Beyond the Short Run: The Term Structure of Implied Moment Risk Premia

Meng Zhang and Jose Olmo*

Abstract

This paper investigates the pricing of the term structure of risk-neutral moments. We construct slope factors by differencing six-month against one-month, and twelve-month against six-month implied moments to isolate incremental distributional risks. We document distinct pricing patterns across horizons: while variance commands a persistent negative premium, the pricing power for skewness and kurtosis is concentrated in the medium-term slope rather than the short-term level. Notably, the medium-term skewness slope carries a significant positive premium. Consistent with the ICAPM, these risk premia align with the factors' predictive power for future market returns. We also uncover a fundamental macro-dichotomy: short-term variance tracks real economic activity, whereas the term structure of higher-order moments is linked to monetary conditions and credit stress. These findings are robust to alternative test assets, the exclusion of small-cap stocks, and controls for liquidity and momentum.

*Meng Zhang, University of Southampton, Department of Economics, Highfield Campus, SO17 1BJ, Southampton, UK. E-mail: mz3n23@soton.ac.uk. Jose Olmo (corresponding author), University of Southampton and Universidad de Zaragoza, Department of Economics, Highfield Campus, SO17 1BJ, Southampton, UK, and Department of Economic Analysis, Universidad de Zaragoza, Gran Vía 2, 50005 Zaragoza, Spain. E-mail: J.B.Olmo@soton.ac.uk and joseolmo@unizar.es. We thank Gonzalo Rubio, Universidad CEU Cardenal Herrera, for helpful comments and suggestions. Jose Olmo also acknowledges financial support from Ministerio de Ciencia, Innovación y Universidades de España (Grant PID2023-147798NB-I00) and Fundación Agencia Aragonesa para la Investigación y el Desarrollo (ARAID).

1 Introduction

It is well documented in the empirical asset pricing literature [Ang et al., 2006, Adrian and Rosenberg, 2008, Chang et al., 2013] that the cross section of stock returns carries a negative premium on implied market variance. This is often proxied by the VIX index, which measures the risk-neutral variance of the market portfolio return derived from option prices. Stocks with higher exposures to VIX earn lower average returns than comparable stocks with lower exposures. The standard explanation, grounded in the intertemporal CAPM (ICAPM) of Merton [1973] and Campbell [1993], is that investors value assets that hedge adverse shifts in the investment opportunity set. Because increases in the VIX forecast deteriorations in market conditions, stocks with high exposure to the VIX trade at a premium; they provide hedging benefits precisely when marginal utility is high.

Beyond implied (risk-neutral) variance (IV), the literature has recently extended the analysis to higher-order implied moments, including implied skewness (IS) and kurtosis (IK) [Agarwal et al., 2009, Xing et al., 2010, Conrad et al., 2013, Chang et al., 2013]. These measures, constructed from option prices using the model-free approach of Bakshi and Madan [2000], Carr and Madan [2001], and Bakshi et al. [2003], are forward-looking and avoid the need to estimate higher moments from historical returns. Because they are derived under risk-neutral expectations, implied moments reflect not only the physical distribution of returns but also investors' risk aversion. Empirical evidence shows that implied skewness carries a negative premium in the cross section of stock returns, in contrast to the theoretical predictions discussed in Chang et al. [2013]. Although negative skewness corresponds to adverse states in standard asset pricing theory and should command a positive premium, stocks with higher exposure to skewness risk earn lower average returns. This puzzle has been rationalized by showing that innovations in skewness forecast future increases in variance, making skewness an effective "bad news" state variable. Evidence on implied kurtosis is comparatively limited, but it has been investigated as a proxy for tail risk, reflecting investors' concerns about extreme market events [Dittmar, 2002].

Most of the existing empirical asset pricing literature focuses on short-term (monthly) implied distributions and their relation to contemporaneous stock returns. However, investors' expectations about longer-term shifts in the investment opportunity set may also contain information relevant to short-term pricing. Several channels can generate such effects. First, risk aversion is inherently forward looking, and expectations of deteriorating long-term conditions are priced into current returns. Second, distinct horizons capture different

dimensions of risk dynamics, distinguishing between transitory shocks and persistent shifts in the return distribution. These mechanisms suggest that information contained in the term structure of implied moments is priced in the cross section of stock returns.

Our analysis builds on this insight by studying the term structure of risk-neutral moments of the market return distribution. Extending the ICAPM framework of Merton [1973] and Campbell [1993], we examine how implied variance, skewness, and kurtosis affect the pricing of stock returns across different horizons. To capture the information embedded in the term structure, we focus on the incremental changes in risk perceptions across maturities. Specifically, we use the one-month moment to represent the short-term risk level and construct term-structure factors defined as the spreads between different maturities (six-month minus one-month, and twelve-month minus six-month). This specific construction allows us to isolate the information specific to medium- and long-term horizons, enabling us to assess whether horizon-specific variation in higher-order risks is systematically priced and whether investor attitudes toward downside risk and tail events differ across maturities.

The empirical strategy followed in this study is straightforward. We compute these risk-neutral moments from a surface of S&P 500 index options across strikes and maturities, interpolating to obtain standardized values. Since implied moments are not directly tradable assets, we construct mimicking portfolios to capture the risk premia associated with these factors. Specifically, we sort stocks based on their estimated exposures to each moment—short-term level and term-structure spreads—and form high-minus-low quintile portfolios. We then implement asset pricing tests using Fama-MacBeth (FM) cross-sectional regressions, comparing models that augment the Fama-French five-factor benchmark (FF-5) with our moment-based factors. This framework allows us to estimate horizon-specific risk premia and assess whether higher-order risks across maturities contain incremental explanatory power beyond standard sources of systematic risk.

Several key findings stand out from the empirical analysis. First, implied variance (IV) carries a significant negative risk premium at both the short and medium horizons, while evidence for the long horizon is weak. This aligns with Dew-Becker et al. [2017], who document a negative short-term volatility premium that vanishes at longer maturities, a pattern also observed in foreign exchange [Della Corte et al., 2021] and bond markets [Choi et al., 2017]. Our results extend this evidence by showing that the hedging role of variance is not confined to the immediate month but persists strongly into the medium term. Second, and in sharp contrast to variance, the pricing of higher-order moments exhibits a distinct pattern from the level to the slope. We find that for implied skewness (IS) and implied kurtosis (IK), pricing power is concentrated in

the term-structure slopes. Specifically, the medium-term implied skewness (IS^{ms}) commands a significant positive premium, while the medium-term kurtosis slope (IK^{ms}) carries a significant negative premium. Although implied skewness and kurtosis are statistically linked through the geometry of the volatility surface, this horizon-dependent pattern suggests that they capture distinct economic risks: while investors may tolerate transitory spikes in tail risk, they require substantial compensation for exposure to persistent deteriorations in the shape of the market return distribution.

We further provide economic interpretation for the empirical results along three dimensions. First, we validate the pricing mechanism through the lens of intertemporal hedging and risk aversion. By examining the joint dynamics of implied moments and market returns, we document a sign reversal: the term-structure factors that command negative risk premia are shown to significantly forecast higher future market excess returns. This confirms that the observed premiums represent rational compensation for bearing risks that signal deteriorating investment opportunities. Second, we address the potential concern that term-structure factors might simply reflect the persistence of short-term shocks. Using Vector Autoregression (VAR) analysis, we demonstrate that the medium- and long-term factors exhibit independent dynamics. Third, we link these risk premia to fundamental economic states. The term structure is anchored to distinct macroeconomic forces: while real economic activity primarily drives the pricing of short-term implied variance, inflationary pressures and credit market conditions are the dominant drivers of medium-term implied skewness and kurtosis. This distinction reveals that while short-term volatility reflects the business cycle, the pricing of tail-risk persistence is linked to monetary and credit environments.

Our empirical findings remain robust to a wide set of additional checks. We first verify that the pricing power of our term-structure factors is not subsumed by standard benchmarks or mechanical correlations. Spanning tests confirm that the medium-term slope factors contain unique pricing information orthogonal to the FF-5 model, while bivariate regressions demonstrate that their significance persists even after controlling for the baseline risk levels. Furthermore, our results are not driven by specific methodological choices or micro-cap anomalies; the horizon-dependent risk premia remain robust across alternative sets of test assets. This robustness holds when using double-sorted portfolios and portfolios formed on book-to-market and operating profitability, as well as when small-cap stocks are excluded from the analysis. Finally, the main findings remain unaltered after controlling for aggregate liquidity risk and momentum, confirming that the term structure of implied moments captures distinct sources of systematic risk beyond established factors.

Our work contributes to the growing literature on the pricing of option-implied information. Following

the seminal work of Ang et al. [2006] on the VIX, subsequent studies have established that higher-order moments, such as implied skewness and kurtosis, contain significant information about future returns [Chang et al., 2013, Conrad et al., 2013, Li et al., 2023]. Concurrently, a separate strand of literature investigates the term structure of risk premia, focusing primarily on variance risk [Dew-Becker et al., 2017, Choi et al., 2017] or individual stock skewness [Borochin et al., 2020]. We bridge these lines of inquiry by analyzing the term structure of higher-order moments. Our paper extends this research by focusing not only on the level of higher moments but also on their term structure, thereby linking risk-neutral information across horizons to the pricing of stock returns.

The remainder of the paper is organized as follows. Section 2 describes the data, defines the implied-moment variables, and explains the construction of the term-structure factors. Section 3 presents the baseline empirical results on the cross-sectional pricing of implied variance, skewness, and kurtosis. Section 4 analyzes the economic mechanisms, examining the predictive relations for market returns, the dynamic interdependencies among term-structure factors, and the macroeconomic determinants of risk premia. Section 5 reports a comprehensive set of robustness checks, including model diagnostics, spanning tests, and controls for alternative test assets and risk factors. Section 6 concludes.

2 Data and Variable Definitions

Our analysis uses all common stocks listed on the NYSE, AMEX, and NASDAQ over the period June 1, 2012, to October 31, 2025. By focusing on the most recent decade of data, we ensure that the analysis reflects contemporary market dynamics. Daily stock prices and market capitalizations are obtained from CRSP via WRDS. To ensure robustness, we filter out extreme observations: for example, stocks trading below one dollar in a given month are excluded from the regressions for that month but remain in the overall cross section, preserving the breadth of the dataset. We complement the stock data with daily observations on the FF-5 factors, Carhart's momentum factor, the risk-free rate (R_F), and the Fama-French 49 industry portfolios, all sourced from Kenneth French's online data library.[1] These factors allow us to control for standard sources of systematic risk when estimating the cross section of returns.

To construct our state variables, we use daily option data from CBOE and OptionMetrics, covering all traded strikes and maturities for SPX options. These data are employed to compute IV, IS, and IK of the

[1] https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

market portfolio returns, proxied by the S&P 500 index returns. In computing implied moments, we exclude strikes with zero trading volume to ensure reliable estimates. The resulting measures serve as state variables in our empirical framework, capturing the expectations and risk preferences of market participants across different horizons and informing our construction of mimicking portfolios for cross-sectional asset pricing.

To study how the risk premia associated with our state variables vary with macroeconomic conditions, we regress the time series of risk premia estimates on a set of representative macroeconomic indicators. The data are obtained from the Federal Reserve Bank of St. Louis FRED database.

2.1 Derivation of Risk-Neutral Moments Using Option Data

Higher-order moments under the risk-neutral measure can be estimated using either model-based or model-free approaches. Model-based methods, such as the stochastic-variance model of Heston [1993] and its jump-diffusion extensions [Bates, 1996], capture stochastic variance and fat-tailed return distributions but require estimation of numerous parameters (e.g., jump intensity and amplitude), making them sensitive to strong prior assumptions [Bakshi et al., 1997].

We adopt a model-free approach, following Bakshi and Madan [2000] and Carr and Madan [2001], which directly extracts higher-order moments from observed option prices without imposing any parametric distribution on returns. This market-driven method is widely used in recent literature [Chang et al., 2012, 2013, Borochin et al., 2020] and has the advantage of relying solely on traded prices. However, it can perform poorly when data are sparse, necessitating interpolation of option prices [Bates, 2012]. Our implementation proceeds in three steps. First, we compute implied volatilities across all traded strikes. Second, these volatilities are interpolated to construct a complete volatility surface, filling in missing prices at each strike. Third, we input these interpolated prices into the model-free method of Bakshi and Madan [2000] to compute the higher-order moments. Appendix A provides technical details on the calculation procedure.

For comparability across horizons, option maturities are aligned to standardized one-, six-, and twelve-month terms. Consistent with the methodology employed by the CBOE in constructing the VIX index, we identify the two option maturities that straddle the target horizon and compute risk-neutral moments separately at these maturities using the model-free approach. We then obtain the standardized maturity values by linearly interpolating the two moment estimates with respect to time to maturity.

2.2 Term Structure Factors

Two main approaches exist for incorporating option-implied information into asset pricing. The first, used by Ang et al. [2006], treats implied variance as a forward-looking state variable and estimates the prices of risk for assets' exposures to these moments through cross-sectional regressions. The second, represented by Carr and Wu [2009], computes variance risk premia as the difference between risk-neutral and realized moments, focusing on their predictive power for aggregate market returns. However, extending this second approach to higher-order moments is far from straightforward. As noted by Chang et al. [2013], obtaining reliable realized skewness and kurtosis is empirically challenging, and identifying their corresponding risk premia requires specifying structural models for volatility and jump-risk dynamics. Such parametric assumptions introduce model dependence and estimation noise that can obscure the economic interpretation of the results. In contrast, risk-neutral higher-order moments extracted directly from option prices are forward-looking, model-free, and can be estimated daily using a single cross section of option data. Therefore, we follow the first approach and employ option-implied variance, skewness, and kurtosis as the state variables in our analysis.

The information underlying the term structure of risk-neutral moments is captured by three state variables defined for each moment. The short-term factor (s) corresponds to the one-month moment, serving as the baseline level of the term structure. The medium-term slope factor (ms) is defined as the six-month moment minus the one-month moment, representing the slope at the front end of the curve. The long-term slope factor (lm) is defined as the twelve-month moment minus the six-month moment. This construction differs from standard approaches that define the long-term factor relative to the short end. By measuring the spread relative to the medium term, the lm factor serves as a proxy for the incremental risk information specific to the long horizon. Throughout the paper, "short-term", "medium-term", and "long-term" refer to the state variables s, ms, and lm, respectively.

Our construction of the slope factors relies on the spread between implied moments at different maturities. This approach is motivated by the objective of isolating the incremental risk information specific to medium and long horizons, akin to extracting forward rates from the yield curve. For implied variance, under the standard assumption of serially uncorrelated returns, the difference between the six-month and one-month variance mathematically corresponds to the expected variance of the future return between these dates. For standardized higher-order moments like skewness and kurtosis, exact additivity does not hold due to the

non-linear normalization by variance. We construct these moment spreads to capture the slope of the term structure. As detailed in Appendix B, while the observed spread does not strictly equal the theoretical incremental moment due to variance normalization, the discrepancy is systematic. It explicitly captures the interaction between incremental distributional risks and intertemporal volatility dependence (e.g., volatility clustering). Consequently, significant premia on these slope factors reflect compensation for these distinct horizon-dependent risks.

This factor construction approach distinguishes our analysis from prior studies, such as Ang et al. [2006] and Chang et al. [2013], in two main respects. First, we define the long-term factor relative to the medium-term horizon (lm) rather than the short-term baseline. By measuring the long term relative to the medium term (12-month minus 6-month), we isolate the pricing information specific to the distant end of the term structure, ensuring it is distinct from the information already captured by the medium-term slope. Second, we employ the level values of these slope variables rather than their time-series innovations. This specification aligns with the ICAPM framework because the slope variables themselves represent innovations along the maturity dimension. They reflect how investors update their hedging demands across different horizons, making the level of the slope a direct measure of the prevailing term structure of risk.

Table 1: Summary Statistics of the Nine State Variables Derived from Implied Moments. This table reports the maximum, minimum, mean, and standard deviation (across daily observations) of nine state variables derived from implied moments. s denotes the one month maturity; ms and lm denote the six month minus one month and twelve month minus six month, respectively. IV, IS, and IK are extracted from SPX option prices using the method of Bakshi and Madan [2000] and Carr and Madan [2001] with an interpolated volatility surface. See Section 2.1 and Appendix A for construction details. The sample covers from June 1, 2012 to October 31, 2025.

	IV^s	IV^{ms}	IV^{lm}	IS^s	IS^{ms}	IS^{lm}	IK^s	IK^{ms}	IK^{lm}
Max	0.0535	0.21	0.38	-0.70	4.70	5.02	149.41	46.53	93.37
Min	0.0006	0.01	0.00	-8.11	-2.94	-4.45	3.42	-108.93	-84.60
Mean	0.0030	0.02	0.03	-3.13	0.18	0.59	26.10	-6.58	-6.40
Std.	0.0034	0.01	0.01	1.09	1.01	0.79	16.50	14.08	9.42

Augmented Dickey-Fuller (ADF) tests confirm that all state variables are stationary at the 10% significance level, supporting the subsequent regression analyses. Table 1 reports summary statistics for the nine variables derived from the risk-neutral moments. IV^s, IV^{ms}, and IV^{lm} are strictly positive. In contrast, implied skewness shows distinct behavior: IS^s is consistently negative, reflecting persistent downside risk, whereas IS^{ms} and IS^{lm} take both positive and negative values, with the long-term variable displaying larger scale. Implied kurtosis also differs: IK^s is positive, consistent with fat-tailed return distributions, while IK^{ms}

and IK^{lm} fluctuate in sign and exhibit greater magnitude than the short-term level. Figure 1 complements these observations by plotting the daily series of the nine variables over the evaluation period.

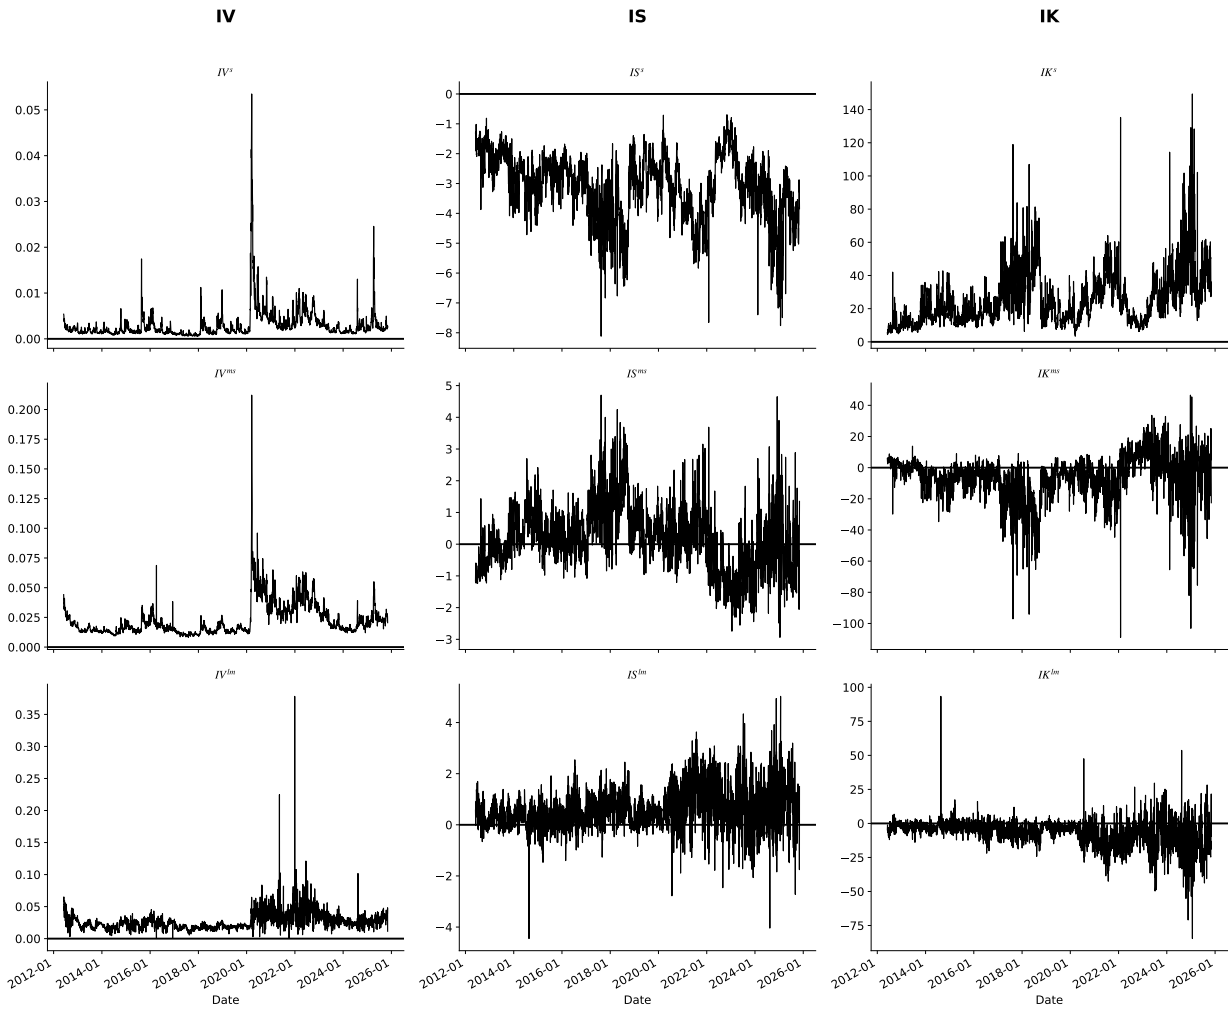

Figure 1: Daily Time Series of the Nine State Variables Derived from Implied Moments. The nine subfigures correspond to IV^s, IV^{ms}, IV^{lm}, IS^s, IS^{ms}, IS^{lm}, IK^s, IK^{ms}, and IK^{lm}. The horizontal zero line is drawn in bold. The sample period is from June 1, 2012 to October 31, 2025.

We examine the pairwise correlations among the state variables, with a particular focus on the correlations among s, ms, and lm within each implied moment. This focus is driven by our specific research objective: to isolate the incremental information contained in the term structure of each moment. While prior literature has extensively examined the interactions between implied variance, skewness, and kurtosis (e.g., Chang et al. [2013]), our contribution lies in identifying the pricing content of the horizon dimension. Consequently, the primary concern regarding multicollinearity arises from the dependencies between different maturities of the same moment, rather than the relationships across different moments. For implied variance, the

correlation between s and ms is 0.81, while the correlation between s and lm is 0.30. For implied skewness, the correlation between s and ms is -0.65, and the correlation with lm is -0.25. Similarly, implied kurtosis shows a correlation of -0.69 between s and ms, and -0.32 between s and lm. Although the results indicate no perfect multicollinearity, the correlation between s and ms is relatively high.

To address this issue and verify that the factors capture distinct pricing information, we provide four sets of empirical evidence. First, we construct tradable mimicking portfolios and document that the term-structure factors often command risk premia with distinct signs and significance levels (Section 3), implying that they compensate for different sources of risk. Second, the predictive regressions in Section 4 show that these factors capture different predictive information regarding future market returns. Third, the VAR analysis in Section 4 demonstrates that ms and lm exhibit independent dynamics that are not fully explained by shocks to s. Finally, the bivariate analysis in Section 5 confirms that the pricing ability of ms remains significant after controlling for s.

3 Empirical Analysis

This section examines the empirical relevance of the nine term-structure state variables for pricing the cross section of returns. It reports the estimated risk premia on the state variables over the full evaluation period, based on three complementary empirical strategies: (i) time-series estimation of stock return exposures to the implied-moment factors, (ii) cross-sectional FM regressions, and (iii) analysis of the time variation in the estimated risk premia. To avoid redundancy, the methodology is presented once using the IV state variables as an illustration. The empirical results for IS and IK are directly reported in their respective subsections.

3.1 Methodology

Let $R^e_{it} = R_{it} - R_{ft}$ and $R^e_{mt} = R_{mt} - R_{ft}$ denote the excess returns on stock i and the market portfolio R_{mt}, respectively, with R_{ft} the risk-free rate. For each $IV^{(j)}_t \in \{IV^s_t, IV^{ms}_t, IV^{lm}_t\}$, with $j = s, ms, lm$, factor loadings are obtained from the time-series regression

$$R^e_{it} = \alpha_i + \beta_{m,i} R^e_{mt} + \beta^{(j)}_{v,i} IV^{(j)}_t + \varepsilon_{it}, \tag{1}$$

where ε_{it} is the idiosyncratic error term. These regressions are estimated at the end of each month using daily returns from that month, updating exposures $\beta_{v,i}^{(j)}$ on a monthly basis.

At each month-end, stocks are sorted into quintiles according to their estimated exposure $\beta_{v,i}^{(j)}$. Quintile 1 contains stocks with the lowest exposure, while Quintile 5 contains those with the highest. To avoid look-ahead bias, the value-weighted returns of these quintiles are computed from daily returns in the subsequent month. Denoting these returns by $R_{Q1,t}^{(j)}, \ldots, R_{Q5,t}^{(j)}$, the tradable mimicking factor is constructed as a high-minus-low (HML) portfolio:

$$Q_t^{(j)} = R_{Q5,t}^{(j)} - R_{Q1,t}^{(j)}, \quad j = s, ms, lm. \tag{2}$$

We evaluate the pricing ability of the constructed mimicking portfolios using the FM two-pass procedure, with the Fama–French 49 industry portfolios as test assets. These portfolios provide broad cross-sectional coverage of U.S. equities while effectively diversifying idiosyncratic noise relative to individual stocks. In the robustness analysis section, we further verify our results using alternative sets of test assets, partly addressing the criticism in Lewellen et al. [2010] and Boons [2016] that asset pricing inferences can depend on the specific choice of test portfolios.

In the first stage, factor exposures are estimated under CAPM and FF-5 as baseline models. For CAPM, we estimate

$$R_{it}^e = \alpha_i + \beta_{m,i} R_{mt}^e + \beta_{Q,i}^{(j)} Q_t^{(j)} + \varepsilon_{it}, \tag{3}$$

where $\beta_{m,i}$ is the loading on the market excess return R_{mt}^e, and $Q_t^{(j)}$ is the tradable factor associated with $IV^{(j)}$ with loading $\beta_{Q,i}^{(j)}$. For FF-5, we estimate

$$R_{it}^e = \alpha_i + \tilde{\beta}_i' \tilde{F}_t + \beta_{Q,i}^{(j)} Q_t^{(j)} + \varepsilon_{it}, \tag{4}$$

where $\tilde{F}_t = (R_{mt}^e, SMB_t, HML_t, RMW_t, CMA_t)'$ is the vector of the FF-5 risk factors, and $\tilde{\beta}_i$ are the corresponding factor loadings. Note that the intercepts α_i appearing in (1), (3), and (4) are regression-specific constants and do not carry over across different stages of the procedure.

In the second stage, the estimated factor loadings are used in cross-sectional regressions to obtain risk premia:

$$R^e_{i,t} = \lambda^{(j)}_{0,t} + \lambda'_{F,t}\widehat{\tilde{\beta}}_i + \lambda^{(j)}_{Q,t}\widehat{\beta}^{(j)}_{Q,i} + \varepsilon_{i,t}, \quad j = s, ms, lm, \tag{5}$$

where $R^e_{i,t}$ is the excess return of portfolio i in month t. The terms $\widehat{\tilde{\beta}}_i$ and $\widehat{\beta}^{(j)}_{Q,i}$ are the estimated factor loadings from either (3) or (4), associated with the benchmark factors (CAPM or FF-5) and the implied-moment state variable, respectively. The coefficients $\lambda_{F,t}$ and $\lambda^{(j)}_{Q,t}$ represent the estimated price of risk at month t. The time-series average of the monthly coefficients,

$$\bar{\lambda}_k = \frac{1}{T}\sum_{t=1}^{T}\lambda_{k,t}, \tag{6}$$

represents the estimated monthly risk premium for factor k, corresponding to the unconditional risk premia of the CAPM and FF-5 risk factors and the implied moment, respectively.

3.2 Implied Variance

Implied variance has been extensively studied as a priced factor. Ang et al. [2006] use innovations to the VIX index, given by its first differences (ΔVIX), as a proxy for economic uncertainty, documenting significantly negative factor returns and risk premia. Chang et al. [2013] extend this framework by incorporating innovations to IS and IK over longer horizons.

Table 2 reports portfolio statistics for quintiles sorted on the beta loadings of IV^s (Panel A), IV^{ms} (Panel B), and IV^{lm} (Panel C). Each panel presents six columns: Quintiles 1 through 5, and the 5-1 portfolio. The 5-1 portfolio, denoted as Q_t above, represents the difference in portfolio returns between the highest and lowest quintiles. Row 1 of each panel reports pre-formation betas for each quintile. Row 2 reports the average monthly returns, computed as the time-series mean of monthly portfolio returns, where each month's return is based on daily returns in the month following portfolio formation. Rows 3 and 4 present risk-adjusted returns: Row 3 reports CAPM alphas, defined as the time-series alpha relative to the CAPM, and Row 4 reports FF-5 alphas, defined as the time-series alpha relative to the FF-5 model. The t-statistics for the mean return, CAPM alpha, and FF-5 alpha, computed using Newey–West standard errors, are reported in parentheses below each estimate. Table 3 reports the estimated risk premium (λ), computed as the time-series average of the monthly prices of risk from the second-pass FM regressions. The CAPM columns report estimates controlling for the market excess return, while the FF-5 columns control for the FF-5 factors.

Table 2 shows that the 5-1 portfolio for IV^s exhibits a monthly mean return of −0.21%, a CAPM alpha of −0.29%, and an FF-5 alpha of −0.30%. Consistent with the portfolio-level evidence, Table 3 reports a significantly negative price of risk of −0.56 in the CAPM, whereas the corresponding estimate under the FF-5 model is −0.34 and is not statistically significant. For the medium term, the negative pricing pattern is more pronounced. The 5-1 portfolio of IV^{ms} records a mean return of −0.54%, with corresponding CAPM and FF-5 alphas of −0.57% and −0.55%, respectively. This is corroborated by the FM regressions, where medium-term IV shows a significant negative premium, estimated at −0.89 for CAPM and −0.50 for FF-5. In contrast, evidence for IV^{lm} is weaker, with a 5-1 mean return of 0.33% and CAPM and FF-5 alphas of 0.18% and 0.22%, respectively. Similarly, the estimated FM risk premia are 0.04 and 0.32 for CAPM and FF-5, respectively, neither of which is statistically significant.

Figure 2 illustrates the time variation in the estimated risk premia for implied variance across horizons. To visualize the variation of risk compensation over the sample period, we superimpose a nonparametric trend using Locally Estimated Scatterplot Smoothing (LOESS). This allows us to capture the low-frequency variation in risk premia that corresponds to the changing economic conditions discussed in Section 4. For IV^s, the LOESS-smoothed series is negative for most of the sample period and turns positive toward the end of the sample. For IV^{ms}, the estimated risk premium remains negative throughout the sample, with the smoothed series staying well below zero over time. For IV^{lm}, the estimated risk premium fluctuates around zero for most of the sample and shows an upward movement starting in 2023. This period coincided with heightened market concerns regarding economic deceleration following aggressive monetary tightening. In the context of the ICAPM, such persistent headwinds signal a deterioration in the long-term investment opportunity set. Overall, these time-varying dynamics align with the unconditional risk premia estimates reported in Table 3.

3.3 Implied Skewness

The term structure of implied skewness remains largely unexplored in the empirical asset pricing literature. One notable exception is Borochin et al. [2020], who examine risk-neutral skewness for individual stock options with maturities of one, six, and twelve months.[2]

[2] These authors document significantly positive (negative) premia for short-term (1-month) and long-term (12-month) skewness, respectively, while the medium-term skewness is found to be statistically insignificant. To capture the long-minus-short maturity effect, they construct a stock-level factor defined as the difference between long- and short-term skewness (12m–1m), analogous in spirit to the SMB and HML Fama-French factors. Their factor construction methodology and empirical approach differ fundamentally from ours, and the resulting empirical findings are also markedly different.

Table 2: Portfolio Statistics Sorted on IV Betas. This table presents the summary statistics of quintile portfolios formed at the end of each month by sorting stocks on their estimated beta loadings with respect to IV^s (Panel A), IV^{ms} (Panel B), and IV^{lm} (Panel C). Stocks are ranked into quintiles based on their pre-ranking betas, which are estimated from daily returns within the month. Value-weighted portfolios are then formed at the end of the estimation month, and portfolio returns are computed from daily returns in the subsequent month. Each panel reports six columns: Quintiles 1 through 5 and the 5-1 portfolio, defined as the return difference between the highest- and lowest-beta quintiles. Row 1 reports the average pre-ranking betas used for portfolio formation. Row 2 reports the mean monthly returns, calculated as the time-series average of monthly post-ranking returns, reported in percent. Rows 3 and 4 report risk-adjusted returns in percent: CAPM alpha (Row 3) and FF-5 alpha (Row 4), defined as intercepts from time-series regressions relative to the respective benchmark models. t-statistics, based on Newey–West standard errors, are reported in parentheses.

Quintiles	1	2	3	4	5	5-1
Panel A: β_{IV^s}						
Pre Formation β	-13.09	-2.81	1.23	3.87	12.16	
Mean Return	1.07	0.94	0.92	1.01	0.86	-0.21
	(2.74)	(2.78)	(2.96)	(3.08)	(2.30)	(-0.84)
CAPM Alpha	-0.24	-0.26	-0.21	-0.17	-0.40	-0.29
	(-1.51)	(-1.77)	(-2.83)	(-2.11)	(-2.97)	(-1.29)
FF-5 Alpha	-0.17	-0.27	-0.27	-0.20	-0.35	-0.30
	(-1.11)	(-2.10)	(-3.38)	(-2.34)	(-2.71)	(-1.38)
Panel B: $\beta_{IV^{ms}}$						
Pre Formation β	-17.85	-4.12	2.02	3.82	14.68	
Mean Return	1.37	0.95	0.88	0.95	0.81	-0.54
	(3.44)	(2.84)	(2.88)	(2.88)	(2.20)	(-2.25)
CAPM Alpha	0.03	-0.26	-0.23	-0.23	-0.44	-0.57
	(0.19)	(-2.80)	(-2.00)	(-1.95)	(-3.07)	(-2.59)
FF-5 Alpha	0.07	-0.29	-0.25	-0.23	-0.38	-0.55
	(0.43)	(-3.42)	(-2.74)	(-2.44)	(-3.04)	(-2.50)
Panel C: $\beta_{IV^{lm}}$						
Pre Formation β	-9.47	-4.52	-0.26	8.69	9.54	
Mean Return	0.82	1.05	0.95	0.97	1.18	0.33
	(2.19)	(3.05)	(3.03)	(2.93)	(3.03)	(1.28)
CAPM Alpha	-0.45	-0.16	-0.19	-0.21	-0.12	0.18
	(-3.24)	(-1.00)	(-1.83)	(-2.14)	(-0.62)	(0.69)
FF-5 Alpha	-0.39	-0.17	-0.23	-0.22	-0.03	0.22
	(-2.89)	(-1.25)	(-3.00)	(-2.19)	(-0.15)	(0.95)

Table 3: Risk Premium of IV. This table reports the estimated risk premium (λ) for the 5-1 mimicking portfolios of IV^s, IV^{ms}, and IV^{lm}. The risk premium is obtained from monthly FM cross-sectional regressions. The CAPM columns report estimates controlling for the market excess return, and the FF-5 columns report estimates controlling for the FF-5 factors. Reported values are the time-series averages of the monthly estimates of λ. t-statistics, based on Newey–West standard errors, are shown in parentheses.

	CAPM			FF-5		
	IV^s	IV^{ms}	IV^{lm}	IV^s	IV^{ms}	IV^{lm}
λ	-0.56	-0.89	0.04	-0.34	-0.50	0.32
	(-2.17)	(-3.97)	(0.11)	(-1.32)	(-2.01)	(1.21)

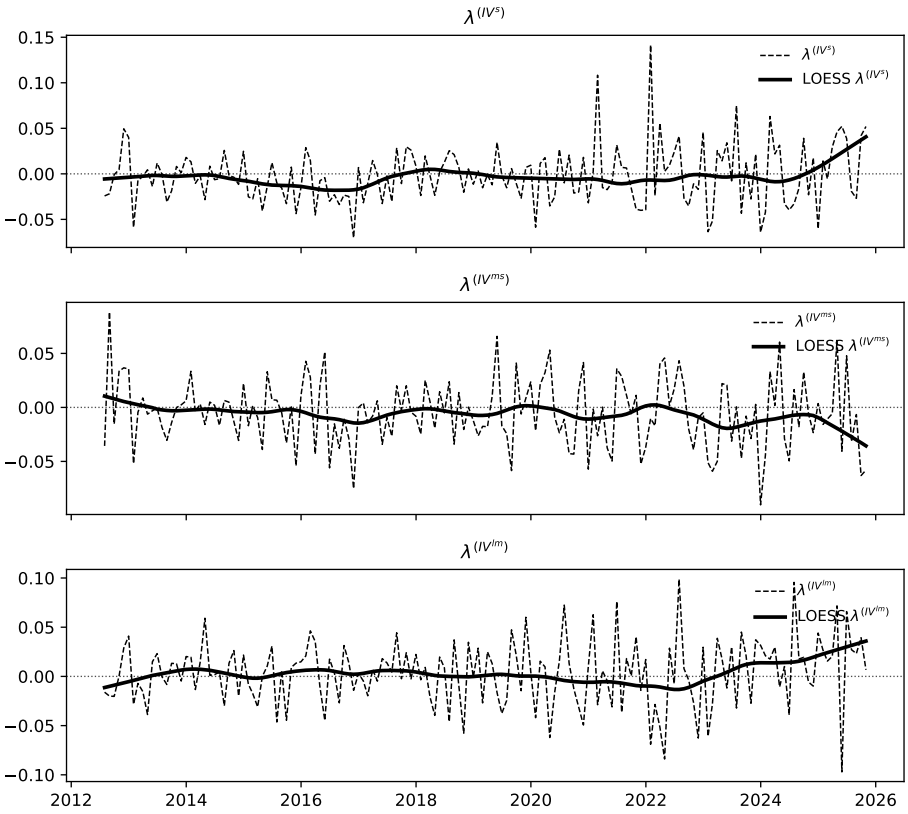

Figure 2: Time-Varying Risk Premium of IV. This figure plots the monthly estimates of the risk premium (λ) for IV^s, IV^{ms}, and IV^{lm}. The risk premium is obtained from monthly FM cross-sectional regressions controlling for the FF-5 factors. The figure reports the raw monthly λ estimates and their corresponding LOESS-smoothed series.

Following the structure adopted in the preceding section, we present the portfolio statistics and the corresponding FM risk premia for IS in Table 4 and Table 5, respectively. For IS^s, Panel A of Table 4 provides only limited evidence of a negative mean return for the 5-1 quintile portfolio, and this pattern is consistent with the CAPM and FF-5 alphas. This finding is corroborated by the statistically insignificant FM risk premium reported in Table 5. By contrast, Panel B of Table 4 reveals novel evidence that the 5-1 quintile portfolio constructed on IS^{ms} exhibits a positive and significant mean return of 0.53%, with a CAPM alpha of 0.40% and an FF-5 alpha of 0.31%. Consistent with these portfolio results, the FM regressions yield a positive and statistically significant risk premium for the medium-term slope. Finally, Panel C of Table 4 shows that the 5-1 quintile return associated with the IS^{lm} factor has a significantly negative mean return of −0.60%, while the corresponding CAPM and FF-5 alphas are −0.60% and −0.64%, respectively. In the cross-sectional tests, the long-term IS is associated with a negative but statistically insignificant risk

premium.

Figure 3 illustrates the time variation in the estimated risk premia for IS across horizons. For IS^s, the estimated risk premium fluctuates around zero over the sample period, taking both positive and negative values. For IS^{ms}, the risk premium is positive for a large portion of the sample, with the LOESS-smoothed series remaining predominantly above zero. For IS^{lm}, the estimated risk premium also varies around zero for most of the sample but becomes more negative in the later period, particularly from 2023 onward. This period coincided with the onset of a restrictive monetary regime and tightening liquidity conditions. In the context of the ICAPM, such monetary constraints signal a specific deterioration in the financial investment opportunity set, driving the risk premium into deeper negative territory. Overall, these time-series patterns are consistent with the risk premia reported in Table 5.

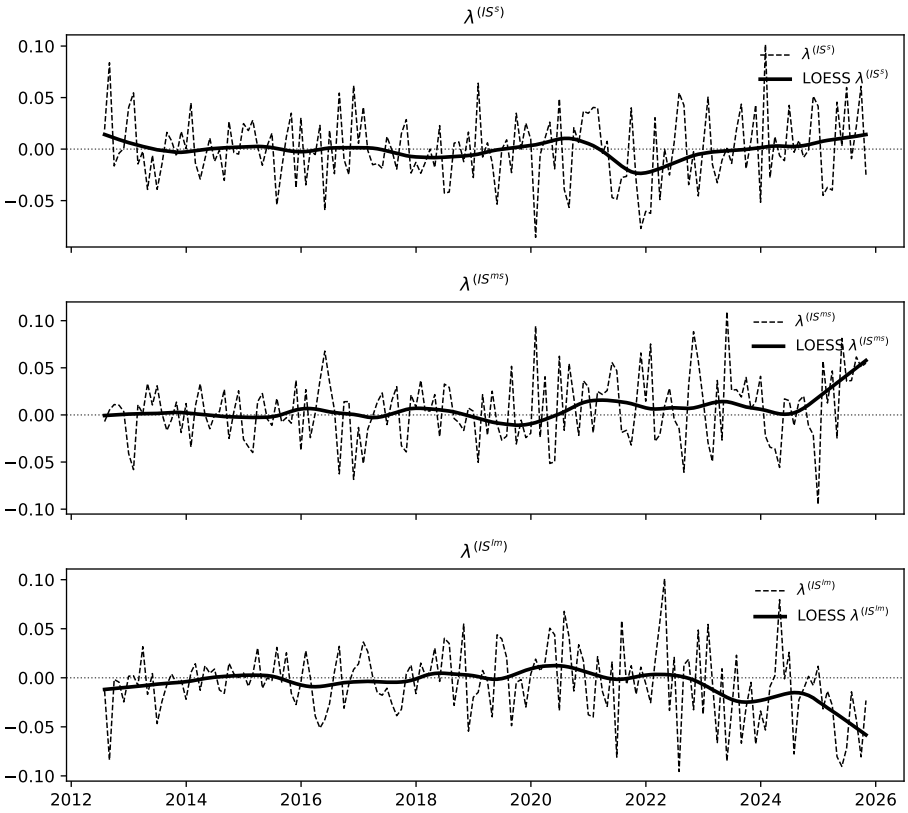

Figure 3: Time-Varying Risk Premium of IS. This figure plots the monthly estimates of the risk premium (λ) for IS^s, IS^{ms}, and IS^{lm}. The risk premium is obtained from monthly FM cross-sectional regressions controlling for the FF-5 factors. The figure reports the raw monthly λ estimates and their corresponding LOESS-smoothed series.

Table 4: Portfolio Statistics Sorted on IS Betas. This table presents the summary statistics of quintile portfolios formed at the end of each month by sorting stocks on their estimated beta loadings with respect to IS^s (Panel A), IS^{ms} (Panel B), and IS^{lm} (Panel C). Portfolios are constructed following the same procedure as in Table 2. Each panel reports six columns: Quintiles 1 through 5 and the 5-1 portfolio. Row 1 reports the average pre-ranking betas used for portfolio formation. Row 2 reports the mean monthly returns in percent, and Rows 3 and 4 report CAPM and FF-5 alphas in percent. t-statistics, based on Newey–West standard errors, are reported in parentheses.

Quintiles	1	2	3	4	5	5-1
Panel A: Portfolios sorting based on β_{IS^s}						
Pre Formation β	-13.83	-4.12	-0.67	4.53	13.11	
Mean Return	0.85	1.10	0.98	1.02	0.84	-0.00
	(2.36)	(3.34)	(3.09)	(3.12)	(2.07)	(-0.01)
CAPM Alpha	-0.37	-0.07	-0.16	-0.15	-0.52	-0.26
	(-2.51)	(-0.76)	(-1.20)	(-1.72)	(-3.04)	(-1.30)
FF-5 Alpha	-0.38	-0.13	-0.20	-0.13	-0.37	-0.12
	(-2.96)	(-1.62)	(-1.57)	(-1.42)	(-2.73)	(-0.80)
Panel B: Portfolios sorting based on $\beta_{IS^{ms}}$						
Pre Formation β	-25.20	-2.86	-1.00	2.73	22.31	
Mean Return	0.61	0.83	0.99	1.19	1.19	0.53
	(1.61)	(2.59)	(3.01)	(3.61)	(3.14)	(1.98)
CAPM Alpha	-0.64	-0.32	-0.20	0.03	-0.07	0.40
	(-3.66)	(-2.65)	(-1.98)	(0.29)	(-0.43)	(1.57)
FF-5 Alpha	-0.51	-0.34	-0.25	0.01	-0.04	0.31
	(-3.68)	(-3.06)	(-3.31)	(0.06)	(-0.24)	(1.20)
Panel C: Portfolios sorting based on $\beta_{IS^{lm}}$						
Pre Formation β	-10.73	-7.64	0.62	5.63	11.39	
Mean Return	1.25	1.22	0.96	0.92	0.64	-0.60
	(3.09)	(3.70)	(2.96)	(2.85)	(1.75)	(-2.23)
CAPM Alpha	-0.11	0.04	-0.21	-0.23	-0.58	-0.60
	(-0.51)	(0.42)	(-1.80)	(-1.34)	(-3.41)	(-1.66)
FF-5 Alpha	-0.02	0.04	-0.26	-0.25	-0.53	-0.64
	(-0.09)	(0.50)	(-2.92)	(-1.59)	(-3.57)	(-2.04)

Table 5: Risk Premium of IS. This table reports the estimated risk premium (λ) for the 5-1 mimicking portfolios of IS^s, IS^{ms}, and IS^{lm}. The risk premium is obtained from monthly FM cross-sectional regressions. The CAPM columns report estimates controlling for the market excess return, and the FF-5 columns report estimates controlling for the FF-5 factors. Reported values are time-series averages of the monthly λ estimates. t-statistics, based on Newey–West standard errors, are shown in parentheses.

	CAPM			FF-5		
	IS^s	IS^{ms}	IS^{lm}	IS^s	IS^{ms}	IS^{lm}
λ	-0.02	0.86	-0.17	-0.05	0.82	-0.44
	(-0.06)	(2.35)	(-0.42)	(-0.18)	(2.64)	(-1.23)

3.4 Implied Kurtosis

Chang et al. [2013] report limited empirical evidence that innovations in kurtosis are priced in the cross section of stock returns. Our analysis partially confirms the findings of Chang et al. [2013], in that the short-term IK factor does not exhibit a pronounced premium. However, our term-structure factors, particularly the ms factor, display a strong premium, which stands in sharp contrast to most existing results in the implied kurtosis literature. Table 6 shows that medium-term IK exhibits significant negative returns: the 5-1 portfolio delivers a monthly mean return of -0.70%, with a CAPM alpha of -0.86% and an FF-5 alpha of -0.77%, all of which are statistically significant. The lm factor exhibits a positive return, with the 5–1 portfolio delivering a mean return of 0.49%, and CAPM and FF-5 alphas of 0.26% and 0.29%, respectively, though the statistical significance is relatively weak.

Table 7 shows that the risk premium of medium-term IK is consistent with its portfolio-level statistics and cumulative returns, and it is significantly negative under both the CAPM and FF-5 models, at –0.83 and –0.88, respectively. For IK^s, the estimated risk premium is 0.31 under CAPM and 0.44 under FF-5. For IK^{lm}, the estimated risk premium is 0.10 under CAPM and 0.34 under FF-5.

Figure 4 illustrates the time variation in the estimated risk premia for implied kurtosis across horizons. For IK^s, the LOESS-smoothed series remains predominantly positive over time. For IK^{ms}, the risk premium varies around zero in the earlier part of the sample but becomes persistently negative in the later period. For IK^{lm}, the estimated risk premium fluctuates around zero for much of the sample and shows an upward movement starting in 2023. The market demands higher compensation not only for general uncertainty (variance) but also for extreme tail risks (kurtosis). Thus, the post-2023 rise reflects a consistent repricing of risk across the term structure in response to economic deterioration.

4 Economic Interpretation

4.1 Predictive Relations and Economic Mechanism

The ICAPM implies that state variables capturing changes in investment opportunities should covary negatively with contemporaneous market returns but positively with future returns. When such variables rise in bad states, investors' marginal utility is high, and the expected compensation for bearing these risks subsequently increases. Empirical studies have documented this sign pattern across different types of state

Table 6: Portfolio Statistics Sorted on IK Betas. This table presents the summary statistics of quintile portfolios formed at the end of each month by sorting stocks on their estimated beta loadings with respect to IK^s (Panel A), IK^{ms} (Panel B), and IK^{lm} (Panel C). Portfolios are constructed following the same procedure as in the IV and IS analyses. Each panel reports six columns: Quintiles 1 through 5 and the 5-1 portfolio. Row 1 reports the average pre-ranking betas used for portfolio formation. Row 2 reports the mean monthly returns in percent, and Rows 3 and 4 report CAPM and FF-5 alphas in percent. t-statistics, based on Newey–West standard errors, are reported in parentheses.

Quintiles	1	2	3	4	5	5-1
Panel A: Portfolios sorting based on β_{IK^s}						
Pre Formation β	-13.62	-2.68	1.52	2.45	13.52	
Mean Return	0.75	1.04	0.97	1.03	0.99	0.20
	(1.90)	(3.14)	(3.03)	(3.19)	(2.60)	(0.79)
CAPM Alpha	-0.56	-0.15	-0.18	-0.11	-0.30	0.09
	(-2.69)	(-1.73)	(-1.64)	(-0.90)	(-2.08)	(0.38)
FF-5 Alpha	-0.41	-0.14	-0.21	-0.17	-0.28	-0.04
	(-2.51)	(-1.67)	(-2.19)	(-1.53)	(-2.45)	(-0.21)
Panel B: Portfolios sorting based on $\beta_{IK^{ms}}$						
Pre Formation β	-20.41	-3.06	0.36	4.40	22.54	
Mean Return	1.31	1.13	1.00	0.86	0.59	-0.70
	(3.48)	(3.48)	(3.04)	(2.65)	(1.51)	(-2.56)
CAPM Alpha	0.06	-0.02	-0.18	-0.31	-0.69	-0.86
	(0.29)	(-0.21)	(-1.76)	(-2.63)	(-3.32)	(-2.61)
FF-5 Alpha	0.10	-0.04	-0.24	-0.33	-0.56	-0.77
	(0.48)	(-0.43)	(-3.00)	(-2.79)	(-3.25)	(-2.32)
Panel C: Portfolios sorting based on $\beta_{IK^{lm}}$						
Pre Formation β	-8.69	-8.03	-1.39	7.29	9.22	
Mean Return	0.69	0.88	0.91	1.25	1.23	0.49
	(1.85)	(2.69)	(2.88)	(3.73)	(3.00)	(1.84)
CAPM Alpha	-0.56	-0.26	-0.24	0.05	-0.14	0.26
	(-3.06)	(-1.61)	(-2.34)	(0.50)	(-0.66)	(0.80)
FF-5 Alpha	-0.50	-0.29	-0.26	0.03	-0.04	0.29
	(-3.10)	(-1.85)	(-3.38)	(0.34)	(-0.22)	(1.08)

Table 7: Risk Premium of IK. This table reports the estimated risk premium (λ) for the 5-1 mimicking portfolios of IK^s, IK^{ms}, and IK^{lm}. The risk premium is obtained from monthly FM cross-sectional regressions. The CAPM columns report estimates controlling for the market excess return, and the FF-5 columns report estimates controlling for the FF-5 factors. Reported values are time-series averages of the monthly estimates of λ. t-statistics, based on Newey–West standard errors, are shown in parentheses.

	CAPM			FF-5		
	IK^s	IK^{ms}	IK^{lm}	IK^s	IK^{ms}	IK^{lm}
λ	0.31	-0.83	0.10	0.44	-0.88	0.34
	(0.89)	(-2.03)	(0.28)	(1.76)	(-2.49)	(1.03)

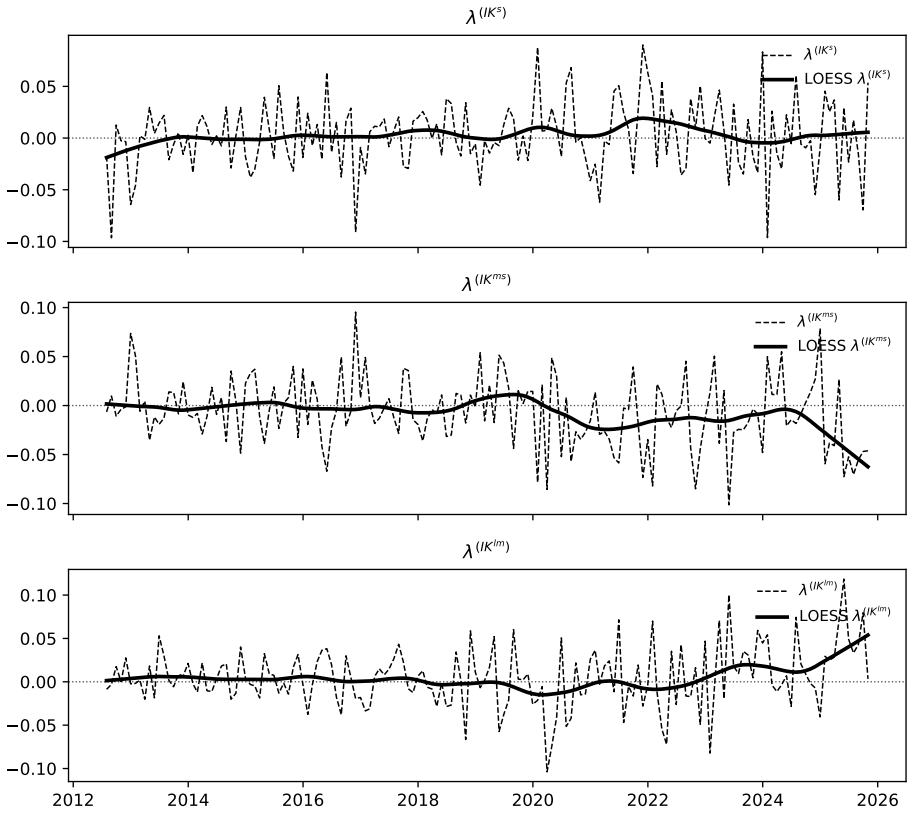

Figure 4: Time-Varying Risk Premium of IK. This figure plots the monthly estimates of the risk premium (λ) for IK^s, IK^{ms}, and IK^{lm}. The risk premium is obtained from monthly FM cross-sectional regressions controlling for the FF-5 factors. The figure reports the raw monthly λ estimates and their corresponding LOESS-smoothed series.

variables. For example, Boons [2016] shows that for real and financial state variables, such as the dividend yield, term spread, and credit spread, sign consistency generally holds: these variables tend to forecast higher future market returns when their contemporaneous realizations are low, and their estimated risk premia share the same sign as the predictive coefficients. In contrast, Ang et al. [2006] report that innovations in the VIX index display a negative λ while positively forecasting future market excess returns, and Gonzalez-Urteaga et al. [2024] document a comparable result for principal components derived from option-implied risk-premium term structures. These findings indicate that for risk factors reflecting market uncertainty or higher-order moments, the price of risk and the predictive coefficients can exhibit opposite signs, reflecting intertemporal adjustments in risk compensation.

We estimate the following time-series regression:

$$MKT_{t,t+h} = \alpha_{j,h} + \beta_{j,h} X_t^j + \varepsilon_{t+h}^j, \tag{7}$$

where $MKT_{t,t+h}$ denotes the compounded market excess return over the next h months, and X_t^j represents one of the implied-moment factors (IV^j, IS^j, or IK^j) for horizon $j \in \{s, ms, lm\}$. The contemporaneous relation corresponds to $h = 0$, and predictive horizons are $h = 1, 3, 6$, and 12 months. All regressions are estimated with Newey–West standard errors using $h - 1$ lags (for the appropriate bandwidth) to adjust for overlapping observations. Table 8 reports the results.

Table 8: Contemporaneous and Prediction Results. This table reports results from time-series regressions examining the contemporaneous and predictive relations between market excess returns and the state variables for IV, IS, and IK across short-, medium-, and long-term horizons (s, ms, lm). In the contemporaneous regressions, the monthly market excess return (MKT_t) is regressed on each state variable observed in the same month ($h = 0$). In the predictive regressions, the cumulative market excess return over the next h months ($h = 1, 3, 6, 12$) is regressed on the lagged state variable from month t. Panel A presents contemporaneous coefficients, while Panel B reports predictive coefficients. For ease of interpretation, coefficients on IS and IK are multiplied by 100, while coefficients on IV are reported in original units. t-statistics, computed using Newey–West standard errors, are shown in parentheses.

	IV			IS (*100)			IK (*100)		
	IV^s	IV^{ms}	IV^{lm}	IS^s	IS^{ms}	IS^{lm}	IK^s	IK^{ms}	IK^{lm}
Panel A: Contemporaneous Results									
	-5.99	-0.54	-0.28	-0.50	0.13	0.68	0.02	0.00	-0.07
	(-5.19)	(-1.48)	(-0.78)	(-1.61)	(0.42)	(1.71)	(0.76)	(0.11)	(-2.26)
Panel B: Prediction Results									
1 *month*	2.42	0.57	0.17	0.41	-0.33	0.04	-0.03	0.03	-0.00
	(1.22)	(1.85)	(0.47)	(1.56)	(-1.15)	(0.09)	(-1.76)	(1.87)	(-0.11)
3 *month*	6.11	1.02	0.35	1.53	-1.38	-0.36	-0.10	0.12	0.02
	(3.00)	(1.34)	(0.43)	(2.99)	(-2.26)	(-0.41)	(-3.56)	(3.75)	(0.32)
6 *month*	9.67	1.33	0.09	2.59	-2.05	-1.20	-0.15	0.16	0.12
	(2.39)	(0.86)	(0.07)	(3.24)	(-2.35)	(-0.70)	(-3.57)	(3.67)	(1.07)
12 *month*	16.46	2.28	1.04	5.37	-3.91	-3.92	-0.33	0.32	0.33
	(2.17)	(0.79)	(0.41)	(3.72)	(-3.23)	(-1.22)	(-3.84)	(4.55)	(1.10)

(i) Implied Variance: The contemporaneous regressions ($h = 0$) show that all variance-related factors carry negative coefficients, indicating that increases in short-, medium-, and long-term implied variance tend to coincide with market declines. The effect is particularly pronounced for IV^s, which exhibits a large and highly significant coefficient, whereas the coefficients on IV^{ms} and IV^{lm} are smaller in magnitude and statistically insignificant. In the predictive regressions, the coefficients on IV^s turn positive across all forecast horizons, consistent with volatility mean reversion and a positive variance risk premium. While the one-month-ahead coefficient is relatively imprecisely estimated, the predictive relationship strengthens and becomes statistically significant at medium and longer horizons. This suggests that high short-term uncertainty forecasts higher subsequent market excess returns over extended horizons. The coefficients on

IV^{ms} and IV^{lm} are also predominantly positive, although generally insignificant.

(ii) Implied Skewness: IS^s shows a negative contemporaneous coefficient, consistent with its negative λ in the cross section. Assets more exposed to declines in IS^s tend to outperform when market downside risk deepens, reflecting their hedging value. This relationship is captured by a negative loading in the contemporaneous regression, which indicates that an increase in the implied skewness of the market return distribution is associated with negative contemporaneous market returns. The predictive coefficient of IS^s is positive and becomes statistically significant at horizons beyond one month, implying that more negative current skewness signals costly downside insurance. The medium-term slope factor, IS^{ms}, exhibits the reverse pattern: its contemporaneous coefficient is positive, but its predictive coefficients turn negative and are statistically significant at medium and long horizons. This suggests that when investors anticipate downside risk at medium horizons relative to the short term, expected future returns fall. The long-term slope, IS^{lm}, displays a similar sign pattern but remains statistically insignificant in the predictive regressions.

(iii) Implied Kurtosis: IK^s shows a positive but statistically insignificant contemporaneous relation with market returns. In the predictive regressions, however, its coefficient is negative and highly significant across all horizons. This negative predictability indicates that high short-term kurtosis, reflecting heightened tail-risk expectations, is associated with lower future market excess returns. In contrast, IK^{ms} exhibits a significantly positive predictive coefficient across all horizons. This sign reversal suggests that distinct economic forces drive the pricing of tail risk at different maturities: while short-term kurtosis levels tend to precede lower returns, a steepening of the medium-term kurtosis slope forecasts higher future market risk premia. At the long horizon, IK^{lm} shows a significant negative contemporaneous association with market returns but does not display significant predictive power.

4.2 Dynamic Interdependencies

The analysis in the previous section documents distinct risk premia across the term structure of implied moments, most notably the sign reversals in risk premia for both implied skewness and kurtosis from short to medium horizon. In this section, we examine whether the medium- and long-term factors contain independent pricing information or primarily reflect the persistence of short-term shocks by estimating VAR models for each implied moment.

We construct three separate VAR systems using monthly observations of the raw state variables in levels, each consisting of the short- (s), medium- (ms), and long-term (lm) factors for a given moment.

Moreover, to study the transmission mechanism between moments, we estimate a fourth system focusing on the interactions among the medium-term factors (IV^{ms}, IS^{ms}, IK^{ms}). For each system, we select the optimal lag length based on the Akaike Information Criterion (AIC). We evaluate the dynamic relationships using two complementary tools: (i) Granger causality tests, which assess whether innovations in one component help predict movements in others, and (ii) Forecast Error Variance Decomposition (FEVD) at a 12-month horizon, which quantifies the contribution of each component to the forecast error variance of the others.

Panel A of Table 9 reports the dynamics of implied volatility. As expected, volatility shocks exhibit strong persistence at short and medium horizons. IV^{s} strongly Granger-causes IV^{ms}, and the variance decomposition indicates that IV^{s} explains a substantial share of the variation in IV^{ms} (85.16%). Both short- and medium-term IV also significantly predict long-term IV. At the same time, the long-term component displays pronounced intrinsic persistence: IV^{lm} does not Granger-cause shorter maturities and retains nearly half (49.40%) of its own forecast error variance. These results indicate that, while volatility shocks propagate across horizons, long-term IV is not merely a mechanical reflection of short-term fluctuations.

Panel B reports the results for implied skewness. The Granger causality tests show that IS^{s} does not significantly predict either IS^{ms} or IS^{lm}. Consistent with this weak predictive linkage, the variance decomposition shows that IS^{lm} retains over 52% of its own forecast error variance at a 12-month horizon. These results suggest that long-horizon IS captures information that is not mechanically inherited from short-term downside movements, providing a structural basis for the empirical finding that medium-term IS commands a risk premium with a different sign from short-term IS.

Panel C reports the corresponding results for implied kurtosis. While short-term IK Granger-causes long-term IK ($p = 0.04$), the forecast error variance decomposition indicates that IK^{s} explains only 8.82% of the variation in IK^{lm}. In contrast, the medium-term factor IK^{ms} explains 54.88% of the variation, while IK^{lm} retains 36.31% of its own forecast error variance. These results indicate that the dynamics of long-term kurtosis are not primarily determined by short-term fluctuations.

Panel D examines the economic mechanism linking these moments by analyzing the interplay within the medium-term horizon. The results reveal a strong causal flow from downside risk (IS^{ms}) to tail risk (IK^{ms}). IS^{ms} is a powerful predictor of IK^{ms}, explaining 79.44% of its forecast error variance. Conversely, we do not find evidence of reverse causality, as IK^{ms} does not Granger cause IS^{ms} ($p = 0.32$). In contrast, IV^{ms} has limited explanatory power for the higher moments. These findings suggest a specific economic narrative: during periods of stress, investors first adjust their expectations regarding fundamental downside

deterioration (IS), which subsequently shapes their pricing of extreme tail events (IK).

The VAR evidence supports the view that the term-structure factors employed in our asset pricing tests are not mechanical transformations of short-term risk. In particular, the medium- and long-term components of higher moments exhibit substantial horizon-specific dynamics that are not fully explained by short-term fluctuations.

Table 9: VAR Granger Causality Tests and Variance Decompositions. This table reports results from Vector Autoregression (VAR) models estimated on the raw state variables. Panels A, B, and C examine the dynamic interactions across the term structure of implied volatility (IV), skewness (IS), and kurtosis (IK), respectively, at short- (s), medium- (ms), and long-term (lm) horizons. Panel D analyzes the interactions among the medium-term components (IV^{ms}, IS^{ms}, IK^{ms}). The left-hand side of each panel reports p-values from Granger causality tests, where smaller values indicate stronger predictive relationships across horizons. The right-hand side reports forecast error variance decompositions at a 12-month horizon, measuring the contribution of shocks at different horizons to the variability of each component.

	Granger Causality (p-values) *Target*			Variance Decomposition (%) *Target*		
Source	IV^s	IV^{ms}	IV^{lm}	IV^s	IV^{ms}	IV^{lm}
Panel A: IV						
IV^s	—	0.0000	0.0212	87.36	85.16	37.55
IV^{ms}	0.0002	—	0.0322	11.08	13.95	13.05
IV^{lm}	0.9713	0.9180	—	1.56	0.88	49.40
Panel B: IS						
	IS^s	IS^{ms}	IS^{lm}	IS^s	IS^{ms}	IS^{lm}
IS^s	—	0.1539	0.2394	93.08	64.20	13.01
IS^{ms}	0.8127	—	0.3908	5.27	32.24	34.90
IS^{lm}	0.7315	0.2647	—	1.65	3.56	52.09
Panel C: IK						
	IK^s	IK^{ms}	IK^{lm}	IK^s	IK^{ms}	IK^{lm}
IK^s	—	0.0388	0.0445	82.28	76.01	8.82
IK^{ms}	0.6657	—	0.1882	15.35	18.96	54.88
IK^{lm}	0.4595	0.1495	—	2.38	5.03	36.31
Panel D: Medium-Term Mechanism						
	IV^{ms}	IS^{ms}	IK^{ms}	IV^{ms}	IS^{ms}	IK^{ms}
IV^{ms}	—	0.3675	0.5512	91.04	10.38	4.12
IS^{ms}	0.9544	—	0.0507	8.11	75.81	79.44
IK^{ms}	0.3171	0.0364	—	0.85	13.81	16.44

4.3 Macroeconomic Drivers of Risk Premia

This section examines the macroeconomic forces underlying the estimated risk premia. We relate the estimated monthly risk premia to a set of macroeconomic state variables that capture distinct sources of risk.

The empirical specification focuses on the levels of risk premia, allowing us to examine how compensation for uncertainty, downside asymmetry, and tail risk varies systematically with the macroeconomic environment. All macroeconomic variables are lagged to mitigate concerns regarding simultaneity and to reflect the information set available to investors at the time of pricing risk.

The macroeconomic variables are chosen to represent different dimensions of the economic environment. (i) Real economic activity is proxied by the Chicago Fed National Activity Index (CFNAI). (ii) Inflationary dynamics are captured by both realized inflation, measured by changes in the Consumer Price Index (CPI), and inflation expectations, measured by the ten-year Breakeven Inflation Rate (T10YIE). (iii) Policy-related uncertainty is measured by the Economic Policy Uncertainty (EPU) index [Baker et al., 2016]. (iv) Financial conditions are represented by the spread between Moody's Baa- and Aaa-rated corporate bond yields (Baa–Aaa), which reflects credit market stress and default risk premia.

We examine the stationarity of all variables using the ADF test. Based on the test results, CFNAI is included in levels. To induce stationarity, we apply first differences to T10YIE and Baa-Aaa. Furthermore, CPI and EPU are transformed using log-differences to represent the inflation rate and relative changes in uncertainty, respectively. The dependent variable is the estimated risk premium of factor k, denoted as λ_t^k. These series are obtained from the FM cross-sectional regressions based on the augmented FF-5 model, as described in Section 3. The empirical model includes the lagged dependent variable to account for persistence, and lagged macroeconomic variables to mitigate simultaneity bias:

$$\lambda_t^k = \alpha^k + \rho\lambda_{t-1}^k + \sum_{m=1}^{M} \beta_m^k X_{t-1}^m + \varepsilon_t^k, \tag{8}$$

where X_{t-1}^m represents the one-period lag of the transformed macroeconomic variable m, β_m^k measures the exposure of the risk premium to macroeconomic conditions, and ρ captures the serial dependence of the risk premium.

Table 10 reports the regression results. First, regarding the short-term variance premium $\lambda^{(IV^s)}$, the significant negative association with inflation expectations (T10YIE) aligns with the inflation risk mechanism documented by Boons et al. [2020]. They show that in the post-2000 period, inflation exhibits a positive covariance with real consumption growth (a positive nominal-real covariance). Consequently, a rise in inflation expectations serves as a signal of improved future economic prospects ("good news"), leading to a reduction in the variance risk premium. Additionally, we observe a positive association with CFNAI,

Table 10: Macro Regression Results. This table reports results from time-series regressions of the estimated risk premia (λ_t^k) on lagged transformed macroeconomic state variables. The predictors include CFNAI (levels), T10YIE and Baa–Aaa (first differences), and CPI and EPU (log-differences). The regression specification is $\lambda_t^k = \alpha^k + \rho\lambda_{t-1}^k + \beta M_{t-1} + \varepsilon_t$, where M_{t-1} denotes the vector of macroeconomic variables. All independent variables are lagged by one month. Reported coefficients are multiplied by 100. t-statistics, computed using Newey–West standard errors, are reported in parentheses. Coefficients significant at the 10% level are shown in bold.

	$\lambda^{(IV^s)}$	$\lambda^{(IV^{ms})}$	$\lambda^{(IV^{lm})}$	$\lambda^{(IS^s)}$	$\lambda^{(IS^{ms})}$	$\lambda^{(IS^{lm})}$	$\lambda^{(IK^s)}$	$\lambda^{(IK^{ms})}$	$\lambda^{(IK^{lm})}$
CFNAI	**0.25**	-0.03	**-0.23**	**-0.39**	0.12	**0.42**	**0.36**	-0.11	**-0.24**
	(2.04)	(-0.23)	**(-1.74)**	**(-2.79)**	(0.93)	**(2.89)**	**(2.66)**	(-0.86)	**(-1.79)**
T10YIE	**-2.79**	**-4.12**	-4.13	-1.50	0.91	**5.68**	0.56	0.53	-3.39
	(-2.46)	**(-2.00)**	(-1.55)	(-0.56)	(0.53)	**(2.57)**	(0.26)	(0.35)	(-1.56)
Baa–Aaa	4.97	-0.24	**-14.98**	**-12.87**	5.31	**11.47**	**13.82**	**-6.66**	-6.23
	(1.23)	(-0.06)	**(-3.71)**	**(-3.15)**	(1.46)	**(2.27)**	**(3.37)**	**(-1.65)**	(-1.44)
CPI	84.66	56.99	-39.28	**-307.81**	**321.91**	-178.17	**325.07**	**-359.44**	43.46
	(0.74)	(0.52)	(-0.34)	**(-2.97)**	**(3.65)**	(-1.36)	**(3.27)**	**(-4.37)**	(0.58)
EPU	**-0.77**	**-0.88**	0.25	-0.13	-0.59	0.05	0.14	0.60	-0.49
	(-1.73)	**(-2.23)**	(0.51)	(-0.27)	(-1.32)	(0.07)	(0.28)	(1.39)	(-0.89)
λ_{t-1}^k	-6.86	3.02	**-11.62**	**-11.13**	-1.42	-4.20	**-23.73**	1.22	7.96
	(-1.05)	(0.46)	**(-2.30)**	**(-1.71)**	(-0.13)	(-0.36)	**(-3.94)**	(0.14)	(1.19)

suggesting that real economic activity plays a significant role in predicting increases in the risk premium on short-term market uncertainty. By contrast, the price of risk associated with the medium-term variance ($\lambda^{(IV^{ms})}$) is not predicted by variables proxying real activity. Instead, behavioral variables such as inflation expectations and EPU predict a negative relation with such risk premium. This evidence indicates that the mechanism extends to the medium horizon: inflationary signals reduce the compensation required for bearing medium-term uncertainty. Finally, $\lambda^{(IV^{lm})}$ is driven by a distinct set of forces. Credit conditions, proxied by the Baa-Aaa spread, carry a negative and statistically significant predictive coefficient. This evidence indicates that the long-term variance premium is closely tied to credit market conditions and financial stress. When credit conditions deteriorate, investors place greater weight on long-term financial uncertainty, leading to a reduction in the compensation for bearing long-term variance risk. Additionally, the negative coefficient on CFNAI suggests a counter-cyclical behavior at the long end.

Next, IS risk captures the asymmetric nature of market returns and connects to macroeconomic conditions that influence the probability and severity of downside events. $\lambda^{(IS^s)}$ is significantly negatively related to CFNAI, Baa-Aaa, and CPI. The negative coefficient on CFNAI indicates that the premium for short-term downside protection becomes more pronounced in periods of weak economic activity. Furthermore, the strong negative response to credit spreads and realized inflation suggests that short-term downside risk is

closely linked to deteriorations in credit conditions and inflationary shocks. By contrast, $\lambda^{(IS^{ms})}$ is unrelated to real activity or credit spreads. Instead, it shows a significant positive correlation with CPI. This implies that the pricing of medium-term downside asymmetry is driven by realized inflation shocks. In this context, high inflation heightens concerns about downside outcomes, leading to a repricing of skewness risk over the medium horizon. Finally, $\lambda^{(IS^{lm})}$ exhibits a reversal in its relationship with credit conditions compared to the short end. Specifically, it is significantly positively related to Baa-Aaa, T10YIE, and CFNAI. The positive coefficient on credit spreads contrasts sharply with the negative coefficient observed at the short horizon. This suggests that persistent credit risks are priced as a long-term constraint, where investors demand higher compensation for bearing long-term downside exposure. The positive association with T10YIE further confirms that long-term skewness risk is shaped by inflation expectations and long-term financial health rather than transient business cycle fluctuations.

While skewness and kurtosis denote different statistical features, they are economically linked as both reflect investor concerns about downside tail risks. Specifically, macroeconomic conditions that influence the probability of adverse outcomes often simultaneously affect their asymmetry and severity. Consequently, skewness and kurtosis premia may respond to similar underlying forces, albeit through distinct channels. Thus, the IK risk premium captures compensation for tail thickness and extreme outcomes in the return distribution. It is significantly related to real activity, inflation, and credit conditions. Its positive relationship with CFNAI contrasts with the behavior of short-term skewness, suggesting that short-term tail risk is not primarily driven by recessionary fears. Instead, when economic conditions are robust, investors demand higher compensation for exposure to rare but severe short-term shocks, even if the perceived probability of such events is low. Additionally, the statistically significant predictive power of CPI and Baa-Aaa indicate that inflationary shocks and tightening financial conditions exacerbate concerns about the severity of immediate tail events. In contrast, $\lambda^{(IK^{ms})}$ shows strong negative correlations with CPI and Baa-Aaa. Finally, $\lambda^{(IK^{lm})}$ is most closely tied to real economic fundamentals. Its significant negative relationship with CFNAI indicates a counter-cyclical behavior. Unlike the short-term premium, which rises in good times, the long-term tail risk premium increases when the real economy weakens. This pattern suggests that $\lambda^{(IK^{lm})}$ reflects concerns about the severity of potential outcomes arising from prolonged economic downturns, rather than transient market shocks.

5 Robustness Checks

5.1 Model Diagnostics and Factor Spanning Tests

In this section, we perform a series of diagnostic tests to evaluate the pricing performance of our implied-moment factors and ensure that their explanatory power is not subsumed by standard risk factors. Following the critique in Lewellen et al. [2010], we report the cross-sectional R^2 and the GRS statistic [Gibbons et al., 1989] to assess the goodness of fit of the model. Furthermore, we conduct spanning tests, as suggested by Barillas and Shanken [2017], to determine whether our factors provide distinct pricing information that is not spanned by the existing benchmark models.

Table 11 reports the diagnostic statistics for the FF-5 benchmark and augmented models using 49 industry portfolios as test assets. We first examine the spanning alphas from time-series regressions of each implied-moment factor on the FF-5 factors. Consistent with our main findings, IV^{ms} and IK^{ms} exhibit statistically significant alphas (t-statistics of -2.05 and -1.99, respectively), confirming that the term structure of implied moments provides incremental information beyond standard systematic risk factors. In terms of cross-sectional goodness of fit, the augmented models consistently deliver higher adjusted R^2 values relative to the FF-5 benchmark (59.8%). Furthermore, the models appear well-specified: the cross-sectional intercept λ_0 (the zero-beta rate) is statistically indistinguishable from zero across all specifications (t-statistics < 1.65), consistent with the theoretical restriction for excess returns. Finally, the GRS tests yield large p-values (ranging from 0.68 to 0.79), indicating no evidence of systematic pricing errors across the industry portfolios.

5.2 Bivariate Model

To further disentangle the independent pricing information of the term-structure factors from the baseline levels, we extend our framework to a bivariate specification. Specifically, we focus on the horizon pairs (s, ms) and (ms, lm) to explicitly control for the underlying risk level when evaluating the slope. For each pair, we construct two tradable portfolios using the double-sorting procedure of Fama and French [2008]. These portfolios serve as mimicking factors in the FM two-pass regressions to estimate the risk premia associated with each state variable.

To illustrate the procedure, consider IV and the pair (s, ms). Define $R_{IV^s_{i,t}/IV^{ms}_{j,t}}$ as the return on the portfolio of stocks in quintile i sorted on β_{IV^s} and, within each quintile i, simultaneously sorted into quintile

Table 11: Model Diagnostics and Spanning Tests. This table reports the diagnostic statistics for the FF-5 model (Benchmark) and models augmented with implied-moment factors (IV, IS, IK). The test assets are the 49 Fama-French industry portfolios. The sample period is from June 1, 2012 to October 31, 2025. The left panel reports the spanning test results, where $t(\alpha_{FF5})$ denotes the t-statistic of the alpha from regressing the specific implied moment factor on the FF-5 factors. The right panel reports the cross-sectional regression diagnostics for the 49 portfolios: Adj. R^2 is the cross-sectional adjusted R^2; $t(\lambda_0)$ is the t-statistic of the cross-sectional intercept (pricing error); and GRS p-val is the p-value of the GRS test for the null hypothesis that the alphas of all 49 portfolios are jointly zero. t-statistics are based on Newey–West standard errors.

	Spanning Test	Model Diagnostics		
Augmented Factor	$t(\alpha_{FF5})$	Adj. R^2	$t(\lambda_0)$	GRS p-val
Benchmark (FF-5)	—	0.598	1.59	0.70
Panel A: IV				
IV^s	-0.79	0.610	1.62	0.74
IV^{ms}	-2.05	0.604	1.35	0.70
IV^{lm}	1.39	0.625	1.59	0.75
Panel B: IS				
IS^s	0.05	0.621	0.75	0.68
IS^{ms}	1.65	0.607	1.52	0.76
IS^{lm}	-1.77	0.623	1.63	0.76
Panel C: IK				
IK^s	0.50	0.620	1.19	0.72
IK^{ms}	-1.99	0.616	1.60	0.79
IK^{lm}	1.43	0.608	1.47	0.73

j based on $\beta_{IV^{ms}}$. The mimicking factor for short-term IV, denoted as $R_{\widehat{IV}_t^s}$, is then constructed as the return difference between the average of $R_{IV^s_{5,t}/IV^{ms}_{\cdot,t}}$ and the average of $R_{IV^s_{1,t}/IV^{ms}_{\cdot,t}}$. More formally,

$$R_{\widehat{IV}_t^s} = \frac{1}{5}\sum_{j=1}^{5} R_{IV^s_{5,t}/IV^{ms}_{j,t}} - \frac{1}{5}\sum_{j=1}^{5} R_{IV^s_{1,t}/IV^{ms}_{j,t}}. \tag{9}$$

Similarly, conditioning on short-term exposures yields the medium-term mimicking factor $R_{\widehat{IV}_t^{ms}}$. Thus, for each state variable, we obtain two mimicking factors, each corresponding to a different horizon. We include each factor individually in the first-stage time-series regressions, specified by (3) for the CAPM and (4) for the FF-5 model. The resulting factor loadings are then used in the second-stage cross-sectional regression (5) to estimate the risk premia. Consequently, we report the risk premia associated with both mimicking factors under both specifications. This procedure is consistently applied to IS and IK. Specifically, we examine the following pairs:

$$(IV_t^s, IV_t^{ms}), (IV_t^{ms}, IV_t^{lm}), (IS_t^s, IS_t^{ms}),$$

$$(IS_t^{ms}, IS_t^{lm}), (IK_t^{s}, IK_t^{ms}), (IK_t^{ms}, IK_t^{lm}).$$

The results from the bivariate analysis are reported in Table 12. These estimates provide a direct test of whether the pricing information in the slope factors (ms and lm) is subsumed by s and ms, respectively.

For IV, the results vary across the term structure. In the (s, ms) specification, both the short-term level IV^s (-0.53, $t = -2.34$) and the medium-term slope IV^{ms} (-0.66, $t = -3.41$) carry statistically significant negative risk premia. This suggests that at the front end of the curve, both level and slope contain independent pricing information. However, in the (ms, lm) specification, while IV^{ms} remains robustly significant (-0.63, $t = -3.46$), the long-term slope IV^{lm} is statistically insignificant (-0.14, $t = -0.53$). This indicates that once the medium-term slope is controlled for, the incremental information provided by the long-term slope is limited.

For IS, the bivariate regressions show that the pricing power is concentrated in the medium-term slope. In the (s, ms) pair, the short-term factor IS^s remains statistically insignificant (0.11, $t = 0.39$) similar to the univariate model, whereas the medium-term slope IS^{ms} retains a large and significant positive premium (0.79, $t = 3.06$). A similar pattern holds for the (ms, lm) pair: IS^{ms} remains significant (0.69, $t = 2.87$), while the long-term factor IS^{lm} yields a statistically insignificant estimate (-0.14, $t = -0.43$). These findings confirm that the positive risk premium documented for skewness is primarily driven by the medium-term slope, with little independent contribution from the long end after controlling for the medium-term factor.

For IK, the evidence also points to the dominance of the medium-term factor. In the (s, ms) specification, the short-term factor IK^s is insignificant (0.10, $t = 0.38$), while IK^{ms} commands a strongly significant negative premium (-0.73, $t = -2.93$). In the (ms, lm) specification, IK^{ms} remains significant (-0.70, $t = -3.00$), whereas IK^{lm} is indistinguishable from zero (0.15, $t = 0.42$).

In summary, this analysis confirms the robustness of our baseline findings. These results reinforce the conclusion that the term structure reflects distinct risk dimensions that are not subsumed by short-term fluctuations.

5.3 Alternative Choice of Test Assets for FM Regressions

The cross-sectional asset pricing equation (5) requires a suitable set of test assets to evaluate the predictive content of the proposed factors. Our main analysis relies on the 49 industry portfolios from Kenneth French's data library.

Table 12: Risk Premium for Double-Sorted Mimicking Factors. This table reports risk premium estimates (λ) for bivariate models based on mimicking factors constructed via double sorting. Stocks are grouped into 5×5 portfolios using conditional double sorts: in each month, stocks are first ranked by the control variable, and subsequently sorted into quintiles within each control portfolio based on the target variable. The mimicking factor is formed as the 5-1 HML spread averaged across the five control portfolios, following the methodology of Fama and French [2008]. Reported coefficients are the time-series averages of monthly risk premium estimates from second-pass FM regressions, with factor loadings estimated each month. Results are shown under the CAPM and FF-5 specifications. t-statistics, based on Newey–West standard errors, are reported in parentheses.

	IV			*IS*			*IK*		
	$\lambda^{(IV^s)}$	$\lambda^{(IV^{ms})}$	$\lambda^{(IV^{lm})}$	$\lambda^{(IS^s)}$	$\lambda^{(IS^{ms})}$	$\lambda^{(IS^{lm})}$	$\lambda^{(IK^s)}$	$\lambda^{(IK^{ms})}$	$\lambda^{(IK^{lm})}$
Panel A: CAPM									
(s, ms)	-0.53	-0.66		0.11	0.79		0.10	-0.73	
	(-2.34)	(-3.41)		(0.39)	(3.06)		(0.38)	(-2.93)	
(ms, lm)		-0.63	-0.14		0.69	-0.14		-0.70	0.15
		(-3.46)	(-0.53)		(2.87)	(-0.43)		(-3.00)	(0.42)
Panel B: FF-5									
(s, ms)	-0.24	-0.49		0.03	0.61		0.23	-0.60	
	(-1.19)	(-2.77)		(0.13)	(2.62)		(1.08)	(-2.73)	
(ms, lm)		-0.46	0.10		0.51	-0.32		-0.66	0.35
		(-2.61)	(0.44)		(2.14)	(-1.12)		(-2.82)	(1.13)

To verify robustness, we also consider two alternative sets of test assets in the second-stage FM regressions. The first alternative set consists of 25 portfolios constructed using a double-sorting procedure based on market beta and the factor beta. Following the methodological approach of Ang et al. [2006], this design is intended to sharpen the cross-sectional dispersion of factor exposures while controlling for market risk. Specifically, in the first step, assets are ranked into quintiles based on their estimated market beta. Within each market-beta quintile, assets are subsequently sorted into quintiles according to their loading on the given state variable, yielding 5×5 value-weighted portfolios. Daily returns for month $t + 1$ are computed as value-weighted averages, using market capitalization at the end of month t as weights. The second alternative is the 25 portfolios formed on book-to-market and operating profitability (BM–OP) from Kenneth French's data library. Book-to-market and profitability have been shown to explain systematic variation in average returns (e.g., Fama and French [1993, 2015]; Hou et al. [2015]).

Table 13 reports estimates of risk premia obtained using the double-sorted portfolios (Panel A) and the BM–OP portfolios (Panel B). Overall, the results provide support for the robustness of the premium of our factors, while the evidence for implied volatility is more sensitive to the choice of test assets. For IS, the medium-term slope consistently commands a large and statistically significant positive risk premium across

Table 13: Risk Premium Under Alternative Test Assets. This table reports risk premium estimates (λ) for IV, IS, and IK. In each case, factor loadings are estimated monthly in the first-stage time-series regressions under CAPM and FF-5, and the second-stage cross-sectional regressions use these loadings to obtain λ. Panel A uses 25 portfolios constructed using a double-sorting procedure. Specifically, stocks are sorted into quintiles by market beta and subsequently sorted into quintiles within each market-beta portfolio based on their loading with respect to the given variable, producing 5×5 value-weighted portfolios. Panel B uses the 25 portfolios formed on book-to-market and operating profitability from Kenneth French's data library. t-statistics, based on Newey–West standard errors, are reported in parentheses.

Panel A: Double-Sorted Test Assets

	IV^s	IV^{ms}	IV^{lm}	IS^s	IS^{ms}	IS^{lm}	IK^s	IK^{ms}	IK^{lm}
CAPM									
λ	0.03	-0.42	0.15	-0.09	0.87	-0.57	0.33	-0.85	0.32
	(0.10)	(-1.96)	(0.55)	(-0.28)	(2.38)	(-1.66)	(1.14)	(-2.58)	(0.84)
FF-5									
λ	-0.20	-0.19	0.19	0.11	0.86	-0.46	0.43	-0.66	0.33
	(-0.71)	(-0.96)	(0.92)	(0.47)	(2.71)	(-1.23)	(1.83)	(-2.30)	(0.76)

Panel B: 25 BM–OP Test Assets

	IV^s	IV^{ms}	IV^{lm}	IS^s	IS^{ms}	IS^{lm}	IK^s	IK^{ms}	IK^{lm}
CAPM									
λ	-0.20	-0.25	0.35	-0.24	0.83	-0.35	0.22	-0.79	0.46
	(-0.77)	(-0.61)	(1.38)	(-0.67)	(2.61)	(-0.99)	(0.52)	(-2.09)	(1.36)
FF-5									
λ	-0.38	0.03	0.41	0.02	0.45	-0.36	0.25	-0.56	0.33
	(-1.82)	(0.09)	(1.99)	(0.08)	(1.71)	(-1.02)	(0.82)	(-2.09)	(1.00)

both sets of test assets. In Panel A, IS^{ms} is 0.87 ($t = 2.38$) under CAPM and 0.86 ($t = 2.71$) under FF-5. Similarly, in Panel B, the premium remains positive and significant (0.83 under CAPM and 0.45 under FF-5). The short- and long-term factors generally do not exhibit significant premia. For IK, the medium-term slope consistently exhibits a significant negative risk premium. In Panel A, the estimates for IK^{ms} are -0.85 ($t = -2.58$) and -0.66 ($t = -2.30$) under CAPM and FF-5, respectively. This pattern holds in Panel B, where the factor remains significant across specifications (-0.79 and -0.56). In contrast, the results for IV are less stable. While IV^{ms} shows a significant negative premium under CAPM in Panel A (-0.42, $t = -1.96$), it loses statistical significance after controlling for the FF-5 factors or when using the BM–OP portfolios.

5.4 Excluding Small-Cap Stocks

Small-cap stocks play a key role in cross-sectional asset pricing patterns. Ang et al. [2006] find that small-cap stocks are highly sensitive to economic uncertainty, and excluding them reduces both the economic

magnitude and statistical significance of portfolio returns constructed on ΔVIX beta exposures. Likewise, Hou et al. [2015] show that small stocks exhibit strong investment and profitability anomalies, motivating size-based sorting in the q-factor model to limit disproportionate effects from small-cap firms.

In this robustness check, we examine whether our main findings remain after excluding small-cap stocks. At the end of each month, we rank the cross section by market capitalization and remove all firms in the bottom 20%. Portfolios are then constructed from the remaining stocks, following the same sorting procedures as in the baseline analysis, and the estimation of risk premia is carried out through the two-pass FM regressions.

Table 14 reports the estimated risk premia from this robustness test. For IV, IV^{ms} retains its negative risk premium, which is statistically significant under CAPM (-0.86, $t = -3.95$) and FF-5 (-0.48, $t = -1.96$). In contrast, IV^{s} loses statistical significance ($t = -1.60$ and $t = -0.77$). For IS, IS^{ms} continues to command a strong and significant positive premium (0.84, $t = 2.21$ under CAPM; 0.81, $t = 2.57$ under FF-5). The short- and long-term factors remain statistically insignificant. For IK, the pricing power of the medium-term slope remains intact. IK^{ms} exhibits a significant negative premium across both models (-0.81, $t = -2.06$ under CAPM; -0.86, $t = -2.64$ under FF-5). The long-term factor remains insignificant here. In summary, excluding small-cap stocks does not alter our main conclusion: the medium-term slope factors contain robust pricing information that permeates the broader cross section of stocks.

Table 14: Risk Premium Excluding Small-Cap Stocks. This table reports estimated risk premium (λ) of IV, IS, and IK after excluding the bottom 20% of stocks by market capitalization. Test portfolios are re-formed each month from the remaining stocks, and risk premia are obtained from FM two-pass regressions. Reported coefficients are the time-series averages of monthly estimates under the CAPM and FF-5 specifications. t-statistics, based on Newey–West standard errors, are reported in parentheses.

	IV^{s}	IV^{ms}	IV^{lm}	IS^{s}	IS^{ms}	IS^{lm}	IK^{s}	IK^{ms}	IK^{lm}
CAPM									
λ	-0.43	-0.86	0.10	-0.03	0.84	-0.14	0.17	-0.81	0.05
	(-1.60)	(-3.95)	(0.30)	(-0.07)	(2.21)	(-0.34)	(0.49)	(-2.06)	(0.13)
FF-5									
λ	-0.21	-0.48	0.32	-0.04	0.81	-0.48	0.36	-0.86	0.33
	(-0.77)	(-1.96)	(1.24)	(-0.13)	(2.57)	(-1.32)	(1.43)	(-2.64)	(0.99)

5.5 Robustness to Liquidity and Momentum Effects

Prior studies show that liquidity risk is an important determinant of the cross section of stock returns (e.g., Pástor and Stambaugh, 2003, Acharya and Pedersen, 2005), raising the possibility that liquidity-related

factors may overlap with the information content of our implied moments. Momentum, documented by Jegadeesh and Titman [1993] and incorporated into the Carhart [1997] model, is another widely recognized driver of cross-sectional returns. We therefore augment our pricing models with controls for liquidity and momentum in separate exercises.

5.5.1 Liquidity

Our pricing strategy uses daily data through October 2025, while the Pastor–Stambaugh liquidity factor from Lubos Pastor's library is monthly and ends in December 2024. To maintain consistency across empirical exercises, we construct a daily liquidity measure up to October 2025 by replicating Pástor and Stambaugh [2003] with our dataset. For completeness, we briefly review the methodology:

(1) Compute the stock-level liquidity factor of Pástor and Stambaugh [2003] by running the following regression for each stock i at each month-end using that month's daily data:

$$r_{i,d} = \theta_i + \phi_i r_{i,d-1} + \gamma_i \,\mathrm{sign}(r_{i,d-1})\, V_{i,d-1} + \varepsilon_{i,d}, \tag{10}$$

where $r_{i,d}$ is the daily return of stock i on day d, $\mathrm{sign}(r_{i,d-1})$ is the sign of the previous day's return, and $V_{i,d}$ is the standardized trading volume, measured as the ratio of daily shares traded to shares outstanding. The coefficient γ_i is estimated separately for each month and provides a monthly measure of stock-level sensitivity to liquidity, denoted as LIQ_i.

(2) Construct the portfolio using a standard double-sorting procedure [Fama and French, 2008]. At each month-end, stocks are first sorted into quintiles by LIQ_i, and then within each quintile, they are further sorted into quintiles by the state variable (e.g., IV^s), forming 25 value-weighted portfolios. The mimicking factor is the average return spread between the highest and lowest quintiles across the five liquidity groups.

Following Ang et al. [2006], we construct the liquidity factor in a manner consistent with the state variables analyzed above. Table 15 reports the estimated risk premia for IV, IS, and IK after controlling for liquidity. The results remain largely unchanged in sign and magnitude. The medium-term slopes (IV^{ms}, IS^{ms}, IK^{ms}) retain their statistical significance across specifications, indicating that our main findings are robust to liquidity effects, although the statistical significance of IV^s is weaker under the FF-5 specification.

Table 15: Risk Premium Controlling for Liquidity. This table reports estimated risk premium (λ) for IV, IS, and IK after controlling for stock-level liquidity. At each month-end, stocks are double-sorted into quintiles by liquidity sensitivity and the state variable to form portfolios that construct the mimicking factors. Coefficients are time-series averages of monthly FM two-pass regressions under CAPM and FF-5. Rows denote short (s), medium (ms), and long (lm) horizons, and columns give the corresponding risk premium for IV, IS, and IK. t-statistics, based on Newey–West standard errors, are reported in parentheses.

	IV^s	IV^{ms}	IV^{lm}	IS^s	IS^{ms}	IS^{lm}	IK^s	IK^{ms}	IK^{lm}
CAPM									
λ	-0.53	-0.76	-0.05	-0.07	0.84	-0.16	0.26	-0.75	0.28
	(-2.15)	(-4.13)	(-0.17)	(-0.25)	(3.04)	(-0.44)	(0.90)	(-2.63)	(0.81)
FF-5									
λ	-0.28	-0.48	0.16	-0.09	0.75	-0.41	0.37	-0.75	0.48
	(-1.22)	(-2.48)	(0.76)	(-0.36)	(2.60)	(-1.31)	(1.48)	(-2.51)	(1.55)

5.5.2 Momentum

The momentum factor is included as an additional factor in the CAPM and FF-5 specifications, using the Carhart [1997] momentum factor obtained from Kenneth French's data library. Table 16 shows that the estimated risk premia for IS and IK remain largely unchanged after controlling for momentum, indicating that the pricing of higher-order moments is robust. For IV, while the estimates maintain their signs, the statistical significance becomes weaker under the FF-5 specification.

Table 16: Risk Premium Controlling for Momentum. This table reports the estimated risk premia (λ) for IV, IS, and IK with the Carhart [1997] momentum factor added to the CAPM and FF-5 specifications. Coefficients are time-series averages of monthly FM two-pass regressions. Rows denote short (s), medium (ms), and long (lm) horizons, and columns give the corresponding premium for IV, IS, and IK. t-statistics, based on Newey–West standard errors, are reported in parentheses.

	IV^s	IV^{ms}	IV^{lm}	IS^s	IS^{ms}	IS^{lm}	IK^s	IK^{ms}	IK^{lm}
CAPM									
λ	-0.56	-0.87	0.20	0.09	0.87	-0.42	0.30	-0.92	0.27
	(-2.07)	(-3.33)	(0.69)	(0.26)	(2.38)	(-1.02)	(0.87)	(-2.26)	(0.77)
FF-5									
λ	-0.41	-0.44	0.40	0.04	0.75	-0.51	0.26	-0.88	0.37
	(-1.51)	(-1.63)	(1.53)	(0.14)	(2.19)	(-1.36)	(0.91)	(-2.27)	(1.08)

6 Conclusion

This paper investigates the term structure of risk-neutral moments—variance, skewness, and kurtosis—and their role in pricing the cross section of stock returns. Unlike prior literature that predominantly focuses

on short-term implied moments, we isolate the pricing content of the maturity dimension by constructing slope factors representing the spread between medium-term (six-month) and long-term (twelve-month) expectations relative to the short end. These moment spreads serve as effective proxies for the incremental distributional risks priced by investors, capturing information distinct from the baseline risk levels.

Our empirical analysis documents strong horizon dependence in the pricing of higher-order risks. While implied variance commands a persistent negative risk premium across both short and medium horizons, the pricing of skewness and kurtosis exhibits significant horizon dependence. We find that the pricing power for these higher-order moments is concentrated in the medium-term slope rather than the short-term level. Specifically, the medium-term skewness slope commands a significant positive risk premium, while the medium-term kurtosis slope carries a significant negative premium. These findings suggest that investors distinguish between immediate risks and persistent shifts in the return distribution, requiring distinct compensation for exposure to medium-term deteriorations in downside asymmetry and tail thickness.

We identify economic mechanisms supporting these asset pricing patterns through multiple perspectives. First, consistent with the ICAPM framework, the estimated risk premia align with the predictive content of these factors. Second, the independence of these term-structure factors is rigorously confirmed. Spanning tests and double-sorted portfolios demonstrate that the slope factors contain unique pricing information orthogonal to the short-term baseline. VAR analysis shows that the medium- and long-term slope factors exhibit intrinsic persistence that is not mechanically driven by short-term shocks.

We link these horizon-specific risk premia to distinct macroeconomic fundamentals. Our analysis reveals a dichotomy in the drivers of risk compensation: short-term variance premia are primarily driven by real economic activity and business cycle fluctuations, whereas the premia associated with medium-term skewness and kurtosis are linked to monetary conditions, inflation, and credit spreads. This indicates that while short-term volatility reflects immediate economic conditions, the term structure of higher-order moments captures deeper concerns regarding the underlying macro-financial environment and policy uncertainties.

Overall, this paper highlights the critical importance of the investment horizon dimension in understanding risk premia. Relying solely on short-term implied moments obscures the rich pricing information contained in the term structure. By decomposing risk expectations across maturities, we provide a more comprehensive view of how investors price the temporal evolution of variance, downside risk, and tail events. Future research may extend this framework to other asset classes or international markets to further explore the global pricing of horizon-dependent distributional risks.

Appendix A Extracting Risk-Neutral Moments From Options

Let S_0 denote the current underlying asset price, r the annualized risk-free rate, q the annualized dividend yield, T the option maturity (in years), and K the strike price. Let $C(K)$ and $P(K)$ denote observed market prices of call and put options at time t, with maturity T and strike K. For convenience, we set $\tau = T$. The option dataset is first filtered to retain out-of-the-money (OTM) options with economically meaningful prices: calls with $K > S_0$ and $C(K) > 0.1$, and puts with $K < S_0$ and $P(K) > 0.1$. Arbitrage-free conditions are imposed by requiring

$$C(K) \geq \max(S_0 - Ke^{-r\tau}, 0), \quad P(K) \geq \max(Ke^{-r\tau} - S_0, 0).$$

For each selected option, the implied volatility $\sigma(K)$ is computed by numerically inverting the Black–Scholes formula:

$$C(K) = BS\big(S_0, K, r, q, \tau, \sigma(K)\big), \quad P(K) = BS\big(S_0, K, r, q, \tau, \sigma(K)\big),$$

where $BS(\cdot)$ denotes the standard Black–Scholes pricing function. The inversion is performed using the Newton–Raphson method, with the Brent method as a fallback in case of non-convergence.

The obtained IVs from call and put options are then merged and sorted by moneyness, K/S_0. A cubic spline interpolation is applied to construct a continuous IV curve $\sigma(K/S_0)$, while values outside the interpolation boundaries are held constant at the nearest boundary IV.

Finally, option prices are computed on a dense moneyness grid $K_i = (\text{moneyness})_i \cdot S_0$. For $K_i > S_0$, call prices are computed using the interpolated IV, and for $K_i < S_0$, put prices are computed similarly:

$$C(K_i) = BS\big(S_0, K_i, r, q, \tau, \sigma(K_i/S_0)\big), \quad P(K_i) = BS\big(S_0, K_i, r, q, \tau, \sigma(K_i/S_0)\big).$$

This procedure ensures a smooth, arbitrage-free IV surface and allows the computation of option prices across a dense range of strikes for subsequent model-free moment extraction.

Subsequently, numerical integration is performed to compute the intermediate integrals for moment calculations, with the integrand kernel functions for call options defined as:

$$V_c(K) = \frac{2[1 - \ln(K/S_0)]}{K^2},$$

$$W_c(K) = \frac{6\ln(K/S_0) - 3[\ln(K/S_0)]^2}{K^2},$$

$$X_c(K) = \frac{12[\ln(K/S_0)]^2 - 4[\ln(K/S_0)]^3}{K^2}$$

For put options, the integrand kernel functions are:

$$V_p(K) = \frac{2[1 + \ln(S_0/K)]}{K^2},$$

$$W_p(K) = \frac{6\ln(S_0/K) + 3[\ln(S_0/K)]^2}{K^2},$$

$$X_p(K) = \frac{12[\ln(S_0/K)]^2 + 4[\ln(S_0/K)]^3}{K^2}$$

Numerical integrals for call ($K > S_0$) and put ($K < S_0$) regions are calculated respectively to obtain the quantities:

$$V = e^{r\tau}\left[\int_0^{S_0} V_p(K)P(K)dK + \int_{S_0}^{\infty} V_c(K)C(K)dK\right]$$

$$W = e^{r\tau}\left[\int_{S_0}^{\infty} W_c(K)C(K)dK - \int_0^{S_0} W_p(K)P(K)dK\right]$$

$$X = e^{r\tau}\left[\int_0^{S_0} X_p(K)P(K)dK + \int_{S_0}^{\infty} X_c(K)C(K)dK\right]$$

Finally, given the integrals above, the risk-neutral moments are calculated as follows. First, define the risk-neutral mean of the returns as:

$$\mu = e^{r\tau} - 1 - \frac{V}{2} - \frac{W}{6} - \frac{X}{24}$$

Then the risk-neutral variance is given by:

$$\text{Variance} = V - \mu^2$$

The risk-neutral skewness is:

$$\text{Skewness} = \frac{W - 3\mu V + 2\mu^3}{(\text{Variance})^{3/2}}$$

The risk-neutral kurtosis is:

$$\text{Kurtosis} = \frac{X - 4\mu W + 6\mu^2 V - 3\mu^4}{(\text{Variance})^2}$$

Appendix B Theoretical Derivation for Term Structure Factors

Appendix B.1 Setup and Definitions

We assume the underlying asset price follows a strictly stationary log-return process $\{r_t\}$ with finite fourth moments. To ensure consistency in notation and improve readability, we denote the variance, skewness, and excess kurtosis operators as $\text{Var}(\cdot)$, $\text{Skew}(\cdot)$, and $\text{Kurt}(\cdot)$, respectively.

We focus on three specific horizons corresponding to the empirical analysis: 1 month, 6 months, and 12 months. Let $R^{(1)} = r_1$ denote the realized return over the first month. The cumulative returns over the six-month and twelve-month horizons are defined as $R^{(6)} = \sum_{i=1}^{6} r_i$ and $R^{(12)} = \sum_{i=1}^{12} r_i$, respectively. In the main text, the moments of $R^{(1)}$ serve as the short-term baseline factors (denoted as s).

Our objective is to capture the *incremental* risk information specific to the medium and long horizons. Theoretically, this corresponds to the moments of the forward returns between these maturities. We define these incremental returns as:

$$\widetilde{R}^{(6,1)} \equiv \sum_{i=2}^{6} r_i = R^{(6)} - R^{(1)}, \qquad \widetilde{R}^{(12,6)} \equiv \sum_{i=7}^{12} r_i = R^{(12)} - R^{(6)}. \tag{11}$$

These variables correspond to the ideal state variables underlying our medium slope (ms) and long slope (lm) factors, respectively.

We seek to recover the statistical moments of these incremental returns, which we denote as the Incremental Moments:

- Incremental Variance: $\mathcal{V}^{(6,1)} \equiv \text{Var}(\widetilde{R}^{(6,1)})$ and $\mathcal{V}^{(12,6)} \equiv \text{Var}(\widetilde{R}^{(12,6)})$.
- Incremental Skewness: $\mathcal{S}^{(6,1)} \equiv \text{Skew}(\widetilde{R}^{(6,1)})$ and $\mathcal{S}^{(12,6)} \equiv \text{Skew}(\widetilde{R}^{(12,6)})$.
- Incremental Kurtosis: $\mathcal{K}^{(6,1)} \equiv \text{Kurt}(\widetilde{R}^{(6,1)})$ and $\mathcal{K}^{(12,6)} \equiv \text{Kurt}(\widetilde{R}^{(12,6)})$.

However, these incremental moments are not directly observable from current option prices, which only provide the implied moments for cumulative horizons (i.e., moments of $R^{(1)}$, $R^{(6)}$, and $R^{(12)}$). To bridge this gap, we construct our empirical state variables as the Differenced Moments (spreads). For the medium-term horizon (ms), the observed slope factors for variance (IV), skewness (IS), and kurtosis (IK) are constructed as:

$$IV^{ms} = \text{Var}(R^{(6)}) - \text{Var}(R^{(1)}), \tag{12}$$

$$IS^{ms} = \text{Skew}(R^{(6)}) - \text{Skew}(R^{(1)}), \tag{13}$$

$$IK^{ms} = \text{Kurt}(R^{(6)}) - \text{Kurt}(R^{(1)}). \tag{14}$$

The long-term slope factors (lm) are defined analogously as the spread between the 12-month and 6-month moments. The following sections analyze the relationship between these observable spreads (Eq. 12-14) and the theoretical incremental moments.

Appendix B.2 Variance: Exact Additivity

We start by examining the relationship between the observed variance spread and the incremental variance. We consider the general benchmark relevant for financial returns where the log-return process $\{r_t\}$ is serially uncorrelated, albeit possibly dependent (e.g., exhibiting conditional heteroskedasticity or stochastic volatility). Specifically, we assume:

$$\text{Cov}(r_t, r_{t-k}) = 0 \quad \text{for all } k \neq 0. \tag{15}$$

This assumption preserves the martingale difference property of returns while allowing for time-varying second moments, $\text{Var}(r_t) = \sigma_t^2$.

Under this assumption, the variance of a cumulative return over any horizon n satisfies exact additivity:

$$\text{Var}(R^{(n)}) = \text{Var}\left(\sum_{t=1}^{n} r_t\right) = \sum_{t=1}^{n} \text{Var}(r_t) = \sum_{t=1}^{n} \sigma_t^2. \tag{16}$$

Applying this property to our construct for the medium slope factor (IV^{ms}), the differenced variance is:

$$\begin{aligned} IV^{ms} &= \mathrm{Var}(R^{(6)}) - \mathrm{Var}(R^{(1)}) \\ &= \sum_{i=1}^{6} \sigma_i^2 - \sigma_1^2 \\ &= \sum_{i=2}^{6} \sigma_i^2. \end{aligned} \tag{17}$$

Comparing this result with the definition of the incremental variance $\mathcal{V}^{(6,1)}$ in Appendix B.1:

$$\mathcal{V}^{(6,1)} \equiv \mathrm{Var}(\widetilde{R}^{(6,1)}) = \mathrm{Var}\left(\sum_{i=2}^{6} r_i\right) = \sum_{i=2}^{6} \sigma_i^2. \tag{18}$$

Thus, under the standard assumption of serial uncorrelation, we obtain the identity:

$$IV^{ms} = \mathcal{V}^{(6,1)}. \tag{19}$$

The same logic applies to the long-term slope factor, where $IV^{lm} = \mathcal{V}^{(12,6)}$. This result demonstrates that, for variance, the slope factors constructed from option-implied moments provide a mathematically precise measure of the forward expected variance over the corresponding medium and long horizons, unaffected by the baseline short-term volatility level.

Appendix B.3 Higher Moments: Relationship between Spreads and Incremental Risks

Unlike variance, standardized moments such as skewness and kurtosis do not satisfy exact additivity. This is because standardized moments are nonlinear functions of cumulants; while cumulants are additive across independent blocks (or under specific dependence structures), standardization by the variance destroys this property. Consequently, for skewness and kurtosis, the observed term structure spreads defined in Appendix B.1 do not strictly equal the theoretical incremental moments:

$$IS^{ms} \neq \mathcal{S}^{(6,1)}, \tag{20}$$

$$IK^{ms} \neq \mathcal{K}^{(6,1)}. \tag{21}$$

However, the discrepancy between the observed spread and the incremental moment is informative. We

can analytically decompose this relationship to understand the information content of our constructed factors. Taking skewness as an example, the relationship between the incremental skewness $\mathcal{S}^{(6,1)}$ and our empirical factor IS^{ms} is given by:

$$IS^{ms} = \mathcal{S}^{(6,1)} - \underbrace{\left[\Delta_{\text{Norm}} + \Delta_{\text{Dependence}}\right]}_{\text{Adjustment Terms}}, \tag{22}$$

where the adjustment terms capture two specific market dynamics:

- Nonlinear Variance Normalization (Δ_{Norm}): This term arises because the normalization factor for the six-month horizon differs from that of the sub-periods due to volatility heterogeneity.
- Intertemporal Dependence ($\Delta_{\text{Dependence}}$): This term involves higher-order cross-cumulants (e.g., $\text{Cov}(r_t^2, r_{t+k})$) that link the short-term returns to the medium-term incremental returns.

Specifically, in the general case with finite fourth moments, the spread IS^{ms} captures not only the marginal contribution of returns between months 2 and 6 (the pure incremental risk) but also the interaction between short-term volatility shocks and subsequent return distributions. As shown in the derivation of the skewness gap (referencing the property that $\Delta_{\text{Dependence}}$ contains terms such as $\text{Cov}(r_1^2, \widetilde{R}^{(6,1)})$), non-zero adjustment terms provide direct evidence of asymmetric volatility dynamics and volatility clustering.

Therefore, while IS^{ms} and IK^{ms} are not direct measures of forward moments in a strict mathematical sense, they serve as valid proxies for the *slope* of the term structure. A significant premium on IS^{ms} implies that the market prices the joint information of incremental downside risk and the persistence of volatility dynamics across horizons.

References

Viral V Acharya and Lasse Heje Pedersen. Asset pricing with liquidity risk. *Journal of financial Economics*, 77(2):375–410, 2005.

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

Vikas Agarwal, Gurdip Bakshi, and Joop Huij. Do higher-moment equity risks explain hedge fund returns? Technical report, CFR working paper, 2009.

Andrew Ang, Robert J Hodrick, Yuhang Xing, and Xiaoyan Zhang. The cross-section of volatility and expected returns. *The journal of finance*, 61(1):259–299, 2006.

Scott R Baker, Nicholas Bloom, and Steven J Davis. Measuring economic policy uncertainty. *The quarterly journal of economics*, 131(4):1593–1636, 2016.

Gurdip Bakshi and Dilip Madan. Spanning and derivative-security valuation. *Journal of financial economics*, 55(2):205–238, 2000.

Gurdip Bakshi, Charles Cao, and Zhiwu Chen. Empirical performance of alternative option pricing models. *The Journal of finance*, 52(5):2003–2049, 1997.

Gurdip Bakshi, Nikunj Kapadia, and Dilip Madan. Stock return characteristics, skew laws, and the differential pricing of individual equity options. *The Review of Financial Studies*, 16(1):101–143, 2003.

Francisco Barillas and Jay Shanken. Which alpha? *The Review of Financial Studies*, 30(4):1316–1338, 2017.

David S Bates. Jumps and stochastic volatility: Exchange rate processes implicit in deutsche mark options. *The Review of Financial Studies*, 9(1):69–107, 1996.

David S Bates. Us stock market crash risk, 1926–2010. *Journal of Financial Economics*, 105(2):229–259, 2012.

Martijn Boons. State variables, macroeconomic activity, and the cross section of individual stocks. *Journal of Financial Economics*, 119(3):489–511, 2016.

Martijn Boons, Fernando Duarte, Frans De Roon, and Marta Szymanowska. Time-varying inflation risk and stock returns. *Journal of Financial Economics*, 136(2):444–470, 2020.

Paul Borochin, Hao Chang, and Yangru Wu. The information content of the term structure of risk-neutral skewness. *Journal of Empirical Finance*, 58:247–274, 2020.

John Y Campbell. Intertemporal asset pricing without consumption data. *American Economic Review*, 83 (3):487–512, 1993.

M.M. Carhart. On persistence in mutual fund performance. *Journal of Finance*, pages 57–82, 1997.

Peter Carr and Dilip Madan. Optimal positioning in derivative securities. 2001.

Peter Carr and Liuren Wu. Variance risk premiums. *The Review of Financial Studies*, 22(3):1311–1341, 2009.

Bo-Young Chang, Peter Christoffersen, Kris Jacobs, and Gregory Vainberg. Option-implied measures of equity risk. *Review of Finance*, 16(2):385–428, 2012.

Bo Young Chang, Peter Christoffersen, and Kris Jacobs. Market skewness risk and the cross section of stock returns. *Journal of Financial Economics*, 107(1):46–68, 2013.

Hoyong Choi, Philippe Mueller, and Andrea Vedolin. Bond variance risk premiums. *Review of Finance*, 21 (3):987–1022, 2017.

Jennifer Conrad, Robert F Dittmar, and Eric Ghysels. Ex ante skewness and expected stock returns. *The Journal of Finance*, 68(1):85–124, 2013.

Pasquale Della Corte, Roman Kozhan, and Anthony Neuberger. The cross-section of currency volatility premia. *Journal of Financial Economics*, 139(3):950–970, 2021.

Ian Dew-Becker, Stefano Giglio, Anh Le, and Marius Rodriguez. The price of variance risk. *Journal of Financial Economics*, 123(2):225–250, 2017.

Robert F Dittmar. Nonlinear pricing kernels, kurtosis preference, and evidence from the cross section of equity returns. *The Journal of Finance*, 57(1):369–403, 2002.

Eugene F Fama and Kenneth R French. Common risk factors in the returns on stocks and bonds. *Journal of financial economics*, 33(1):3–56, 1993.

Eugene F Fama and Kenneth R French. Dissecting anomalies. *The journal of finance*, 63(4):1653–1678, 2008.

Eugene F Fama and Kenneth R French. A five-factor asset pricing model. *Journal of financial economics*, 116(1):1–22, 2015.

Michael R Gibbons, Stephen A Ross, and Jay Shanken. A test of the efficiency of a given portfolio. *Econometrica: Journal of the Econometric Society*, pages 1121–1152, 1989.

Ana Gonzalez-Urteaga, Gonzalo Rubio, Pedro Serrano, and Antoni Vaello-Sebastià. The term structure of the expected market risk premium, and the cross-section of stock returns. *Available at SSRN 5018422*, 2024.

Steven L Heston. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The review of financial studies*, 6(2):327–343, 1993.

Kewei Hou, Chen Xue, and Lu Zhang. Digesting anomalies: An investment approach. *The Review of Financial Studies*, 28(3):650–705, 2015.

Narasimhan Jegadeesh and Sheridan Titman. Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of finance*, 48(1):65–91, 1993.

Jonathan Lewellen, Stefan Nagel, and Jay Shanken. A skeptical appraisal of asset pricing tests. *Journal of Financial economics*, 96(2):175–194, 2010.

Junye Li, Lucio Sarno, and Gabriele Zinna. Skewness risk premia and the cross-section of currency returns. *Available at SSRN 4580189*, 2023.

Robert C Merton. An intertemporal capital asset pricing model. *Econometrica: Journal of the Econometric Society*, pages 867–887, 1973.

L'uboš Pástor and Robert F Stambaugh. Liquidity risk and expected stock returns. *Journal of Political economy*, 111(3):642–685, 2003.

Yuhang Xing, Xiaoyan Zhang, and Rui Zhao. What does the individual option volatility smirk tell us about future equity returns? *Journal of Financial and Quantitative Analysis*, 45(3):641–662, 2010.