{ij,t+1}$ would acquire a nonzero risk-neutral expectation, and an additional priced factor would be required. Such a factor would, however, be partially spanned by $rs_{j,t+1}^2$, since total realized variance of security j reflects both systematic and idiosyncratic variance.

we arrive at the factor model:

$$r_{ij,t+1} = \gamma_{ij,t} + \Delta_{ij,t} \cdot rs_{j,t+1} + \mathcal{V}_{ij,t} \cdot f_{v,t+1} + \Gamma_{ij,t} \cdot rs_{j,t+1}^2 + \epsilon_{ij,t+1}. \quad (\text{IA.41})$$

Gamma factor in the presence of jumps and idiosyncratic variances. The squared excess return $rs_{j,t+1}^2$ in equation (IA.34) emerges from the pure diffusion setting, but it also captures the leading-order effect of systematic jumps when these are present. To see this, consider the exact jump revaluation from Section A:

$$\frac{1}{S_{j,t}} [O^{ij}(S_{j,t}(1 + \beta_j J), \boldsymbol{\sigma}_t^2) - O^{ij}(S_{j,t}, \boldsymbol{\sigma}_t^2)]. \quad (\text{IA.42})$$

Applying a second-order Taylor expansion around $S_{j,t}$:

$$\begin{aligned} \frac{1}{S_{j,t}} [O^{ij}(S_{j,t}(1 + \beta_j J), \boldsymbol{\sigma}_t^2) - O^{ij}(S_{j,t}, \boldsymbol{\sigma}_t^2)] &\approx \frac{1}{S_{j,t}} \left[\frac{\partial O^{ij}}{\partial S_j} S_{j,t} \beta_j J + \frac{1}{2} \frac{\partial^2 O^{ij}}{\partial S_j^2} (S_{j,t} \beta_j J)^2 \right] \\ &= \underbrace{\beta_j J \frac{\partial O^{ij}}{\partial S_j}}_{\text{delta}} + \underbrace{\frac{1}{2} S_{j,t} (\beta_j J)^2 \frac{\partial^2 O^{ij}}{\partial S_j^2}}_{\text{gamma}} \\ &= \Delta_{ij,t} \beta_j J + \Gamma_{ij,t} (\beta_j J)^2, \end{aligned} \quad (\text{IA.43})$$

where $\Delta_{ij,t}$ and $\Gamma_{ij,t}$ are the delta and gamma exposures. The delta component is absorbed into the first factor $rs_{j,t+1}$, since jumps contribute to the total underlying return. The gamma component, however, loads on $(\beta_j J)^2 (dN_t - E_t^Q[dN_t])$ —a squared price movement driven by jump arrivals.

In discrete time, jumps produce large realizations of $rs_{j,t+1}^2$: when a jump of size $\beta_j J$ occurs during $[t, t+1]$, the squared return is dominated by $(\beta_j J)^2$, which deviates substantially from $E_t^Q[rs_{j,t+1}^2]$. The factor $\Gamma_{ij,t} \cdot rs_{j,t+1}^2$ in the discrete model (IA.41) therefore captures this jump-induced convexity exposure alongside the discretization effect from the pure diffusion setting.

This observation clarifies the dual role of $rs_{j,t+1}^2$ in the factor model. First, it captures the option's convexity (gamma) exposure to large price movements, whether these arise from the accumulation of diffusive shocks over a discrete interval or from jumps. Second,

since $rs_{j,t+1}^2$ reflects the total realized variance of security j —including both systematic and idiosyncratic components—it controls for underlying-specific variance innovations that would otherwise contaminate inference on the systematic variance risk premium $f_{v,t+1}$. In this sense, the squared return provides a model-free proxy for the exact but model-dependent jump revaluation in (IA.42), while simultaneously absorbing idiosyncratic variance risk that is not spanned by the systematic factor.

Appendix B. Estimation and Inference

Following Fournier, Jacobs, and Orłowski (2024) (FJO, hereafter), we model the $\mathcal{V}_t^{ij}$, Δ_t^{ij} , and Γ_t^{ij} as nonparametric functions of option characteristics (i.e., time to maturity, moneyness, and log of implied volatility). This Section is an abridged version of their Internet Appendix C with notation specialized to our modeling choices. $\mathcal{V}_t^{ij}$ is the systematic variance factor exposure. Δ_t^{ij} , and Γ_t^{ij} are the delta and gamma exposures to the underlying-specific excess return and its square. We have:

$$r_{ij,t+1} = \gamma_t^{ij} + \mathcal{V}_t^{ij} \cdot f_{v,t+1} + \Delta_t^{ij} \cdot rs_{t+1}^j + \Gamma_t^{ij} \cdot (rs_{t+1}^j)^2 + \varepsilon_{ij,t+1}. \quad (\text{IA.1})$$

Setting $\gamma_t^{ij} = -\mathcal{V}_t^{ij} E_t^Q[f_{v,t+1}] - \Gamma_t^{ij} E_t^Q[(rs_{t+1}^j)^2]$ in equation (IA.1) imposes the no-arbitrage condition in the model. This restriction is obtained by taking the risk-neutral expectation of (IA.1) using the fact that: (i) $E_t^Q[r_{ij,t+1}] = E_t^Q[rs_{t+1}^j] = 0$ as both $r_{ij,t+1}$ and rs_{t+1}^j are expressed in excess returns and are tradable assets, and (ii) $\varepsilon_{ij,t+1}$ in equation (IA.1) is assumed not to be priced and therefore satisfies $E_t^Q[\varepsilon_{ij,t+1}] = 0$.³⁴

We make the following assumptions which will allow us to estimate the model via two-stage linear regression:

1. The risk-neutral and physical expectations of the variance factor f_v are affine in VIX_t^2 :
 $\mathbf{g}_{v,t} = [1 \quad VIX_t^2]'$ and $E_t^Q[f_{v,t+1}] = \mathbf{g}_{v,t}' \boldsymbol{\mu}_v^Q$;
2. The risk-neutral and physical expectations of $(rs_{t+1}^j)^2$ are the same across all assets:
 $E_t^Q[rs_{t+1}^j] = g_\Gamma \mu_\Gamma^Q$ with, implicitly, $g_\Gamma \equiv 1$; This assumption restricts the potential risk

³⁴If $Cov_t\left(\frac{\xi_{t+1}}{\xi_t} \left(1 + r_t^f\right); \varepsilon_{ij,t+1}\right) = 0$ (i.e., $\varepsilon_{ij,t+1}$ is not priced) then we have $E_t^Q[\varepsilon_{ij,t+1}] = E_t^P[\varepsilon_{ij,t+1}] = 0$ in the absence of arbitrage opportunities.

premium on $(rs_{t+1}^j)^2$ to be equal across all stocks and we make it because in unreported results we consistently obtained small and insignificant estimates of this premium in line with FJO;

3. We model exposures using basis expansions of a vector of option characteristics; To simplify notation, we use the same characteristics for each factor; Given the characteristic vector $\mathbf{c}_{ij,t}$ we obtain its Thin Plate Regression Spline basis (Wood, 2003) expansion $\phi^j(\mathbf{c}_{ij,t}) =: \boldsymbol{\phi}_{ij,t}$ and hence exposure to factor l is linear in asset j 's parameter vector $\mathbf{b}_{j,l}$: $\boldsymbol{\phi}_{ij,t}' \mathbf{b}_{j,l}$; We construct the bases separately for each underlying asset with the procedure described in detail in Internet Appendix B to FJO.

B.1. Option Exposures

We rewrite equation (IA.1) as linear in appropriately redefined parameters and explanatory variables. Using our assumptions, we first obtain

$$\gamma_t^{ij} = \boldsymbol{\phi}_{ij,t}' \mathbf{a}_{j,0}^Q + VIX_t^2 \boldsymbol{\phi}_{ij,t}' \mathbf{a}_{j,1}^Q \quad (\text{IA.2})$$

with $\mathbf{a}_{j,0}^Q = -\mathbf{b}_{j,v} \mu_{v,0}^Q - \mathbf{b}_{j,\Gamma} \mu_{\Gamma}^Q$ and $\mathbf{a}_{j,1}^Q = -\mathbf{b}_{j,v} \mu_{v,1}^Q$. We further define $\mathbf{G}_{ij,t} \equiv [\boldsymbol{\phi}_{ij,t}' \quad VIX_t^2 \boldsymbol{\phi}_{ij,t}']'$ and $\mathbf{a}_j \equiv [\mathbf{a}_{j,0}^{Q'} \quad \mathbf{a}_{j,1}^{Q'}]'$ to obtain

$$\gamma_t^{ij} = \mathbf{G}_{ij,t}' \mathbf{a}_j. \quad (\text{IA.3})$$

Next, we replace factor exposures with their basis expansions using assumption 3, and we rewrite the product of each basis expansion and factor realization using the following notation

$$\mathbf{F}_{ij,l,t+1} \equiv f_{l,t+1} \boldsymbol{\phi}_{ij,t}, \quad (\text{IA.4})$$

with $l \in \{\Delta, \Gamma, \mathcal{V}\}$. Equation (IA.1) then becomes linear in the extended parameter vector $[\mathbf{a}_j' \quad \mathbf{b}_{j,\mathcal{V}}' \quad \mathbf{b}_{j,\Delta}' \quad \mathbf{b}_{j,\Gamma}']'$:

$$r_{ij,t+1} = \mathbf{G}_{ij,t}' \mathbf{a}_j + \mathbf{F}_{ij,\mathcal{V},t+1}' \mathbf{b}_{j,\mathcal{V}} + \mathbf{F}_{ij,\Delta,t+1}' \mathbf{b}_{j,\Delta} + \mathbf{F}_{ij,\Gamma,t+1}' \mathbf{b}_{j,\Gamma} + \varepsilon_{ij,t+1}. \quad (\text{IA.5})$$

The extended parameter vector $[\mathbf{a}_j' \quad \mathbf{b}_{j,\mathcal{V}}' \quad \mathbf{b}_{j,\Delta}' \quad \mathbf{b}_{j,\Gamma}']'$ is estimated by OLS for each underlying separately. Note that the risk-neutral expectation parameters $\mu_{\mathcal{V}}^Q$ and μ_{Γ}^Q are not identified. The estimated exposures are given by the products of the basis expansion of

option characteristics with the estimated parameter vectors; for example, an option's vega is calculated as:

$$\widehat{\mathcal{V}}_{ij,t} = \phi'_{ij,t} \widehat{\mathbf{b}}_{j,\nu}. \quad (\text{IA.6})$$

B.2. Risk Premiums

Using the first-stage estimates, we define the projection of the delta-hedged option return on the conditional intercept and realizations of $f_{v,t+1}$ and rs_{t+1}^j :

$$rp_{ij,t+1}^\Delta \equiv E[r_{ij,t+1}^\Delta | rs_{j,t+1}, f_{v,t+1}, \mathbf{c}_{ij,t}] = \mathbf{G}'_{ij,t} \mathbf{a}_j + \mathbf{F}'_{ij,\nu,t+1} \mathbf{b}_{j,\nu} + \mathbf{F}'_{ij,\Gamma,t+1} \mathbf{b}_{j,\Gamma} \quad (\text{IA.7})$$

$$= -E_t^Q[f_{v,t+1}] \mathcal{V}_t^{ij} - E_t^Q[(rs_{t+1}^j)^2] \Gamma_t^{ij} + \mathcal{V}_t^{ij} f_{v,t+1} + \Gamma_t^{ij} (rs_{t+1}^j)^2. \quad (\text{IA.8})$$

Taking the time- t P -expectation of this projection leads to an expression for the option's conditional delta-hedged risk premium,

$$E_t^P[r_{ij,t+1}^\Delta] = (E_t^P - E_t^Q)[f_{v,t+1}] \cdot \mathcal{V}_t^{ij} + (E_t^P - E_t^Q)[(rs_{t+1}^j)^2] \cdot \Gamma_t^{ij}. \quad (\text{IA.9})$$

Following Assumptions 1–2, we take the physical expectations of $f_{v,t+1}$ and (rs_{t+1}^j) to be, respectively, affine in VIX_t^2 and constant. This allows us to express the projection (IA.7) as

$$rp_{ij,t+1}^\Delta = [\mathcal{V}_t^{ij} \quad \mathcal{V}_t^{ij} \cdot VIX_t^2] \boldsymbol{\lambda}_v + \Gamma_t^{ij} \lambda_\Gamma + \nu_{ij,t+1} \equiv \mathbf{d}'_{ij,t} \boldsymbol{\lambda} + \nu_{ij,t+1}. \quad (\text{IA.10})$$

where we introduced the notation

$$\mathbf{d}_{ij,t} \equiv [\mathcal{V}_t^{ij} \quad \mathcal{V}_t^{ij} \cdot VIX_t^2 \quad \Gamma_t^{ij}]' \quad \boldsymbol{\lambda} = [\lambda_{v,0} \quad \lambda_{v,1} \quad \lambda_\Gamma]'. \quad (\text{IA.11})$$

In what follows, $\widehat{rp}^\Delta$ and $\widehat{\mathbf{d}}$ denote the quantities obtained using first-stage estimates of exposures and the conditional intercept parameters. It is also useful to stack the vectors of explanatory variables and parameters in equation (IA.7):

$$\mathbf{H}_{ij,t+1} \equiv [\mathbf{G}'_{ij,t} \quad \mathbf{F}'_{ij,\nu,t+1} \quad \mathbf{F}'_{ij,\Gamma,t+1}]' \quad \mathbf{B}_j \equiv [\mathbf{a}'_j \quad \mathbf{b}'_{j,\nu} \quad \mathbf{b}'_{j,\Gamma}]' \quad \mathbf{b}_j \equiv [\mathbf{b}'_{j,\nu} \quad \mathbf{b}'_{j,\Gamma}]'. \quad (\text{IA.12})$$

Using this notation we can define the unbiased estimator of the vector of risk premiums:

$$\hat{\lambda} = \hat{\lambda}_{OLS} - \mathbf{Q}_{\hat{\mathbf{d}}}^{-1} \frac{1}{N} \sum_{t,j,i} \hat{\mathbf{d}}_{ij,t} \hat{\Sigma}_{j,\hat{\mathbf{b}}_j,\hat{\mathbf{B}}_j} \mathbf{H}_{ij,t+1}. \quad (\text{IA.13})$$

In equation (IA.13), $\hat{\lambda}$ is the OLS estimator of the risk premium parameter vector in equation (IA.10), which is biased due to the presence of first-stage estimation error in the exposures, which is correlated with the estimation error in the return projection on nontraded factors. $\mathbf{Q}_{\hat{\mathbf{d}}} \equiv \frac{1}{N} \sum_{t,j,i} \hat{\mathbf{d}}_{ij,t} \hat{\mathbf{d}}'_{ij,t}$. The debiasing term arises from Corollary C.1 and Proposition C.6 in Internet Appendix C to FJO. Furthermore, the total number of observations in our sample is $N = \sum_t \sum_j N_{jt}$, with N_{jt} the number of option returns observed for underlying j on day t ; and $\hat{\Sigma}_{j,\hat{\mathbf{b}}_j,\hat{\mathbf{B}}_j}$ is the sub-matrix of the variance-covariance matrix of the first-stage parameters of underlying asset j given in Proposition C.3 therein. In our implementation of the estimator, we additionally weight each observation by the inverse of the total number of options on all assets in the sample available on day t : $w_t = 1/N_t$, $N_t = \sum_j N_{jt}$.

Propositions C.8–C.9 in Internet Appendix C to FJO give the asymptotic variance-covariance matrix of $\hat{\lambda}$ and its feasible estimator. The propositions can be applied directly, replacing T therein with the total number of observations N , and η_{t+1}^Ω with $\nu_{ij,t+1}$. This inference scheme assumes that first-stage estimation errors are independent across underlying assets, but allows for arbitrary correlation of residuals in regression (IA.10) within and across underlying assets and time.

The theoretical results in the Internet Appendix to FJO also make use of a matrix $\mathbf{D}_{ij,t}$ such that

$$\mathbf{d}_{ij,t} = \mathbf{D}_{ij,t} \mathbf{b}_j \text{ and } \hat{\mathbf{d}}_{ij,t} = \mathbf{D}_{ij,t} \hat{\mathbf{b}}_j. \quad (\text{IA.14})$$

In our case, it is explicitly given as ($\otimes$ denotes the Kronecker product)

$$\mathbf{D}_{ij,t} \equiv \begin{bmatrix} \boldsymbol{\phi}'_{ij,t} \otimes \begin{bmatrix} 1 \\ \text{VIX}_t^2 \end{bmatrix} & 0 \\ 0 & \boldsymbol{\phi}'_{ij,t} \end{bmatrix}. \quad (\text{IA.15})$$

Appendix C.

Omitted Variable Bias in Premium Estimates from Black-Scholes Hedged Option Returns

We examine the consequences of using Black-Merton-Scholes deltas in first-stage estimation on our estimates of risk premium parameters. We start with a simplified setting with constant exposures and risk premiums, in which option returns are driven by the excess underlying return $rs_{j,t+1}$ (assuming the asset's beta is equal to one and its risk premium is equal to the equity market risk premium) and variance shock $f_{v,t+1}$. We further assume only one derivative per underlying asset. Let $r_{j,t+1}$ denote the BS delta-hedged excess option return.

If the BS delta-hedge is incorrect, then the option return has some residual exposure to the underlying. Therefore, estimating a single-factor model (with $f_{v,t+1}$ only) leads to a biased estimate of $\mathcal{V}$ and the risk premium for $f_{v,t+1}$.

The true asset dynamics is

$$r_{j,t+1} = \gamma_j + \mathcal{V}_j \cdot f_{v,t+1} + \Delta_j \cdot rs_{j,t+1} + \varepsilon_{j,t+1},$$

where the no-arbitrage restriction implies

$$\gamma_j = -\mathcal{V}_j \mu_v^Q$$

under the correct specification.

Consider the alternative model which assumes that the BS delta-hedge is correct and drops $rs_{j,t+1}$:

$$r_{j,t+1} = \tilde{\gamma}_j + \tilde{\mathcal{V}}_j f_{v,t+1} + \eta_{j,t+1}.$$

We can derive the limiting value of the estimators $\hat{\gamma}$ and $\hat{\mathcal{V}}$ in this regression in terms of the parameters of the true model:

$$\text{plim } \hat{\gamma}_j = \gamma_j - \Delta_j (\mu_m^P - \rho_j \mu_v^P),$$

$$\text{plim } \hat{\mathcal{V}}_j = \mathcal{V}_j + \rho_j \Delta_j,$$

where ρ_j is the coefficient in the regression

$$rs_{j,t+1} = \delta_j + \rho_j f_{v,t+1} + u_{j,t+1}.$$

In what follows, we assume that $\rho_j = \rho$ for all assets.

We consider the second stage regression. To further simplify, rather than estimating the risk premium parameter $\lambda_v \equiv \mu_v^P - \mu_v^Q$, we consider estimating μ_v^Q from the cross-sectional regression of the first-stage estimate of γ_j on $\mathcal{V}_j$. We denote the estimand μ in the true model and $\tilde{\mu}$ in the mis-specified model to distinguish it from the true μ_v^Q . In the true model, we have:

$$-\gamma_j = \mathcal{V}_j \mu_v^Q = \mathcal{V}_j \mu + e_j,$$

where e denotes the cross-sectional error in the risk-neutral moment estimation. The OLS estimator of μ is

$$\hat{\mu} = \left(\sum_j \mathcal{V}_j^2 \right)^{-1} \sum_j \mathcal{V}_j^2 \mu_v^Q \quad \text{with } \text{plim} \hat{\mu} = \mu_v^Q,$$

that is, the population value of $\hat{\mu}$ is μ_v^Q .

Now consider the mis-specified model and recall that we can write the population value of $\hat{\gamma}_j$, using the fact that $-\gamma_j = \mathcal{V}_j \mu_v^Q$, as

$$\text{plim } \hat{\gamma}_j = \gamma_j - \Delta_j (\mu_m^P - \rho \mu_v^P).$$

Our second stage regression to estimate $\tilde{\mu}$ is then

$$\mathcal{V}_j \mu_v^Q + \Delta_j (\mu_m^P - \rho \mu_v^P) = (\mathcal{V}_j + \Delta_j \rho) \tilde{\mu} + \tilde{e}_j.$$

The population value of the OLS estimator is

$$\text{plim } \hat{\tilde{\mu}} = \mu_v^Q + \xi [\rho(\mu_v^P - \mu_v^Q) - \mu_m^P].$$

where

$$\xi = \left[\sum_j (\mathcal{V}_j + \rho \Delta_j)^2 \right]^{-1} \sum_j (\mathcal{V}_j + \rho \Delta_j) \Delta_j,$$

is the coefficient of the regression of the exposures to the omitted factor on the exposures to the VAR factor estimated in the mis-specified model.

Examining the above equation, we see there are multiple sources of bias. First of all, the bias disappears if $\Delta_j = 0$, which is when there is no residual exposure to $rs_{j,t+1}$ after hedging out $f_{v,t+1}$. But if some $\Delta_j \neq 0$, a bias term remains even if we consider $\rho = 0$, as long as $\mathcal{V}_j$ is correlated with Δ_j :

$$\text{plim } \widehat{\mu}_{\rho=0} = \mu_v^Q - \left[\sum_j (\mathcal{V}_j)^2 \right]^{-1} \sum_j \mathcal{V}_j \Delta_j \mu_m^P.$$

Note that in the $\rho = 0$ case, the bias term is proportional to the negative of the equity premium, and the constant of proportionality is equal to the coefficient of the regression of Δ_j (residual exposure to underlying) on $\mathcal{V}_j$.

Appendix D. IPCA Estimation and Performance

We follow Büchner and Kelly (2022) and Goyal and Saretto (2024) and estimate a 1, ..., 5-factor IPCA model with a broad set of option and stock level characteristics. Furthermore, we use both the Black-Scholes (BS) and our nonparametrically estimated delta-hedged delevered option returns in the estimation. We estimate models across SPX, stock and ETF equity option returns using the same set of option-only characteristics. Thereafter we include an additional set of equity-level characteristics as a robustness exercise. Our vector of option characteristics includes a total of 10 variables. The robustness exercise conducted for stock and ETF options augments this list with 16 stock-based variables. A detailed description, ID and citation related to the computation of each characteristic is provided in Table A.2 of the Appendix. We model calls and puts jointly, and follow Büchner and Kelly (2022) by interacting the characteristic variables and the constant with a dummy variable encoded as one for put options and zero otherwise, resulting in a total of 22 candidate characteristics for the baseline model (option characteristics only) and 54 for the robustness (option and equity characteristics) with stock and ETF equity options. All characteristic variables are cross-sectionally ranked and rescaled to the $[-1, 1]$ interval. Table 11 presents the total (pooled) R_p^2 from fitting IPCA models with 1 to 5 factors to the nonparametric (BS)

delta-hedged delevered option returns. Panel A shows results for SPX options, while Panel B covers stock and ETF options. Panels C and D report the R_p^2 for stock and ETF options, respectively, with the inclusion of stock/ETF-level characteristics. Our results, across all panels, strongly confirm the findings by Büchner and Kelly (2022) and Goyal and Saretto (2024) – IPCA is a highly effective model for capturing variation in SPX and stock option returns.

The estimation applied to SPX options consistently exhibit the highest R_p^2 across all specifications (Panels A and B), with five-factor R_p^2 close to 98%. These results strongly confirm findings from Büchner and Kelly (2022) with weekly data and over a different sample, that a small number of latent factors can explain a substantial portion of both the nonparametric or BS delevered delta-hedged option return variation. In contrast, R_p^2 for stock and ETF options are lower, in line with the results of Goyal and Saretto (2024) for stock options. In Panels A (B) the five-factor IPCA model explains 33.35% (36.79%) for the nonparametric (BS) delevered delta-hedged stock put option returns. For calls, the R_p^2 is lower at 26.40% (23.50%) respectively. For ETFs, the values are relatively higher – 56.01% (55.80%) for puts and 57.13% (53.47%) for calls. In Panels C and D, we augment the estimation with additional equity-based instruments, which leads to marginally higher R_p^2 for stock and ETF option types, with larger improvements for ETF options.