The Factor Structure of 0DTE Option Returns

Christian Dorion* Piotr Orłowski† Yuhan Song‡§

ABSTRACT

Zero-days-to-expiration (0DTE) options have become half of SPX trading volume. We ask which systematic risks their returns compensate and whether their prices are efficient. Six factors: three realized return moments and three variance-dynamics components, span 30-minute SPX option returns at maturities from 45 minutes up to 15 days and price 0DTEs jointly with longer-dated contracts. Warehousing gamma and absorbing jump-skewness exposure earn premiums; the variance-component premiums differ but sum to zero intraday. Remaining alpha concentrates in 0DTEs, yet a factor-neutral strategy becomes infeasible under minimal transaction costs. Option alphas fall as inventories build and, secondarily, intermediary balance sheet conditions tighten.

JEL Classification: G10; G12; G13.

Keywords: Option Returns; Factor Models; 0DTE; Tail risk.

*HEC Montréal and Canadian Derivatives Institute, christian.dorion@hec.ca.

†HEC Montréal and Canadian Derivatives Institute, piotr.orlowski@hec.ca.

‡HEC Montréal, yuhan.song@hec.ca.

§We thank Diego Amaya (discussant), Torben Andersen, Gustavo Freire (discussant), Kris Jacobs, Robert Korajczyk, Neil Pearson, Viktor Todorov, Liuren Wu and seminar participants at the FMA/Cboe 2024 Conference on Derivatives and Volatility, the 17th SoFiE Conference, the 2026 Derivatives and Asset Pricing Conference, Concordia University, HEC Montréal, Northwestern University, Queen's University, and Western University for their insightful comments and suggestions. This paper was also presented as "Intraday pricing and mispricing of short-maturity SPX options". The three authors are with HEC Montréal; Dorion and Orłowski are with the Canadian Derivatives Institute. Corresponding author: Piotr Orłowski.piotr.orlowski@hec.ca, HEC Montréal, 3000 Chemin de la Côte-Sainte-Catherine, Montréal, Québec H3T 2A7, Canada. We gratefully acknowledge financial support from the Canadian Derivatives Institute and the Social Sciences and Humanities Research Council.

1. Introduction

Zero-days-to-expiration (0DTE) options grew from a rarity to roughly half of S&P 500 option volume in under a decade: 45% of average daily volume by the end of our sample in 2023, above 60% by mid-2025.[1] Whether the pricing of these contracts reflects risk or is exploitation, is unsettled: 0DTE options appear systematically overpriced, and market lore holds that liquidity providers profit at retail speculators' expense. Using intraday data, we show that six factors span short-maturity option returns: three realized moments of the underlying return and three components capturing the dynamics of variance risk explain 99.1% (puts) and 98.2% (calls) of 30-minute return variation. The factor structure recovers the factors' pricing-measure (Q) expectations in every 30-minute period, providing a novel real-time estimator of latent conditional risk prices from which we estimate the factors' risk premiums; the result is a high-frequency view of the option market's stochastic discount factor. Unexplained variation concentrates in 0DTE contracts with pronounced intraday seasonality; this residual is linked to dealer inventories and balance-sheet conditions, yet it cannot be monetized by a factor-neutral hedged strategy once transaction costs are taken into account.

An unhedged index option return is mostly a directional bet: an exposure to the realized first moment of the underlying return, leveraging the equity risk premium. Delta-hedging removes this component, cutting return variation by 95.0% for at-the-money puts (91.7% for calls) and leaving compensation for volatility and higher-moment risk; the delta-hedged return is therefore the natural object for studying the economic determinants of option returns, even at 0DTE.[2] Our return panel spans SPX options with 15 or fewer days to expiration, and a single set of factors and risk prices must price them all, from two-week contracts to options in their final trading hours. Across maturities, we find a sharp transition in the drivers of delta-hedged return variation. For 0DTE contracts, the dominant factor is the realized second moment of the underlying return: an exposure to *prevailing* volatility. As maturity increases, *changes* in spot volatility explain most of the return variation. Jointly, the five factors beyond the underlying return account for 71% of delta-hedged return variation at

[1] The 2023 figure is computed in our sample (Section 2). For 2025, see Cboe Insights, "SPX 0DTE Options Jump to 61% Share on Retail Resurgence" (June 2, 2025) and "The State of the Options Industry: Q3 2025" (October 29, 2025), https://www.cboe.com/insights/.

[2] See, among others, Bakshi and Kapadia (2003) and Broadie, Chernov, and Johannes (2009).

0DTE and 86% at 4–15DTE.

Conditional expectations of powers of S&P 500 index returns under the pricing measure display distinctive yet familiar dynamics: the Q-second moment co-moves with the VIX, while the Q-third moment is negative and spikes more sharply in stress periods. These are the first such estimates from a characteristics-instrumented factor model at intraday frequency and ultra-short maturities; their overall agreement with estimates from static-replication methods supports the approach. Furthermore, in contrast with static-replication methods, our estimates are frequency- rather than horizon-dependent—indeed, they describe Q-expectations over 30 minutes. At such a high frequency, the squared and cubed index returns act as realized variance and skewness factors: when a large index return moves a delta-hedged option, it does so through these two channels (Section 3).

The premium on realized variance is small (19 bp per year) but reliably positive; rather than a high-frequency analog of the negative variance risk premium, we interpret it as compensation for warehousing gamma, the exposure market makers absorb from 0DTE customers (Adams, Dim, Eraker, Fontaine, Ornthanalai, and Vilkov, 2025). The realized skewness factor is a directional pure-jump one and carries the positive premium that asset pricing theory prescribes as compensation for bearing negative skewness risk (Kraus and Litzenberger, 1976; Harvey and Siddique, 2000; Bali and Murray, 2013); scaled by the large exposures of deep out-of-the-money options, it amounts to roughly ±2% per year on each leg of a DOTM risk-reversal. The central new stylized fact is a three-way decomposition of the compensation for variance risk: the jump-variance component is negatively priced, the term-structure component positively, and the two largely offset, while the spot-variance premium is small on average and turns negative around FOMC announcements. The split echoes, at 30-minute frequency, the divide between jump and diffusive variance compensation documented at longer horizons (Bollerslev and Todorov, 2011; Andersen, Fusari, and Todorov, 2015) and the distinct pricing of the variance term structure (Egloff, Leippold, and Wu, 2010).

Despite the factor model's strong fit, we uncover systematic and persistent pricing errors concentrated in 0DTE option returns. We estimate high-frequency time-series of option alphas, the components of expected returns unexplained by factor premiums, and analyze their patterns across maturity, moneyness, and time-of-day. At 0DTE, the model's pricing biases are large and negative, reaching −7.73% p.a. for OTM calls and −4.63% for ATM puts.

For longer-dated options, the magnitudes are below 3.16% p.a. The volatility of 0DTE biases is also four-to-five times higher than at one-day-to-expiration. The alphas display significant autocorrelation and strong intraday seasonality and thus despite their volatility, we consider them informative rather than noise. Relative to the six-factor model, a long position in 0DTE options underperforms in the morning, overperforms around midday, and reverts to negative alpha toward the close; the pattern disappears at longer maturities. On FOMC announcement days, alphas are uniformly negative, consistent with the model's inability to accommodate elevated option prices due to event risk with known timing. Regression evidence links the time variation in the alphas to option open interest (a proxy for aggregate inventory) and, to a lesser extent, to intermediary balance-sheet conditions. These patterns are diagnostically rich and point to specific avenues for model improvement. However, a six-factor-neutral trading strategy designed to exploit the alphas does not survive realistic transaction costs: the cumulative return vanishes when effective half-spreads are just 5% of the quoted half-spread, well below even the most optimistic estimates from the literature (Muravyev and Pearson, 2020). The pricing errors are therefore statistically detectable but economically insignificant, a reassuring indication that the six-factor model captures the bulk of compensated risk in the intraday SPX option market.

Our study contributes to four literatures. First, factor models of option returns: Jones (2006) pioneers nonlinear factor analysis of S&P 500 option returns, Christoffersen, Fournier, and Jacobs (2018) document the factor structure in equity options, and Büchner and Kelly (2022) and Horenstein, Vasquez, and Xiao (2026) estimate conditional factor models on option panels. Building on the semi-parametric, characteristics-instrumented design of Fournier, Jacobs, and Orłowski (2024), we bring this class of models to ultra-short maturities and intraday frequency, where the same exposure surfaces must price contracts hours from expiration alongside two-week options. Second, the pricing of short-horizon market risk (Bollerslev and Todorov, 2011; Andersen et al., 2015): we decompose the compensation for systematic variance innovations into spot-variance, jump-variance, and term-structure components with distinct, largely offsetting prices, estimated within the trading day. Third, the recovery of risk-neutral moments from option prices: static replication is limited to functions of the underlying return and to traded maturities, whereas the factor structure unlocks Q-expectations of any spanned factor at high frequency. The resulting series are natural state variables, akin to the VIX, and combining them with predictive models of the

factors extends our unconditional premium estimates to conditional ones. Fourth, intermediary and demand-based asset pricing (Gârleanu, Pedersen, and Poteshman, 2009; Ni, Pearson, Poteshman, and White, 2021; Chen, Joslin, and Ni, 2019; He, Kelly, and Manela, 2017): the model's pricing errors comove with option open interest and dealer balance-sheet health, and on FOMC days, when scheduled macroeconomic risk commands a premium (Savor and Wilson, 2013), they are uniformly negative, consistent with event-risk compensation that unconditional exposures miss. Throughout, we hold pricing errors to an economic rather than merely statistical standard (Novy-Marx and Velikov, 2016): the residual is economically insignificant, not exploitable mispricing. A factor-neutral strategy designed to exploit it earns a frictionless Sharpe ratio of 1.4 for just 30 minutes of daily market exposure, but its profits vanish at a small fraction of realistic transaction costs.

Our research methodology relies on the foundation of conditional factor models (Jagannathan and Wang, 1996; Nagel and Singleton, 2011; Gagliardini, Ossola, and Scaillet, 2016). When factors are not traded, their Q-expectations encode risk compensation and are unrelated to the risk-free rate. These Q-expectations can be backed out from a factor model's conditional intercepts given the assets' exposures. Option exposures are non-linear and time-varying, but tightly cross-sectionally linked, hence we model them non-parametrically as smooth functions of option characteristics; we model the conditional intercepts similarly. The Black–Scholes–Merton implied volatility is a key variable included in both exposures and intercept functions, as it summarizes market expectations of future volatility and risk. After estimating the conditional intercepts and exposures, we recover the factors' Q-expectations by estimating, at each point in time, cross-sectional regressions of the intercepts on the exposures, as implied by no-arbitrage for models with non-traded factors. To estimate systematic pricing errors, we add to the cross-sectional regressions a set of fixed effects that capture common components of expected returns of options with similar characteristics not spanned by exposures.

Related literature A nascent but rapidly growing literature investigates the implications of 0DTE or ultra-short-maturity (USM) options for market dynamics and asset pricing. Considering options with conventional (non-USM) maturities, Ni, Pearson, Poteshman, and White (2021) find that market-maker rebalancing their hedging positions significantly affects volatility dynamics of individual stocks, implying that intermediaries play a central role in

volatility transmission. Lipson, Tomio, and Zhang (2023) document that retail option trading increases the volatility of the underlying securities, suggesting that the flow of speculative trades can transmit directly into underlying asset prices.

As trading in ultra-short-maturity options grew, market participants raised concerns that same-day expiry options could pose a systemic threat to markets. Careful analysis gives these concerns little credence: Adams, Dim, Eraker, Fontaine, Ornthanalai, and Vilkov (2025) show that the introduction of daily S&P 500 expirations attenuates, rather than amplifies, market volatility, because market makers' hedging flows typically move against contemporaneous price trends. Liquidity provision, rather than speculative activity, thus dominates the impact of 0DTE trading on the underlying index. Fu, Li, Musto, and Pearson (2024) document that realized spreads in the USM market, a measure of liquidity provider profits, are low, at odds with the widespread belief that liquidity providers exploit public customers trading 0DTE SPX options. This evidence suggests that market makers can efficiently hedge the risk factors that, beyond movements in the underlying, drive the returns of USM options; our factor-neutral strategy results (Section 5.6) corroborate this reading, as what liquidity providers earn is compensation for factor risk rather than alpha at customers' expense.

We are the first to estimate a conditional factor model of option returns on USM contracts, at 30-minute frequency and down to hours before expiration; existing option-return factor models price weekly and longer-dated contracts at daily or lower frequency. Our work is nonetheless closely related to two contemporaneous papers. First, Bandi, Fusari, and Renò (2026) develop a novel pricing formula specifically targeted to capture the time-varying skewness and kurtosis implied by 0DTE options. They obtain prices that are within the bid-ask spread for 80% of the options in their sample. They also document significant improvements in delta-gamma hedging when using their model rather than the BMS model. Almeida, Freire, and Hizmeri (2025), on the other hand, study the asset pricing implications of a kernel that would hinge on 0DTE options. One of their important findings is that investors demand a high compensation to bear variance risk over the day.

Our approach, however, fundamentally differs from those studies. In particular, we build on Fournier, Jacobs, and Orłowski (2024) (henceforth FJO) and design a model that interacts intuitive, observable risk factors with exposures estimated using machine learning methods. Such intuitive factors naturally arise in the context of classical no-arbitrage option pricing models, and in particular the Black-Scholes-Merton (BMS) model, which shapes the

language and practice of derivatives trading. Using them mitigates the black-box nature of machine learning models. The decentralized option-valuation framework of Carr and Wu (2020) leverages this insight by associating risk factors to the most common BMS greeks of delta-hedged options: vega, gamma, volga, and vanna. In similar spirit, our baseline empirical model specification relies on the following six factors: the return on the underlying (DLT), its squared (GAM) and cubed (JSK) value, and changes in the market risk-neutral variance, disaggregated into three factors (labelled dSV, dJV, and dTRM respectively). Some of the labels are chosen to evoke the greeks that are associated with these shocks in the BMS model. However, the role of the BMS model ends at providing intuition. We estimate the exposures to these shocks using machine learning methods rather than differentiating option prices implied by a pricing model.

2. Data

Our dataset comprises high-frequency records of quotes of S&P 500 Index options from September 2016 to December 2023,[3] obtained from the Cboe Global Markets (Cboe) Option Intervals dataset, which aggregates option quotes at 1-minute frequency. We use quote midpoints as option prices. As a liquid and tradable proxy for the underlying we use millisecond-precision E-mini S&P 500 (ES) futures trade records from Tickdata Inc., and risk-free rates from the zero-yield curve of OptionMetrics, assumed constant within each day.

We focus on PM-settled options with at most 15 business days to expiration (DTE), setting the intraday trading window to 9:45–15:45 ET and sub-sampling at a 30-minute frequency to avoid opening and closing anomalies and mitigate microstructure noise. Positions are tracked across the overnight period for entry, exit, and total-loss determination, but all returns entering the analysis are intraday 30-minute returns.

We further filter the options to alleviate the issues of sample survivorship bias and look-ahead bias (Duarte, Jones, Khorram, Mo, and Wang, 2026), and to discipline the zero-bid and stale-quote cases that dominate measurement at ultra-short maturities. An option first enters our sample at the *start* of a 30-minute interval only when actively traded (positive volume and a positive bid) and neither in-the-money (ITM) nor extremely deep out-of-the-money

[3]During March 2020, the S&P 500 index triggered circuit breakers four times, leading to periods of illiquidity and extreme market values. Therefore, we exclude quotes during the four circuit-breaker dates.

(DOTM).[4] Subsequently, we track it without positive volume requirements until a return *moves* it into ITM or extreme DOTM territory. The position is then closed with possible re-entry. When the bid falls to zero at the *end* of an interval, we ues the half its ask price; in the last such interval before expiration, we instead record a total loss on the contract, i.e., a -100% raw return. This specifically addresses the survivorship bias concern. To maximize coverage and avoid discarding observations due to zero bids at exact timestamps, we allow a 5-minute tolerance window.[5] We provide more details in Internet Appendix A.

Our final sample contains roughly 14 million 30-minute returns on SPX puts and 6.8 million on SPX calls. Table I details its composition. Options with 4–15 DTE account for about 85% of observations, in contrast, 0DTE options, despite their dominant share of trading activity, account for only 1.4% of put and 1.9% of call observations. The limited number of 0DTE distinct contracts is mainly because the range of actively quoted strikes narrows as maturity shortens, while the fixed strike grid at \$5 further diminishes the number of strikes inside that range. We label an option as OTM or DOTM based on the distribution of normalized moneyness, $\kappa_{t,\tau}$;[6] by design, after discarding ITM options, ATM options comprise 50% of the sample, OTM options 30%, and DOTM 20% (Figure IA.III in Internet Appendix C).

2.1 Option Return Calculation

Our analysis centers on delevered excess option returns defined as:

$$r^e_{t+1} = \underbrace{\frac{O_{t+1} - O_t}{S_t}}_{\equiv r_{t+1}} - \frac{O_t}{S_t} r^f_t = r^{raw,e}_{t+1} \times \frac{O_t}{S_t}\,, \tag{1}$$

The delevered excess return is calculated as the gain from a long position in the option, in excess of the risk-free rate, divided by the price of the underlying asset at initiation, and is equal to the raw option excess return multiplied by the inverse of option-embedded leverage,

[4] We exclude the options in the 1% extreme left and right tails, due to their illiquidity and lack of meaningful informational content. The tail ratio and strike boundary are defined in Section B.

[5] Such procedures are common in the asset pricing literature, see for instance notes on calculation of bond returns in Dickerson, Mueller, and Robotti (2023).

[6] We define normalized moneyness as $\frac{\ln(K/F_{t,\tau})}{\sqrt{\tau CIV_{t,\tau}}}$, where CIV is a variance proxy (Andersen, Bondarenko, and Gonzalez-Perez, 2015) described in Internet Appendix B.

O_t/S_t. We match the ES future prices with the option quotes at the selected timestamps to ensure consistency under the 5-minute tolerance window in sample collection. This definition has several advantages: first, it prevents the occurrence of bias in aggregate statistics due to measurement error in a return's denominator (Duarte, Jones, and Wang, 2024); second, under this definition, return exposures to risk factors do not depend on the ratio of the stock price to an option's price, which simplifies modeling, and are stationary; third, FJO document that it delivers more robust and plausible estimated exposures, most likely because it down-weights extreme returns on very deep OTM options with very low prices.

We report summary statistics for both the unhedged delevered excess returns of equation (1) and their delta-hedged counterparts,

$$r^{dh}_{i,t+1} = r^{e}_{i,t+1} - \widehat{\beta}_{\text{DLT},i,t}\text{DLT}_{t+1}\,, \tag{2}$$

where DLT is the 30-minute ES futures return and the hedge ratio, $\widehat{\beta}_{\text{DLT},i,t}$, comes from the factor model of Section 3. Table I reports annualized statistics by moneyness and maturity.

Unhedged put returns are uniformly negative across all moneyness and maturity buckets (overall mean of -4.71%), whereas call returns are positive for all maturities except 0DTE and DOTM calls (overall mean of -4.09% for 0DTE). The negative returns on 0DTE calls are surprising because calls are on average positively exposed to the market return, which has a positive risk premium. The return magnitudes are the highest for ATM options and intensify as maturity shortens, making 0DTE ATM options the most extreme: mean returns of -13.25% for puts and -4.43% for calls. After delta-hedging, mean returns for both puts and calls are negative and substantially less dispersed: mean returns are -1.57% for puts and -1.45% for calls, while standard deviations decline from 6.75% to 1.63% for puts and from 5.38% to 1.89% for calls. These intraday patterns contrast with the intraday positive delta-hedged SPX returns documented by Muravyev and Ni (2020). Two design differences reconcile the results: their shortest maturity is 4DTE, whereas our sample includes shorter contracts whose hedged returns are uniformly negative (most so for 0DTE); and their earlier sample (2004–2013) predates the markedly different intraday trading dynamics of USM options in our sample period.

3. Factor structure in option returns

To motivate the specification of our factor model, and better understand the factor loadings and risk premiums, we start with a discussion of a structural option pricing model. We give it the following features: (a) two stochastic variance factors: one capturing the spot variance dynamics, and another capturing its long-run dynamics (see, for instance, Aït-Sahalia, Karaman, and Mancini, 2020), and (b) a jump component in the underlying return process with a time-varying intensity, only partially spanned by variance dynamics.[7] This three-factor specification arises from the stylized facts documented by the literature: such model features help explain option prices. In what follows, we discuss how they can help understand option return dynamics.

We emphasize that our empirical model is not a direct implementation of the structural model. The parametric restrictions imposed by SVJ-type models are well known to limit their ability to describe option returns; our non-parametric factor model is freed from these restrictions while retaining the same structural interpretability.

3.1 Three-factor structural option pricing model

Let S_t denote the level of the S&P 500 index. Under the risk-neutral measure, the dynamics of the index, the variance process, and jumps are given by:

$$\begin{aligned}
\frac{\mathrm{d}S_t}{S_{t^-}} &= (r-\delta-\nu_t)\,\mathrm{d}t + \sqrt{(1-\rho^2)\,v_t}\,\mathrm{d}W_{St} + \rho\sqrt{v_t}\,\mathrm{d}W_{vt} + \left(e^J-1\right)\,\mathrm{d}N_t \qquad (3)\\
\mathrm{d}v_t &= k_v\,(m_t - v_t)\,\mathrm{d}t + \sigma_v v_t^{\delta_v}\,\mathrm{d}W_{vt}\\
\mathrm{d}m_t &= k_m\,(\theta_m - m_t)\,\mathrm{d}t + \sigma_m m_t^{\delta_m}\,\mathrm{d}W_{mt}\\
\mathrm{d}u_t &= k_u\,(\theta_u - u_t)\,\mathrm{d}t + \sigma_u u_t^{\delta_u}\,\mathrm{d}W_{ut}\,.
\end{aligned}$$

with $J_t \sim N(\mu_J, \sigma_J^2)$, intensity $\lambda_t = \lambda_0 + \lambda_1 v_t + \lambda_2 u_t$ and jump compensator $\nu_t = \lambda_t(e^{\mu_J+\sigma_J^2/2} - 1)$. Whereas the natural next step seems to be to define a pricing kernel and study the pricing of the risk factors, this would lead to well-documented hurdles. In this model, the

[7] Christoffersen, Jacobs, and Ornthanalai (2012) show that time-varying jump intensities and the associated risk premiums are essential for jointly fitting S&P 500 returns and options. Christoffersen, Feunou, Jeon, and Ornthanalai (2021) further link the dynamics of crash risk embedded in index options to the liquidity of the underlying index constituents.

stochastic variance and jump intensity processes are not observable and estimating them requires imposing very strict parametric assumptions which likely entail mis-specification. As our goal is to obtain sensitivity and risk premium estimates with less structure, and using the standard asset-pricing framework of a conditional linear factor model, we propose option-derived factors which span the factor space of this structural model. Statements in this subsection are supported by Internet Appendix D, developing that model. For a complementary Taylor-series perspective on the full option return dynamics, see Internet Appendix E.

3.1.1 Observable proxies to span variance dynamics

Our proxies are based on the notions of model-free variance and entropy measures derived from option prices and they target the following quantities: spot variance, risk-neutral jump variance, and the slope of the term structure of risk-neutral variance. These three quantities are not collinear. In the model, they load differently on the three factors and thus span the variance and jump intensity dynamics.

The literature presents us with a direct proxy for $\mathrm{d}v$ in the form of spot variance of Todorov (2019). Therefore, our first variance factor is $\mathrm{dSV}_t := \mathrm{d}v_t$, the change in spot variance. we defer the details of the calculation of this and remaining proxies to section 3.2. To derive the remaining proxies, we combine spot variance with the notion of model-free implied variance (MFIV; Jiang and Tian, 2005; Andersen et al., 2015, the VIX index calculated by Cboe is a specific implementation of this concept), which in the context of model (3) is an affine function of the factors:[8]

$$MFIV_{t,\tau} = -\frac{2}{\tau}\mathbb{E}_t^{\mathbb{Q}}\left[\log F_{t+\tau}/F_t\right] = c_0 + c_v(\tau)v_t + c_m(\tau)m_t + c_u(\tau)u_t\,.$$

Our second proxy, dJV, is the change in the difference between short-maturity (τ_S) model-free implied variance and spot variance, $\mathrm{dJV}_t := \mathrm{d}MFIV_{t,\tau_S} - \mathrm{d}v_t$. At leading order, it captures the dynamics of the expected variance of jumps. Indeed, a Taylor expansion of

[8]To be rigorous, we note that our "variance" proxies are in fact entropy proxies (Martin, 2017; Schneider, 2015). However, we allow ourselves to label them as variance proxies since (a) most of the literature does refer to those as variance proxies, and (b) we explicitly account for higher moments of jumps in Section 3.1.2.

$MFIV_{t,\tau} - v_t$ around $\tau = 0$ yields:

$$MFIV_{t,\tau} - v_t = \sigma_J^2 \lambda_t + O(\mu_J^2) + O(\tau)\,.$$

The third proxy, dTRM, is the change in the difference between model-free implied variances at a long maturity (τ_L) and the short maturity τ_S used in dJV, or the change in the slope of the term structure of risk-neutral expected variance, $\mathrm{dTRM}_t := \mathrm{d}MFIV_{t,\tau_L} - \mathrm{d}MFIV_{t,\tau_S}$. When orthogonalized to the first two variance shock proxies, dTRM is identified by shocks to the long-term variance process m_t.

Notably, the sum of the three factors is equal to the change in the long-maturity implied variance,

$$\mathrm{dSV}_t + \mathrm{dJV}_t + \mathrm{dTRM}_t = \mathrm{d}MFIV_{t,\tau_L}\,.$$

This observation will facilitate the comparison of the estimate of the total price of variance risk in our setting with the estimates in the literature based on different samples and methodologies. In what follows, we will sometimes refer to the sum of dSV, dJV, and dTRM as dVAR.

3.1.2 The non-linear jump term

Equipped with these proxies for the three variance shocks, we can rewrite the dynamics of the delevered option return as:

$$\frac{\mathrm{d}O}{S} - r\frac{O_{t^-}}{S}\mathrm{d}t = \frac{\Theta_{t^-}}{S}\,\mathrm{d}t \;+\; O_{S,t^-}\frac{\mathrm{d}S_t}{S_{t^-}} \;+\; \frac{\widetilde{O}^v_{t^-}}{S_{t^-}}\,\mathrm{dSV} \;+\; \frac{\widetilde{O}^u_{t^-}}{S_{t^-}}\,\mathrm{dJV} \;+\; \frac{\widetilde{O}^m_{t^-}}{S_{t^-}}\,\mathrm{dTRM} + \frac{\widetilde{O}^N_{t^-}}{S_{t^-}}\,\mathrm{d}N_t,$$

where the final term accounts for the non-linear impact of jumps in the index on option returns.[9] To capture this non-linear impact in a discrete-time factor model, we use a Taylor expansion of this term around S_{t^-}, which results in

$$O\left(S_{t^-}e^J\right) - O(S_{t^-}) - O_{S,t^-}S_{t^-}\left(e^J - 1\right) \approx \frac{1}{2}O_{SS,t^-}S^2(e^J-1)^2 + \frac{1}{6}O_{SSS,t^-}S^3(e^J-1)^3$$

Otherwise stated, when the realized return is large enough that it impacts the delta-hedged option return, it does so through variance (the square) and skewness shocks (the cube).

[9] Rather than first derivatives with respect to the latent factor, $\widetilde{O}^x$, $x \in \{v, u, m\}$ are loadings identified by those first derivatives as shown in Internet Appendix D.

We propose the factors GAM and JSK, the squared and cubed returns on the S&P E-Mini futures contract, as proxies for the quadratic and cubic terms, respectively and renormalize the exposures accordingly.

3.2 Empirical factors

We operationalize the variance-spanning factors of Section 3.1.1 with model-free estimators, computed independently of our factor-model estimation. Spot variance (SV, the empirical counterpart of v_t) is estimated following Cboe's SPOTVOL methodology, based on the nonparametric estimator of Todorov (2019);[10] risk-neutral variance is measured by the corridor implied variance (CIV; Andersen et al., 2015). Following the SPOTVOL convention, the short maturity τ_S averages the two shortest available maturities exceeding 2DTE for SV and CIV alike, while the long maturity τ_L uses the longest available contracts (14–15DTE). By construction, no factor input comes from 0DTE options. Internet Appendix B details the estimators and the resulting empirical factor definitions. It is also instructive to discuss the conceptual difference between the variance dynamics factors and GAM, the squared return realization. The GAM factor captures the realized variation in the underlying return over a short interval, a high-frequency realization of the prevailing instantaneous variance at the start. In contrast, the variance dynamics factors capture changes in the variance process itself, i.e., changes in expectations about the future realizations of GAM. From an intertemporal asset-pricing perspective, they are associated with changes in the investment opportunity set, and thus are likely to command risk premiums.

Finalizing our model requires choosing instruments for the risk exposures. We use the same signals for each factor, and the non-parametric estimation delivers different exposures. In the BMS framework, the greeks are functions of stochastic time to expiration ($\tau_{i,t}$, which accounts the underlying intraday varaince profile) and adjusted moneyness ($\kappa_{i,t}$), both defined in Internet Appendix C, and of implied volatility ($IV_{i,t}$), which aggregates market-state and relevant pricing information. As demonstrated by Chen, Cheng, Liu, and Tang (2023), machine learning models benefit from being "taught" structural models. In that spirit, we instrument our model with $\boldsymbol{c}'_{i,t} = [\ln \tau_{i,t},\, \kappa_{i,t},\, \ln IV_{i,t}]$, the determinants of the BMS greeks. The logarithms improve estimation, adding flexibility for 0DTE options as maturity

[10] https://www.cboe.com/us/indices/dashboard/spotvol/

approaches zero and preventing overfitting of extreme values of implied volatility during the COVID-19 pandemic.

3.2.1 Factor dynamics

Table II reports, among others, the physical means of the factors ($\overline{f}$), in annualized percentage points. The average realization of DLT, 5.2%, is the expected excess daytime return on the nearest-maturity S&P 500 futures contract. The 1.6% mean of GAM is consistent with the annualized variance of the S&P 500 intraday index return being around 13%. The sample means of the variance dynamics factors are more difficult to interpret directly. In particular, the fact that after annualization the sum of these means is large in magnitude signals that there are large fluctuations in the levels of variance outside of our intraday sample, but within our sample period. The mean of dJV is large and negative: dividing by 252, it implies an average decrease in the short-term variance premium of about 7.3 basis points between 9:45 ET to 15:45 ET. Over the same period, the means of dSV and dTRM are positive, corresponding to average daily increases of about 2.3 and 3.0 basis points, respectively. These observations are consistent with strong seasonal patterns in intraday variance dynamics documented in the literature.

The factors are not strongly persistent (reported in Table IA.II). Only GAM exhibits moderate persistence (0.35 at 30 minutes and 0.36 at one day); dJV has a one-day autocorrelation of 23%, while the remaining factors have one-day autocorrelations at most 10% in absolute value.

The joint correlation structure of the factors (tabulated in Table IA.III), contains several notable features. First, DLT is negatively correlated with dSV (-0.56) and dJV (-0.33), and positively correlated with dTRM (0.37) and JSK (0.39). Second, GAM—the realization of prevailing instantaneous variance—is, in practical terms, uncorrelated with all factors except for JSK (0.35). This correlation structure has implications for one's expectations about the factors' risk premiums, which we confront with empirical estimates in Section 4.

3.3 Model implementation and estimation

We employ the conditional factor model from FJO to capture the non-linear nature of option exposures to risk factors and potential time-variation in risk premiums. The delevered

excess option returns are modeled as:

$$\begin{aligned} r^e_{i,t+1} = r_{i,t+1} - E^Q_t[r_{i,t+1}] &= \sum_l \left(-E^Q_t[f_{l,t+1}] \times \beta_{l,i,t} + \beta_{l,i,t} \times f_{l,t+1} \right) + \varepsilon_{i,t+1} \\ &= c_{i,t} + \sum_l \beta_{l,i,t} f_{l,t+1} + \varepsilon_{i,t+1} \end{aligned} \tag{4}$$

where $E^Q_t[r_{i,t+1}]$ is the risk-neutral expectation of the option return, $f_{l,t+1}$ represents the risk factor realization between time t and $t+1$, $c_{i,t}$ is the conditional intercept, and $\varepsilon_{i,t+1}$ denotes the model residual that satisfies $E^P_t[\varepsilon_{i,t+1}] = E^Q_t[\varepsilon_{i,t+1}] = 0$. Under no-arbitrage, the conditional intercept, $c_{i,t}$, should equal the negative of $\sum_l E^Q_t[f_{l,t+1}]\beta_{l,i,t}$, the sum of the products of negative risk neutral expectation of the factors and their corresponding exposures. We do not impose this restriction in first-stage estimation, but we exploit it later to estimate conditional risk-neutral moments of the factors.

Following FJO, we model option exposures as functions of option characteristics using the thin plate regression spline (TPRS) basis; instrumenting exposures with characteristics is the standard device in conditional factor models for option returns (see also Büchner and Kelly, 2022). The risk loading of option i can be written as

$$\beta_{l,i,t} \approx \boldsymbol{\phi}'_{i,t} \boldsymbol{b}_l \tag{5}$$

where $\boldsymbol{b}_l$ is a $k \times 1$ vector of factor-specific parameters capturing the basis representation of $\beta_{l,i,t}$ across i and t, $\boldsymbol{\phi}_{i,t}$ is $\phi(\boldsymbol{c}_{i,t})$, a "feature" vector in a higher-dimensional space, and $\boldsymbol{c}_{i,t}$ is the vector of option characteristics. As a result, the exposure to a given risk factor l is linear in $\boldsymbol{b}_l$, the parameter vector. Combining equations (1) and (5), and setting $E^Q[f_{l,t+1}] = \mu_{l,t}$ implies the following conditional factor structure for the delevered option return:

$$r^e_{i,t+1} = \sum_l \left(-\mu^Q_{l,t} \times \boldsymbol{\phi}'_{i,t} \boldsymbol{b}_l + f_{l,t+1} \times \boldsymbol{\phi}'_{i,t} \boldsymbol{b}_l \right) + \varepsilon_{i,t+1} \,. \tag{6}$$

To express equation (6) as linear in parameters, we first define $\boldsymbol{F}_{l,i,t+1} = f_{l,t+1} \times \boldsymbol{\phi}_{i,t}$, a $k \times 1$ vector of scaled factor realizations. The conditional intercept $c_{i,t}$ is the sum of products of interactions between $\boldsymbol{\phi}'_{i,t}$, and the conditional risk-neutral expectations. Because our option characteristics contain the option implied volatility, the conditional intercept contains

information relevant for risk-neutral moment and risk premium estimation, and we set $\mu_{l,t}$ to constants. With $\boldsymbol{\phi}_{i,t}$ identical for each factor, we rewrite the conditional intercept as $c_{i,t} = \boldsymbol{\phi}'_{i,t}\boldsymbol{q}$, where $\boldsymbol{q} = \sum_l -\mu_l \times \boldsymbol{b}_l$. With these definitions, we rewrite equation (6) as:

$$r^e_{i,t+1} = \underbrace{\boldsymbol{\phi}'_{i,t} \sum_l \boldsymbol{q}_l}_{c_{i,t}} + \sum_l \boldsymbol{F}'_{l,i,t+1}\boldsymbol{b}_l + \varepsilon_{i,t+1}\,. \tag{7}$$

In this form, we obtain a scaled factor model with a time-varying intercept. The model is linear in parameters $\boldsymbol{q}$ and $\boldsymbol{b}_l$ but the risk-neutral means are not uniquely identified. They can be estimated separately: we obtain period-by-period estimates of $\mu^Q_{l,t}$ from Fama-MacBeth regressions, whereas FJO impose more structure on risk premium dynamics and run an extra set of regressions.[11] OLS estimation of equation (7) yields coefficients $\widehat{\boldsymbol{q}}$ and $\widehat{\boldsymbol{b}}_l$ which can be used to recover the conditional intercept and option exposures. These two quantities can be coupled with factor realizations to examine the explanatory power of the model.

3.4 Empirical results

To assess various aspects of the model's explanatory power, we compute the R^2 in percentage terms as:

$$R^2\% = \left[1 - \frac{\sum_{i,t}\left(r_{i,t+1} - \widehat{r}_{i,t+1}\right)^2}{\sum_{i,t}(r_{i,t+1})^2}\right] \times 100, \tag{8}$$

pooled over the whole option panel, with fitted values of the general form $\widehat{r}_{i,t+1} = \widehat{c}_{i,t} + \sum_l \widehat{\beta}_{l,i,t} f_{l,t+1}$. The subsections below vary which terms enter the fitted value; the coefficients always come from the single estimation with all six factors. Two remarks clarify what this statistic measures. First, the coefficients $\widehat{\boldsymbol{q}}$ and $\widehat{\boldsymbol{b}}_l$ are estimated once, by pooled OLS on the full panel, and the exposures entering $\widehat{r}_{i,t+1}$ are functions of time-t characteristics only; the R^2 thus measures the contemporaneous spanning of option returns by the factor realizations, the analogue of a market-model R^2, not an out-of-sample forecast. Second, the fit is disciplined

[11] If the set of characteristics for instrumenting exposures differed across factors, and if we had additional state variables for factor risk premiums, in the conditional intercept we could replace $\sum_l -\mu_{l,t} \times \boldsymbol{\phi}'_{i,l,t}\boldsymbol{b}_l$ with $\widetilde{\boldsymbol{\phi}}'_{i,t}\boldsymbol{q}$, where $\widetilde{\boldsymbol{\phi}}_{i,t}$ is a new TPRS basis defined on the complete set of characteristics and state variables.

by parsimony: each factor enters through a single loading surface parameterized by 84 basis coefficients, so the number of free parameters is negligible relative to the millions of option-interval returns being fit, and the basis cannot mechanically absorb return variation.

3.4.1 Option exposures

Figure 3 reports the representative exposures to the factors across selected maturities. Exposures are computed for synthetic options on a fixed moneyness-maturity grid.

The DLT exposure reproduces the textbook delta: positive for calls, negative for puts, with magnitudes empirically bounded by ± 1, and monotone in moneyness at all maturities. Put-call parity implies that higher-order exposures coincide for puts and calls; we impose this restriction in estimation. The GAM exposure traces the familiar gamma: positive throughout, peaking near at-the-money, and decreasing roughly symmetrically into both wings. The JSK exposure largely tracks the BSM shape of the third S-derivative of the option price: negative on the OTM-put side, switching sign near ATM, and largest in magnitude at moderate moneyness on either side. The shapes of the three exposures broadly match their definitions as successive derivatives of the option price with respect to the underlying.[12]

By the same put-call parity argument, the variance-dynamics exposures dSV, dJV, and dTRM coincide for puts and calls. The dSV and dJV exposures are positive across the surface: each factor represents an increase in a component of risk-neutral variance, and option prices respond monotonically to such increases. The dTRM exposure carries no a priori sign—it captures the change in the slope of the variance term structure, which can move either way depending on how short- and long-maturity expectations adjust. Empirically it is positive: after controlling for spot- and jump-variance moves, options remain sensitive to long-run variance shocks. All three profiles are bell-shaped in moneyness, like BSM vega: a single mode with smooth decay into both wings.

The pattern of the estimated variance-dynamics exposures validates the three-factor specification. Under a one-factor variance specification option prices would respond only to

[12] A naive reading of the figure would interpret each curve as the slope of the previous one—GAM $= \partial\text{DLT}/\partial\kappa$ and JSK $= \partial\text{GAM}/\partial\kappa$. Both relations are violated in sign: DLT decreases in κ where GAM is positive, and GAM rises on the OTM-put side where JSK is negative. The puzzle resolves once we recall that $\kappa \propto \log(K/F)$ moves opposite to S at fixed K. In S-space the BSM relations $\Gamma = \partial\Delta/\partial S$ and the third-order counterpart of JSK, $\partial^3 O/\partial S^3 = \partial\Gamma/\partial S$, hold by construction; the chain rule flips the sign each time $\partial/\partial S$ is converted into $\partial/\partial\kappa$.

the level of implied variance, and not to its slope. Such a model could also include jumps as long as their sensitivity was spanned by the variance dynamics. This assumption would force the dSV and dJV exposures to be equal, since their sum gives the change in short-maturity implied variance, and the dTRM exposure to be zero. The data reject both restrictions: the dSV and dJV profiles differ visibly across moneyness in both magnitude and curvature, and the dTRM exposure is consistently positive across the surface rather than vanishing.

The maturity dimension behaves as option pricing theory predicts. Return-power exposures (GAM, JSK) have the largest magnitude at 0DTE and shrink rapidly toward zero with maturity. Variance-dynamics exposures (dSV, dJV, dTRM), on the contrary, grow with maturity and peak on the OTM-put side—mirroring the R^2 decomposition of Section 3 and locating variance risk in the left tail of longer-dated contracts. Interestingly, sensitivity to spot variance (dSV) increases just as much as sensitivity to longer-term variance (dTRM) as option maturities increase—a pattern broadly consistent with the volatility overreaction documented by Vasquez (2017) in the cross-section of equity options.

3.4.2 Delta-hedging efficiency

We begin by examining the model's ability to delta-hedge option returns: the fitted value retains only the delta term, $\widehat{\beta}_{\text{DLT},i,t}\text{DLT}_{t+1}$, without the conditional intercept.

Delta-hedging is highly effective in both puts (94.2%) and calls (87.6%). It remains effective even at the shortest maturities, including 0DTE options, and in fact rises only slightly from 0DTE to 1–2DTE only to drop again for longer maturities (Figure IA.IV).

Table IA.IV reports (in the columns labeled "DLT") results further disaggregated by moneyness. Naturally, delta-hedging is more effective for at-the-money (ATM) options and drops off further out-of-the-money (OTM). In 0DTE options, the drop-off is more pronounced in puts than in calls, and an opposite pattern holds in calls. At longer maturities, the delta-hedging effectiveness falls off faster for calls, with 4–15DTE OTM calls being effectively unhedgeable using only the underlying.

3.4.3 Model explanatory power for delevered excess option returns

The residual variation after delta-hedging can be further explained by the remaining five factors, and hence we report the model's total explanatory ability for option returns: the

fitted value now includes the conditional intercept and all six factors.

In puts, the six factors and conditional intercept explain 99.1% of total return variation; in calls, 98.2%. There is very little difference in explanatory power across maturities, with the R^2 for 0–1DTE options being between one and two percentage points lower than for longer maturities (Figure IA.IV). The results are similarly strong across moneyness levels (cf. Table IA.IV, in the columns labeled "6F"). [13] In OTM options, the R^2 is 93.9% for calls and 97.1% for puts, and the sub-sample with the worst fit—0DTE OTM puts—still has an R^2 of 87.1%. We see a drop-off in explanatory power for DOTM options, from 30.0% for 0DTE puts to 75.8% for 4–15DTE puts, and 74.1% across the whole DOTM put sub-sample; the drop-off in calls is slower than in puts, except for the longest maturities.

3.4.4 Sources of variation in delevered delta-hedged option returns

We examine in detail the sources of variation in delta-hedged option returns: the returns are now the delta-hedged delevered excess returns of equation (2), and the fitted value includes the conditional intercept and all factors except DLT.

Table III reports, for put and call options, by moneyness and maturity, the R^2 of the five-factor model in explaining delta-hedged delevered excess returns (in columns labeled "5F"). Furthermore, the table reports the incremental R^2 of each factor when added to a model that already includes DLT and other factors that are listed before it in the table.

Our first observation is that the five-factor model explains a significant portion of the variation in delta-hedged option returns, the R^2 is 84.3% in puts and 85.0% in calls. It is noteworthy that the explanatory power remains very high even at the shortest maturity and in OTM options (approximately 70% in both puts and calls), only dropping away in DOTM 0DTE contracts (41.9% (50.4%)). For 1+DTE options, the R^2 never drops below 77.9%. Even though delta-hedging removes up to 94% of return variation, the remaining five factors still explain a large portion of the residual variation. Their explanatory power

[13] Table IA.V repeats the analysis for the four-factor model in which the aggregate variance factor dVAR replaces dSV, dJV, and dTRM. The aggregate model attains a nearly identical overall fit (98.8% in puts, 97.7% in calls), but the decomposition matters precisely where explaining returns is relatively hard. It raises the R^2 of 1–3DTE OTM puts (calls) from 91.8% (92.1%) to 95.8% (94.6%), of 1–3DTE DOTM puts (calls) from 57.6% (75.2%) to 69.1% (81.6%), and of 4–15DTE DOTM puts (calls) from 70.6% (66.9%) to 75.8%(76.1%). These gains are large given that DLT is already effective at explaining the return variation, which leaves small residuals for the remaining factors.

for delta-hedged returns compares very favorably to, for instance, factor models for stock returns, which typically explain 30–40% of return variation.

A more significant empirical observation is the dramatic change in the sources of variation in delta-hedged option returns as we move from 0DTE to longer maturities. In 0DTE options, the dominant factor is GAM, the squared return on the underlying, which explains 58.8% (59.0%) of return variation in puts (calls). Already at 1–3DTE, the explanatory power of GAM drops by a factor of three: to 17.5% (17.4%) in puts (calls); the drop continues—albeit more slowly—at longer maturities. This finding is complemented by the opposite pattern in the explanatory power of dSV, the spot-variance change factor, which explains only 3.7% (2.3%) of return variation in 0DTE options, but 28.6% (31.5%) already at 1–3DTE. A similar pattern holds for the other two variance dynamics factors, although their contributions are smaller in magnitude.

These results match economic intuition. At ultra-short maturities, the realized variance over the interval is the main driver of option returns after delta-hedging: the options are primarily exposed to prevailing variance and expire before the onset of future variance. Exposure of 0DTE contracts to the variance-dynamics factors is nonetheless not zero, nor should it be, because revisions in the variance expected over the remaining life reprice the contract until the very end. The estimated exposures are correspondingly small (Section 3.4.1), and the small share of 0DTE return variation attributed to dSV confirms that the maturity-smooth spline basis does not transfer longer-maturity loadings onto 0DTE contracts. As maturity increases, the role of prevailing variance diminishes, and changes in variance expectations become the dominant drivers of option returns.

When moving across the moneyness spectrum, we observe additional patterns. First, the further away from ATM, the less important GAM becomes, even at 0DTE (albeit it remains the dominant factor at the shortest maturity). There are significant changes in the relative importance of the variance dynamics factors. The explanatory power of dSV slightly decreases as we move OTM in puts, but increases in calls. Importantly, the explanatory power of dJV increases rapidly as we move OTM in puts (up to 34.2% (8.4%) in DOTM 4–15 DTE puts (calls)). These patterns are consistent with the idea that dJV captures compensation for short-term left-tail events. The explanatory power of dTRM is small across the board, increases with maturity (from −1.4% to (7.7%)), and remains stable across moneyness.

The role of the JSK factor is negligible except for DOTM 0DTE options, where it explains

13.2% (20.1%) of return variation in puts (calls); once again this is consistent with interpreting JSK as a jump-skewness factor.

4. Compensation for risk

A risk premium is the difference between two components: the physical and risk-neutral expectations of the factor. We estimate risk premiums under minimal assumptions on their dynamics while allowing for time variation and accounting for the co-movement of premiums and exposures. We separate the estimation of the conditional Q-expectations from the time-series estimation of physical means, and we test the former against restrictions implied by the model before turning to the premiums themselves.

4.1 Risk Premiums with non-traded factors

Risk premiums on non-traded factors are typically estimated by cross-sectional regressions of excess asset returns on their factor loadings. Similarly to Gagliardini, Ossola, and Scaillet (2016), we use a cross-sectional approach only to estimate the conditional Q-expectations of the factors, and then combine them with time-series estimates of the physical means to obtain risk premiums.

4.1.1 The conditional intercept and risk-neutral expectations

Let r denote the excess return on an asset, F_T an excess return factor and F_{NT} a non-traded factor, and β_T and β_{NT} the asset's exposures to the respective factors. Then, the following factor model for the asset holds in each period (we omit asset and time subscripts to simplify notation):

$$r = c + \beta_T F_T + \beta_{NT} F_{NT} + \varepsilon \,. \tag{9}$$

Given that both r and F_T are excess returns on traded assets, $E^Q[r] = E^Q[F_T] = 0$. Taking Q-expectations on both sides of equation (9) thus implies that

$$c = \beta_{NT} \times (-E^Q[F_{NT}]). \tag{10}$$

We leverage the multivariate generalization of the no-arbitrage restriction in equation (10) to estimate the Q-expectations of the factors from the conditional intercepts of the factor model for option returns estimated in Section 3.

4.1.2 Implementation

Using the notation of our factor model (4), the no-arbitrage restriction (10) becomes:

$$c_{i,t} = \sum_l \beta_{l,i,t} \times (-E_t^Q[f_{l,t+1}]) = -\sum_l \beta_{l,i,t} \times \mu_{l,t}^Q \,. \tag{11}$$

We estimate $\mu_{l,t}^Q$ period-by-period using cross-sectional regressions of the conditional intercepts estimated in Section 3, $\widehat{c}_{i,t}$, on the estimated factor loadings, $\widehat{\beta}_{l,i,t}$. Furthermore, to allow for similar pricing errors among options with similar characteristics, we group options into buckets based on their type (puts or calls), maturity (0, 1–3 or 4–15 days) and moneyness (ATM, OTM, DOTM), and include bucket fixed effects in the cross-sectional regression:

$$\widehat{c}_{i,t} = \alpha_{b(i),t} - \sum_l \widehat{\beta}_{l,i,t} \times \mu_{l,t}^Q + \nu_{i,t} \,, \tag{12}$$

where $b(i)$ is the bucket to which option i belongs.

Because both the regressand and the regressors in equation (12) derive from the first-stage estimates of Section 3, we make the source of identification explicit. The first stage estimates the conditional-intercept surface without imposing the no-arbitrage restriction (11); nothing in equation (7) ties the intercepts to the exposures. The cross-sectional regression (12) therefore exploits an overidentifying cross-equation restriction, and the bucket fixed effects measure its violations. The slopes $\mu_{l,t}^Q$ are identified by within-bucket variation: each 30-minute cross-section contains thousands of contracts, and the exposures vary continuously with maturity, moneyness, and implied volatility inside each of the nine coarse buckets per option type, keeping the loadings far from collinear with the fixed effects. Moreover, the factors are built without 0DTE options, and 0DTE contracts account for 1.4% of put and 1.9% of call observations, so the estimates of $\mu_{l,t}^Q$ are identified, in effect, from longer-maturity contracts.

We include in the regression the option deltas ($\widehat{\beta}_{\text{DLT},i,t}$), even if the DLT factor is traded.

Theory suggests that the estimated $\mu^Q_{\text{DLT},t}$ should be zero, but in practice it could deviate from zero due to frictions in the option market or misspecification of our model. Section 4.2 confronts the estimates of $\mu^Q_{l,t}$ with three restrictions they should satisfy under the model: a zero Q-expectation for the traded DLT factor, positivity of the Q-second moment $\mu^Q_{\text{GAM},t}$, and agreement of the predictable variance seasonality under P and Q. We further examine the model's pricing errors in Section 5.

4.2 Conditional risk-neutral expectations of the factors

We report on three aspects the conditional Q-moments: (a) their magnitudes, (b) time-series behavior, (c) and intraday seasonality. The six factors in our model can be categorized in two ways: first, traded (DLT) and non-traded, and second, representing the underlying asset dynamics (DLT, GAM, and JSK), and aggregate option price/volatility dynamics (dSV, dJV, dTRM).

Table II contains the averages of the conditional risk-neutral expectations of the factors, $\mu^Q_{l,t}$, alongside the sample averages of factor realizations (i.e., unconditional P-expectations),[14] and their differences, which are estimates of unconditional risk premiums λ_l, discussed in Section 4.3. Figure 1 presents daily time-series of the conditional Q-moments sampled at 09:45 ET (at the beginning of the first trading interval) and at 15:15 ET (at the beginning of the last trading interval of the day). Figure 2 presents the associated intraday seasonal patterns.

4.2.1 Underlying asset dynamics factors

Because we estimate our model using excess returns and DLT is traded, $\mu^Q_{\text{DLT},t}$ should be zero in each period. In a frictionless market, any deviation would reflect only sampling noise; in the presence of frictions, deviations can instead be informative about them. The average of the $\mu^Q_{\text{DLT},t}$ estimates is 0.01% per year. This overidentification check corroborates our estimation approach: the model does not assign, on average, a non-zero mean to a factor with a null Q-expectation, so the restriction is satisfied by the data rather than imposed on them. Figure 1 nonetheless shows that estimated $\mu^Q_{\text{DLT},t}$ exhibits considerable time-series

[14] Table IA.II contains the estimates of 30-minute (1-period) and 1-day autocorrelations for the conditional risk-neutral expectations and unconditional realizations of the factors.

variation. A deviation of $\mu^Q_{\mathrm{DLT},t}$ from zero measures the gap between the financing rate implicit in the cross-section of option prices and the risk-free rate we use to construct excess returns. The estimates deviate symmetrically from zero both in calm and market-stress periods; the magnitude of these deviations spikes in both directions during the COVID crisis. Notably, they are negative from the start of the 2022 bear market until the end of our sample in December 2023. Our estimates of $\mu^Q_{\mathrm{DLT},t}$ are very persistent: they have a positive autocorrelation of -0.02 at the 30-minute horizon and 0.03 at the one-day horizon. The persistence at the one-day horizon, in particular, indicates that these deviations are not pure estimation noise, but rather reflect genuine variation in the implied financing rates. The same conclusion follows from a clear intraday seasonality pattern (presented in Figure 2): $\mu^Q_{\mathrm{DLT},t}$ rises steadily over the trading day, from -0.4% at the open to 0.2% near the close, crossing zero in the late morning.

GAM is the square of DLT, and thus $\mu^Q_{\mathrm{GAM},t} = E^Q_t[\mathrm{DLT}^2_{t+1}]$ is the conditional Q-second moment of the high-frequency underlying excess return; and as $E^Q_t[\mathrm{DLT}_{t+1}] = 0$, $\mu^Q_{\mathrm{GAM},t}$ is the conditional variance. As a consequence, $\mu^Q_{\mathrm{GAM},t}$ should always be positive. The estimates of $\mu^Q_{\mathrm{GAM},t}$ are indeed positive not only on average in our whole sample, but also on average in every intraday interval, and practically everywhere in the sample, as evident in Figure 1. Economically, $\mu^Q_{\mathrm{GAM},t}$ is closely related to measures of option-implied variance, such as the VIX- or *CIV*-type indexes, with two distinctions. First, $\mu^Q_{\mathrm{GAM},t}$ measures the Q-second moment rather than entropy; second, those indexes are maturity-specific, tied to a given option expiration, whereas $\mu^Q_{\mathrm{GAM},t}$ is frequency-specific, defined over our model's 30-minute sampling interval. Recovered from the cross-section of option returns rather than a fixed-maturity price integral, this estimator is, to our knowledge, new to the literature. The average of the $\mu^Q_{\mathrm{GAM},t}$ estimates is 1.4% per year (Table II); given a near-zero first moment, this average corresponds to an annualized risk-neutral volatility of about 12%. The series is low and stable in calm markets but rises sharply in periods of distress, most strongly in March 2020 and again through 2022 (Figure 1). Consistent with its reading as an implied variance, it co-moves closely with the short-maturity *CIV*, with a correlation of 0.84 at the open and 0.79 near the close (Figure 1). The estimates are highly persistent, with an autocorrelation of 0.99 at the 30-minute and 0.76 at the one-day horizon, far above the persistence of the factor's own realizations, as befits a smooth measure of expected variance. Intraday, $\mu^Q_{\mathrm{GAM},t}$ declines steadily over the trading day, from 1.7% at the open to 1.1% near the close (Figure 2).

Unlike annualized implied variances computed from 0DTE options, which rapidly increase at ultra-short maturities, our fixed-horizon estimator remains robust throughout the trading day.[15]

JSK is the cube of DLT, and therefore $\mu^Q_{\text{JSK},t} = E^Q_t[\text{DLT}^3_{t+1}]$ is the conditional Q-third moment of the high-frequency underlying return. At our 30-minute frequency this power is dominated by jumps, so $\mu^Q_{\text{JSK},t}$ measures the risk-neutral skewness of jumps in the S&P 500 index.[16] Its average, $-0.2\,bp$ per year (Table II), is negative, consistent with the well-documented negative risk-neutral skewness of index returns.[17] The series is noisy and mostly negative, with pronounced downward spikes in periods of market distress; close to maturity, however, it is no longer consistently negative (Figure 1) due to an upward level shift. It is persistent within the day (autocorrelation 0.97 at 30 minutes) but far less so across days (0.20 at the one-day horizon). Intraday, $\mu^Q_{\text{JSK},t}$ rises steadily from $-0.5\,bp$ at the open toward zero by the close (Figure 2). Assuming one jump in a thirty-minute period, these magnitudes of $\mu^Q_{\text{JSK},t}$ are consistent with expected jumps returns between -0.54% and -0.25%.[18]

Overall, we recover the risk-neutral variance and skewness of index returns—and the financing rate built into option prices—at high frequency and at a fixed horizon rather than maturity.

FOMC days. Table II also reports the averages of $\mu^Q_{l,t}$ computed separately for FOMC announcement days. The Q-second moment averages $+2.0\%$ on announcement days against $+1.3\%$ otherwise: option prices embed about 50% more variance over intervals containing event risk with known timing. The asymmetry is stronger in the third moment: $\mu^Q_{\text{JSK},t}$ averages -0.6 bp on FOMC days, nearly three times its non-FOMC average of -0.2 bp, so the additional variance priced ahead of announcements is tilted towards negative jumps. The financing-rate wedge $\mu^Q_{\text{DLT},t}$ drops to -0.5% from 0.0% on the remaining days, although neither average is statistically different from zero.

[15] Vilkov (2023) reports average variance swap rates at sub-day maturities without annualizing; because they decline more slowly than the time to expiration, annualizing makes them rise through the day and explode in the final minutes before expiration.

[16] See, for instance, Jacod and Protter (2012): in powers of price increments of a semimartingale sampled at short intervals, jumps dominate.

[17] See, for instance, Dennis and Mayhew (2002), Bakshi, Kapadia, and Madan (2003), Kozhan, Neuberger, and Schneider (2013), and Orłowski, Schneider, and Trojani (2024).

[18] For instance, $-(0.2/10000/(252\times 12))^{\frac{1}{3}} \approx -0.0025$.

4.2.2 Variance dynamics factors

The remaining three factors are changes rather than levels. Their conditional Q-expectations are therefore risk-neutral drifts of variance components, and can take either sign. Unlike the moments of the previous section, they have no direct empirical benchmark: intraday, fixed-horizon risk-neutral drifts of decomposed variance components have, to our knowledge, not been measured before. We therefore evaluate the estimates on internal consistency and on their intraday behavior.

On average, $\mu^Q_{\mathrm{dSV},t}$ is positive (4.1%), while $\mu^Q_{\mathrm{dJV},t}$ (−3.5%) and $\mu^Q_{\mathrm{dTRM},t}$ (−5.4%) are negative (Table II). These are annualized expectations of 30-minute changes, so the annualized percentages are large by construction. We therefore also report each as the implied change in spot volatility, in volatility points, from a 20% reference level.[19] In these units the averages are small: +0.04, −0.03, and −0.05 volatility points per day. Their deviations from zero arise from a few crisis episodes (Figure 1).

The estimates are highly persistent within the day (30-minute autocorrelations of 0.99, 0.96, and 0.97) but transient across days (one-day autocorrelations of 0.19, 0.29, and 0.31; Table IA.II). In this they resemble jump skewness and differ from the variance level $\mu^Q_{\mathrm{GAM},t}$, which is persistent at both horizons: the drifts are governed by a fast, intraday state, not by the slow-moving level of variance.

Figure 2 shows the intraday seasonality of $\mu^Q_{\mathrm{dSV},t}$. It rises monotonically from −11.2% (−0.11 points) at the open to +16.2% (+0.16 points) at the close, crossing zero near midday. This is the time-derivative of the familiar intraday U-shape of physical variance (Andersen and Bollerslev, 1997). Our risk-neutral estimator thus reproduces a feature of physical variance. This serves as another validation of our procedure, because the predictable, seasonal volatility drift does not carry a risk premium and hence must agree under P and Q.

Because the three factors sum to the change in long-maturity implied variance (Section 3), so do their Q-expectations, and their intraday averages differ between the open and the close. At the open the whole surface is expected to fall: the short end by −14.8% (−0.15 points) and the long end by nearly as much, −11.9% (−0.12). The morning's expected decline in variance thus propagates across maturities. At the close the short end is expected to rise (+15.1%, +0.15 points), but the long end barely moves (+2.1%, +0.02); the afternoon increase is almost

[19] A variance change $d(\sigma^2)$ implies a volatility change $d\sigma = d(\sigma^2)/(2\sigma)$; at $\sigma = 20\%$, one annualized percentage point of variance is about 0.01 volatility points per day.

entirely absorbed by a flattening slope ($\mu^Q_{\mathrm{dTRM},t}$ falling to -13.0%, -0.13 points) and does not reach long horizons. Over the trading day, the risk-neutral expected repricing of the surface thus rotates from level to slope, from $\mu^Q_{\mathrm{dSV},t}$ at the open to $\mu^Q_{\mathrm{dTRM},t}$ at the close.

Periods of market distress show the same offsetting (Figure 1). In these episodes, $\mu^Q_{\mathrm{dSV},t}$ and $\mu^Q_{\mathrm{dJV},t}$ turn sharply negative: variance is elevated and is expected to fall. $\mu^Q_{\mathrm{dTRM},t}$ turns sharply positive: the term structure is inverted and is expected to re-steepen. Because the components move in opposite directions, the long end of the surface moves far less than each of them.

On average, short-maturity implied variance is nearly driftless under Q: $\mu^Q_{\mathrm{dSV},t} + \mu^Q_{\mathrm{dJV},t}$ averages $+0.6\%$. This masks a steady rotation in its composition: the jump-variance wedge is expected to decay at every time of day, while spot variance is expected to rise and replace it. The entire net expected decline in the surface, -4.7% (-0.05 vol points) per year, is thus a flattening of the term-structure slope rather than a fall in the short end; this distinction is not visible in maturity-integrated measures such as the VIX.

FOMC days. The risk-neutral drifts change character on FOMC announcement days (Table II). The average of $\mu^Q_{\mathrm{dSV},t}$ reverses sign, from $+4.7\%$ on ordinary days to -12.0% (-0.12 vol points per day): option prices anticipate the resolution of announcement uncertainty and the associated decline in spot variance. The slope drift, by contrast, vanishes ($\mu^Q_{\mathrm{dTRM},t}$ averages $+0.5\%$, against -5.6% otherwise), while $\mu^Q_{\mathrm{dJV},t}$ is somewhat more negative than on ordinary days (-5.4% against -3.4%), albeit imprecisely estimated. Summing the components, the surface is expected to decline by -16.9% (-0.17 vol points)—four times the non-FOMC figure—and, in contrast with the average trading day, the decline is led by the short end and shifts the surface down almost in parallel rather than loading on the slope. The realized decline in spot variance on announcement days, -54.7%, is larger still than the expected risk-neutral one; we return to this gap when discussing risk premiums in Section 4.3.2. A fine contrast with Section 4.2.1 emerges: on announcement days the risk-neutral expectation of imminent return variation and negative skewness rises, while expectations of variation further out in the future decline. Option prices thus place the announcement risk squarely within the trading day and treat it as a transitory burst rather than the onset of persistently higher volatility.

4.3 Risk premiums

Table II reports the annualized unconditional risk premiums in percent—differences between the average factor realizations (unconditional P-expectations) and the corresponding averages of the conditional Q-expectations—together with both averages.

4.3.1 Compensation for underlying moments

Because the sample average of $\mu^Q_{\mathrm{DLT},t}$ is essentially zero (see Section 4.2.1), the DLT risk premium collapses to the sample average of DLT, 5.2% per year.

The estimate of the premium for GAM is small (19 bp per year) but statistically significant at the 5% level. Importantly, our estimate, obtained from intraday data, is positive.[20] This sign is consistent with the structure of the 0DTE market documented in Adams, Dim, Eraker, Fontaine, Ornthanalai, and Vilkov (2025) and Vasquez, Amaya, Pearson, and Garcia-Ares (2025). They show that S&P 500 option market makers carry persistently positive net gamma arising from their absorption of end-user customers' net-short ATM-straddle positions traded predominantly in the 0DTE market. Because gamma exposure is concentrated in market makers' inventories and entails costly, time-sensitive delta rebalancing, the holder of this exposure is likely to earn a positive expected return on it to willingly carry it in equilibrium. A positive $\widehat{\lambda}_{\mathrm{GAM}}$ is consistent with this compensation. Note that although GAM locally captures realized return variance, the premium we estimate is not a high-frequency analog of the standard variance risk premium which reflects insurance demand, and is consequently negatively estimated from longer holding periods; instead $\widehat{\lambda}_{\mathrm{GAM}}$ reflects compensation for intraday liquidity provision through gamma warehousing.

Turning to jump skewness, the negative estimates of $\mu^Q_{\mathrm{JSK},t}$ combined with near-zero average realizations of JSK deliver a positive jump-skewness premium of 0.3 bp per year—the sign that crash-insurance reasoning prescribes. As the cube of a 30-minute return, JSK is mechanically tiny per period, and $\widehat{\lambda}_{\mathrm{JSK}}$ is correspondingly small in raw factor units. Option exposures to JSK are correspondingly massive (Figure 3) and JSK significantly contributes to option returns, especially in DOTM options. Option markets price negative jumps more heavily than they materialize, so investors who absorb jump-skewness exposure earn a positive

[20] While Fournier, Jacobs, and Orłowski (2024) obtain a positive estimate of similar magnitude from daily data, theirs is not statistically significant.

expected return for doing so. The result is consistent with the significant skewness premium documented at a monthly horizon in Kozhan, Neuberger, and Schneider (2013).[21] Orłowski, Schneider, and Trojani (2024) isolate jump skewness via a high-frequency third moment trading strategy and report an annualized 10.8% excess return on position margin capital. We compare this to a DOTM risk-reversal in our setting, where each leg earns roughly ±2% per year of jump-skewness compensation (Figure 4). The spread yields about 4%. This is broadly comparable to Orłowski, Schneider, and Trojani's estimate, but lower because of our return definition, reporting on a fully collateralized basis.

4.3.2 Compensation for variance risk

To assess the compensation for variance dynamics, we first sum the premiums on dSV, dJV, and dTRM to obtain the premium for the change in long-term risk-neutral model-free implied variance (dVAR)—the quantity most directly comparable to existing estimates. The result, −0.4% per year, is essentially zero; as for all our results, this is based on *intraday* option returns. The most closely related estimates in the literature are from FJO, who find a premium on the change in the Cboe VIX^2 index between −6.2% and −15.2% per year in *daily* option returns, and from Dickerson, Fournier, Jacobs, and Orłowski (2025), who find a more negative premium (approximately −20% per year) using weekly returns on index, stock, and ETF puts and calls. Estimates of similar magnitude are implied by option pricing models (see, e.g., Eraker, 2004) and are also obtained via Fama-MacBeth regressions in Duarte, Jones, and Wang (2024). Our near-zero estimate is an order of magnitude smaller than those benchmarks and is consistent with direct evidence that variance compensation in option markets accrues outside regular trading hours. Jones and Shemesh (2018) document that delta-hedged option returns are systematically lower over non-trading periods, particularly weekends. Muravyev and Ni (2020) similarly find that the SPX delta-hedged option return premium is concentrated overnight, with intraday returns slightly positive. Orłowski, Schneider, and Trojani (2024) find that intraday compensation for variance risk is an order of magnitude smaller than overnight and statistically insignificant.

Our approach delivers granular evidence on the pricing of variance risk components under minimal parametric assumptions. We show that the null finding for variance risk pricing in

[21] The key difference between JSK and Kozhan et al.'s skewness measure is that theirs is not pure jump skewness: it conflates the jump component with diffusive leverage.

regular trading hours arises from nearly offsetting compensations across the three components. To our knowledge, this is the first such decomposition of variance risk premiums.[22]

The literature suggests signs for the three premiums. The jump-variance premium should be negative: because dJV is the change in risk-neutral expected jump variance, a negative λ_{dJV} means expected increases in Q-jump variance exceed those that realize on average. This is the pattern in the variance-swap-decomposition literature, where Aït-Sahalia, Karaman, and Mancini (2020) find a significantly priced jump component in variance swap rates and Bollerslev, Todorov, and Xu (2015) attribute a substantial part of the variance premium to left-tail jump variation. The term-structure premium should be positive, since the variance risk premium is concentrated at the short end of the term structure (Dew-Becker, Giglio, Le, and Rodriguez, 2017; Aït-Sahalia et al., 2020): short-horizon variance is priced more aggressively under Q than long-horizon variance, so a long "slope" position earns a positive expected return on average.[23] As an illustration, positive dTRM shocks coincide with sharp drops in short-term uncertainty during market rebounds, which steepen the variance term structure. When it comes to the price of shocks to spot variance, in structural option-pricing models it is uniformly estimated negative, an order of magnitude below the jump-variance price.[24]

For the jump variance and term structure dynamics premiums, our estimates corroborate the existing evidence, if imprecisely: -14.9% on dJV and 13.0% on dTRM, neither significant on its own. Because of the existing evidence suggesting that $\lambda_{\text{dJV}} < \lambda_{\text{dTRM}}$, we evaluate the significance of their difference in a one-sided test: the gap of -27.9%, with a t-statistic of -1.81, rejects equal premiums at the 5% level (10% two-sided). A single variance-dynamics factor would force the two prices to coincide, so this difference rejects the one-factor restriction just as the loading shapes in Section 3.4.1 reject it on the exposure side. The two rejections jointly support our three-factor specification.

The positive estimate of the premium on dSV is somewhat surprising; taken at face value,

[22] Option-pricing models can produce a comparable decomposition, but only under much more restrictive structural assumptions on the underlying dynamics.

[23] Recall that dTRM is the change in the slope of the risk-neutral variance term structure—the change in the long-minus-short-maturity implied-variance differential; see Section 3.2.

[24] See Bakshi, Cao, and Chen (1997), Pan (2002), Eraker (2004), Broadie, Chernov, and Johannes (2007), Broadie, Chernov, and Johannes (2009), and Andersen, Fusari, and Todorov (2015); in each the diffusive premium is small—often indistinguishable from zero—while the jump-variance premium is an order of magnitude larger.

it would imply that investors pay to hedge against negative spot variance shocks. This sign sits at odds with the structural option pricing literature, in which every estimate of the price of risk for the Brownian shock to spot variance is negative—albeit small relative to the price of jump-variance risk.[25] Splitting the sample by FOMC days clarifies the picture. On non-FOMC days, the estimate is close to zero, consistent with the structural literature. On FOMC days, by contrast, the premium turns negative: investors price in increases in diffusive variance ahead of the announcement and these anticipated increases fail to fully materialize. Nonetheless, while the negative pricing of diffusive variance risk on FOMC days is consistent with economic intuition, our estimate of the pricing of jump risk on the same days cuts against it. The dJV premium, strongly negative in the full sample (-14.9%), turns positive on announcement days ($+4.7\%$): investors now price in smaller changes in option-implied jump variance that subsequently materialize, rather than pricing them higher as they do on average.

A plausible source of the asymmetry lies in the structure of option trading. Our factors are constructed from option prices—though not from 0DTEs—and the center of the smile is far more actively traded than the tails (see Figure IA.III in Internet Appendix C), likely because ATM straddles carry much lower transaction costs than deep OTM strikes. Demand therefore concentrates near at-the-money and bids up prices of options on which dSV loads most heavily, rather than in the tails, on which dJV loads. This trading-flow channel highlights a deliberate modeling choice: following the literature, we treat option prices as determined by the underlying in a frictionless market, not by market clearing. Because our estimator interprets all price variation as compensation for factor risk, such demand pressure would be absorbed into the estimated premiums rather than identified as a friction, displacing their signs from a risk-based benchmark and pointing to the same intermediation frictions that Section 5 takes up directly.

[25] See Bakshi, Cao, and Chen (1997), Pan (2002), Eraker (2004), Broadie, Chernov, and Johannes (2007), Broadie, Chernov, and Johannes (2009), and Andersen, Fusari, and Todorov (2015). In each of these studies, the diffusive variance premium is small (often statistically indistinguishable from zero), while the jump-variance premium is an order of magnitude larger and accounts for the bulk of compensation for variance risk.

4.3.3 Implications for expected option returns

To translate these per-factor premiums into option-level expected-return compensation, we examine each factor's contribution to option returns via the product $\widehat{\beta}_{l,i,t} \times \widehat{\lambda}_{l,t}$. The solid (orange) lines in Figure 4 report the time-series average of this product along the moneyness and maturity axes; the dotted (teal) lines instead report the product of unconditional averages, $\overline{\widehat{\beta}}_l \times \overline{\widehat{\lambda}}_l$. By the identity $\overline{\widehat{\beta}_l \widehat{\lambda}_l} = \overline{\widehat{\beta}}_l \, \overline{\widehat{\lambda}}_l + \widehat{\mathrm{Cov}}(\widehat{\beta}_l, \widehat{\lambda}_l)$, the gap between the two lines is the time-series covariance between exposure and premium: if the exposure to a given factor tends to increase (decrease) when its premium increases, the positive (negative) covariance pushes the solid line above (below) the dotted line.

Empirically, the covariance term dominates. For GAM, the product of unconditional means is positive everywhere ($\widehat{\lambda}_{\mathrm{GAM}} > 0$ and $\widehat{\beta}_{\mathrm{GAM}} > 0$ across the surface), yet its time-series average is negative across the moneyness-maturity grid: when the premium rises, exposures fall by enough to overturn the product of the means. For JSK, theory implies a directional pattern: OTM puts have negative exposures and OTM calls positive ones, so crash-averse pricing lowers expected put returns and raises expected call returns. The covariance mostly undoes this pattern however, leaving dollar compensation for JSK much smaller than the unconditional means suggest. The variance-dynamics compensation is sizable at 4–15DTE but largely cancels across the three components, the positive dTRM compensation offsetting the negative dSV and dJV ones, so the total variance premium at these maturities depends on which component dominates at a given time.

Across maturities, variance components carry the bulk of dollar compensation at 4–15DTE; as maturity shortens, the variance contribution dampens and the return-power components (GAM, JSK) grow. At 0DTE, by contrast, every factor's $\widehat{\beta} \times \widehat{\lambda}$ is economically small: short-dated options barely load on non-tradable factors, so whatever is distinctive about their pricing cannot plausibly be a standard risk premium—a puzzle we take up in Section 5. We are the first to present such disaggregate and extensive evidence on the joint pricing of variance and jump dynamics risk in intraday trading of a proxy for the U.S. stock market factor. The rich patterns we document open several avenues for future research.

5. Are the pricing errors economically significant?

In this section, we examine the pricing errors of our factor model across option maturities, moneyness, and time-of-day. At least since Bakshi, Cao, and Chen (1997), the option pricing literature has used the economic patterns in pricing errors as diagnostics of where models could be improved; we follow that tradition here. We estimate time-varying option alphas as described in Section 5.1 below.

5.1 Systematic pricing errors

5.1.1 Conditional intercept and *alpha*

Unless equation (10) in Section 4.1.1 holds, the asset's expected return is no longer solely determined by compensation for systematic risks. The difference between the estimated conditional intercept and $\beta_{NT} \times (-E^Q[F_{NT}])$ is an asset's "alpha"—a pricing error relative to the assumed model:

$$\alpha := c - \beta_{NT} \times (-E^Q[F_{NT}]). \tag{13}$$

We treat the estimated alphas as diagnostics of our model's pricing fit; they describe the model's systematic pricing errors for given type/maturity/moneyness buckets.

Asset-specific alphas are not identified in a single cross-section; since SPX contracts with similar characteristics should have similar alphas, we group options into buckets and estimate the alphas as bucket fixed effects in regression (12).

5.1.2 Implementation

We return to regression (12) in Section 4.1.2 and recall it here:

$$\widehat{c}_{i,t} = \alpha_{b(i),t} - \sum_l \widehat{\beta}_{l,i,t} \times \mu^Q_{l,t} + \nu_{i,t}\,. \tag{14}$$

We now focus on the time-series of estimated $\widehat{\alpha}_{b(i),t}$, the bucket-specific alphas which measure the model's average pricing errors on options with similar characteristics. We obtain these estimates from cross-sectional regressions estimated separately for every 30-minute trading period, using conditional intercepts and option exposures estimated in Section 3. As regression

(12)–(14) does not contain an intercept term (under the null of correct pricing it should be zero), we obtain estimates for every bucket. We further refer to the time-average of $\widehat{\alpha}_{b(i),t}$ as the model's *pricing error* on bucket $b(i)$, and to $\widehat{\alpha}_{b(i),t}$ itself as the bucket *alpha*.

5.2 Pricing errors in the cross-section

As discussed in Section 4.1.2, we estimate systematic pricing errors, $\alpha_{b(i),t}$, by grouping options into buckets based on their type (puts or calls), maturity (0, 1–3 or 4–15 days) and moneyness (ATM, OTM, DOTM). Table IV reports the summary statistics of the estimated alphas in percent per annum. Numerous patterns are of interest. First, the average pricing error (in columns labeled $\overline{\alpha}_{b(i),t}$) is negative in both puts and calls at all maturities. This implies that the model leaves more unexplained return variation on the downside than on the upside.

In 0DTE options, the pricing errors are relatively large and negative—for instance, −4.63% (ATM puts), −3.52% (OTM puts), −3.94% (ATM calls), and −7.73% (OTM calls)—indicating that 0DTE options are more expensive than the factor model implies, taking the pricing of the remaining options across all maturity and moneyness buckets as given. Beyond 0DTE, the pricing errors are smaller in magnitude, generally below 1% in absolute value. The average magnitude of the alphas (in columns labeled $\overline{|\alpha|}_{b(i),t}$) exhibits similar patterns in calls and puts: it is largest in 0DTEs and decreases away from ATM, with the decrease faster in puts. 0DTE alphas are an order of magnitude larger than those at longer maturities. The average magnitude is much higher than the absolute value of the bias, signaling sign changes in the alphas over time, which we explore in the remainder of this section.

The alphas are persistent. In columns of Table IV labeled "$\text{AC}^{30\text{min}}$" and "$\text{AC}^{1\text{day}}$" we report the autocorrelation at 30 minutes (one lag) and at one day (i.e., between values at the same time of day on consecutive trading days). All the one-lag autocorrelations are above 0.90 (the lowest being 0.91 for DOTM 0DTE puts), and the coefficients increase with option maturity. The one-day autocorrelations are lower, with estimates between 0.32 and 0.71 for puts, and between 0.29 and 0.88 for calls. This long-horizon persistence pattern is similar across option maturities and moneyness ranges. Thus, the alphas are not pure noise, but rather they reflect systematic patterns relative to the factor model.

5.3 Time-series variation

Figure 5 plot the time-series of alphas for 0DTE options. The analogous figure for non-0DTE options is Figure IA.VII in Internet Appendix H. In the figure, the left (right) panel presents results for puts (calls). Within each panel moneyness-specific series are distinguished by color. From top to bottom, panels show time-series of alphas based on returns between 09:45–10:15, 12:15–12:45, and 15:15–15:45 ET.

Alphas vary substantially over time, are strongly correlated across option types, maturity and moneyness buckets, and exhibit pronounced negative spikes (when options become particularly expensive) around two types of events: FOMC announcements and periods of market distress. The COVID-19 spike in March 2020 is 2 to 3 times larger than FOMC spikes across most buckets—for instance, for 0DTE options, the annualized alpha reaches 200% of the underlying SPX value. Interestingly, 0DTE alphas become very volatile and spike in both directions starting in March 2022, which coincides with the start of the Fed's rate-hiking cycle, heightened market volatility, and a year-long market decline. This disruption ends in 2023 when a bull market resumes.

As noted in Section 5.2, alphas are more pronounced and more volatile in 0DTE options. For 1+DTE options, the dynamics and average levels of alphas do not vary substantially across different times of the trading day. In contrast, 0DTE alphas exhibit pronounced changes in their time-series patterns under one hour to the options' expiration. In both puts and calls, the alpha corresponding to the 15:15–15:45 ET return is strongly negative and more volatile than earlier in the day.

Overall, we identify two patterns that warrant further investigation: the alphas respond to the macroeconomic environment, and they exhibit intraday seasonality.

5.4 Pricing errors and the economic environment

In this section, we explore the economic mechanisms behind the time-series variation in option alphas. We regress alphas corresponding to 15:15-15:45 ET returns, and the absolute values of these alphas, on variables capturing different aspects of the economic environment:

$$\begin{aligned}\widehat{\alpha}_{b(i)} = \beta_0 + \beta_{HKM} HKM_{d-1} + \beta_{\varepsilon_{HKM}} \varepsilon_{HKM,d} + \beta_{ESF} ESF_{d-1} \\ + \beta_{OI} OI_{b,d-1} + \beta_{CIV} CIV_b + \beta_{SP500} r_{SP500} + \varepsilon_{b(i)} \,. \end{aligned} \tag{15}$$

In the regression (15), variables bearing the subscript "d" are measured at 15:15 ET on the same trading day as $\widehat{\alpha}_{b(i)}$, those with the subscript "$d-1$" are measured at the same time of the preceding trading day, and r_{SP500}, without a subscript, is the return from 09:45 to 15:15 ET on the same day. We estimate the regression separately for each maturity bucket and option type on daily observations, with Newey–West standard errors accounting for the persistence of the alphas documented in Table IV (see Internet Appendix F). In addition to reporting the results from the above multivariate regression for ATM options, the top panel Table V also reports the R^2 from univariate regressions of each variable. The bottom panel of the table reports similar results, but with $\left|\widehat{\alpha}_{b(i),t}\right|$ as the dependent variable.[26]

The definitions of all variables are given below. In brief, HKM, ε_{HKM}, and ESF proxy for funding conditions in financial markets, OI proxies for option demand and inventories, and CIV_b and r_{SP500} we interpret as controls.

5.4.1 Funding conditions

Trading options requires funding to maintain their delta hedge and balance-sheet capacity to bear the residual risk. In this section, we test whether shocks to either pass through to alphas.

First, we follow Siriwardane, Sunderam, and Wallen (2025) and construct an equity spot-futures (ESF) spread from the difference between an ES futures-implied forward interest rate and the 3-month OIS rate (we provide details in Internet Appendix G). A nonzero ESF signals that funding costs in the equity index futures market differ from those implied by the OIS rate. Siriwardane, Sunderam, and Wallen (2025) interpret this as a measure of limits to arbitrage, and Hazelkorn, Moskowitz, and Vasudevan (2023) also document that ESF measures liquidity demand for exposure to global equity markets. While the fact that ESF co-moves only partially with spreads in other segments is central to Siriwardane et al.'s conclusion that arbitrage capital is not freely mobile across markets, ES futures are the primary delta-hedging environment for option traders. Table V reports that as ESF increases, alphas on 0DTE puts and calls become increasingly negative. This is consistent with 0DTE options being more expensive when funding conditions tighten. The quantitative impact and explanatory power are, however, limited: at most 2.5%, and ESF is only significant in

[26] Tables IA.VII to IA.XI, discussed in Internet Appendix H, demonstrate that results are robust across several dimensions.

longer-maturity calls. Similarly, funding costs in the futures market have little impact on the magnitudes of alphas, with a statistically significant impact only in 0 and 1–3DTE ATM calls. In results for (D)OTM options reported in Internet Appendix H, we obtain consistently negative and significant estimates for calls. Because of the close hedging relationship between the SPX option and ES futures markets, *ESF* may itself reflect option hedging demand. We therefore instrument *ESF* using the corresponding spread for the Russell 2000 E-Mini futures and report results in Table IA.VI in Internet Appendix H. The results are broadly similar to the baseline regressions.

Second, we include the lagged aggregate market equity ratio of NY Fed primary dealers, *HKM*, that we refer to as the capital ratio, and the contemporaneous realization of the associated risk factor, ε_{HKM}—the non-predictable innovation to the capital ratio, expressed as a growth rate—defined by He, Kelly, and Manela (2017). While the primary dealers are not market makers in SPX options anymore, some of them played this role for a significant fraction of our sample. They also remain central to the over-the-counter derivatives markets, where they are likely to hold short volatility exposure, and hedge partly through SPX options.[27] We find that the capital ratio, *HKM*, significantly predicts alphas in 1–3DTE options and 4–15DTE calls, with positive coefficient estimates indicating that higher capital ratios—signaling better risk-bearing capacity—are associated with more positive alphas, i.e., cheaper options. The capital ratio has the strongest explanatory power (11.1%–19.6%) at 1–3DTE in both puts and calls; at 0DTE the relationship is insignificant. The contemporaneous evidence is weaker. The coefficient on the intermediary risk factor, ε_{HKM}, is positive at the shortest maturity, but only significant in calls, with a small univariate R^2 of 0.9%.

5.4.2 Open interest

Open interest is informative about two aspects of option trading activity: first, the concentration of positions and inventories in a given bucket; second, bucket-level option demand. We measure open interest in options in a given bucket at the end of the preceding trading day. For instance, for 0DTE options, we take the prior day's *OI* in options with one day to expiration.

[27] Goldman Sachs, for instance, shut down their SPX market-making operation in 2017 [Bloomberg]; the OTC equity derivatives market is about $7T in notional terms, with trading activity split evenly between swaps and options [ISDA, 2024].

Although net option supply is zero (at market as well as bucket level), the resulting risk exposures are supported by specific investor wealth, which is likely differentially aggregated on the buy- and sell-sides. Margin requirements bind on the sell side, and more so on concentrated short positions, pushing option prices up and alphas down. An equivalent rise in OI driven by the buy-side pushes prices in the same way.

Our results indeed support this narrative: the coefficient on open interest is statistically significant in both puts and calls: Options become more expensive when aggregate inventories are higher. Economically, however, the effect is concentrated in puts, largest at 0DTE, and weakens with maturity. The starkest difference between the two option types lies in the explanatory power of open interest: in puts, the univariate R^2 ranges from 12.6% at 1–3DTE to 14.6% at 0DTE, while in calls it is close to zero except for the longest maturity, where it reaches 5.1%. The estimates from regressions of absolute alphas confirm that open interest matters most at the shortest maturity: the univariate R^2 is 16.7% at 0DTE in puts, while it is 6.1% at 1–3DTE. In calls, only the 0DTE coefficient is significant—and positive, with a modest R^2 of 1.5%. To summarize: higher open interest predicts more expensive options in both puts and calls, with economically meaningful magnitudes concentrated in puts at the shortest maturity. The pattern is consistent with both readings of OI—sell-side concentration of standing positions and buy-side demand pressure—which our regression does not separate.

5.4.3 Market controls

Finally, we directly control for the conditions in the equity and the SPX option markets. The variable r_{SP500} captures the overall market return on the day, from open until 15:15 ET (the start of our alpha measurement period), which is likely correlated with delta-hedging activity in the option market. We also include the day's CIV index, computed from options at a given maturity bucket at 15:15 ET. CIV is a proxy for a general level of option prices. Whereas we refrain from interpreting its coefficient, we note that it has significant explanatory power (34.3%–65.1% R^2 in all regressions) for both alphas and their magnitudes across all option types and maturities. While this is to be expected, as CIV is a measure of how expensive options are in general, and hence its calculation integrates the alphas themselves, it also indicates that our results on intermediary variables and open interest are not driven by omitted-variable bias related to overall option expensiveness.

All in all, these results suggest that the pricing errors of short-maturity options are related

to the diverse roles of the market participants, and to their ability to absorb risk.

5.5 Intraday seasonality

In Figure 6, we present a summary of the intraday seasonality of pricing errors. At each time of the trading day, we compute the average pricing error across the sample period for each option bucket. A strong seasonal pattern is only present in 0DTE options. Pricing errors are initially negative for all moneyness categories at market open, except DOTM 0DTE puts, between -5% and -20% per-year basis. They subsequently increase steadily until 13:15 ET, reaching positive values of up to 5%. Then, they decline rapidly to reach values more negative than -20% at the last 30-minute interval before expiration, from 15:15 to 15:45 ET.

Decomposing the 0DTE pattern into FOMC and non-FOMC days (Figure IA.VIII, in the Internet Appendix) shows that pricing errors are negative throughout the trading day on FOMC days (DOTM puts excepted), consistent with elevated demand for downside protection around monetary policy announcements, while the inverted U-shape is entirely a non-FOMC phenomenon. The seasonality is also specific to the shortest maturity: 1–3DTE pricing errors show minimal intraday seasonality in every subsample we examine, including before and after the April 2022 introduction of daily SPX expirations (Figure IA.IX).

5.6 Economic significance of the pricing errors

Finally, we assess the economic significance of 0DTE alphas by constructing self-financing, six-factor-neutral portfolios. As aforementioned, the alphas could be symptomatic of a priced systematic factor being omitted in the model. If, on the other hand, our factor set were complete, the alphas would reflect genuine mispricing by market participants. In both cases, a trading strategy guided by the pricing errors should lead to significant profits. In the omitted factor case, these profits would ensue from reaping the premium associated with the factor. Under the mispricing narrative, if the profits were robust to realistic transaction costs, we could argue that limits to arbitrage prevent participants from undoing the mispricing. Finally, it is also possible that the pricing errors are simply too small to exploit, in which case the model's fit, while imperfect, is economically adequate.

For each 0DTE option at each 30-minute timestamp, we hedge its six factor exposures using a basket of 18 longer-maturity instrument contracts: from each of six buckets defined

by maturity (1–3DTE, 4–15DTE) and moneyness (ATM, OTM, DOTM), we select the three contracts with the highest trading volume of the preceding 30-minute interval, requiring the volume to be positive. The hedging weights are obtained by projecting each target option's exposure vector onto the factor-loading matrix of the instruments:

$$\boldsymbol{w}_{i,t} = \left(\widehat{\boldsymbol{B}}_t \widehat{\boldsymbol{B}}_t'\right)^{\dagger} \widehat{\boldsymbol{B}}_t \boldsymbol{e}_{i,t}\,, \tag{16}$$

where (i) $\widehat{\boldsymbol{B}}_t$ combines the matrix of factor loadings (18×5) with a 6^{th} column containing the time-t cost of opening each instrument's DLT-hedged position, $O_{j,t} - \widehat{\beta}_{\text{DLT},j,t} F_t$, with $O_{j,t}$ the option mid-quote, $\widehat{\beta}_{\text{DLT},j,t}$ the model delta, and F_t the ES futures price, (ii) $\boldsymbol{e}_{i,t}$ is the (6×1)-vector with the additive inverse of (a) the 5 factor loadings and (b) the analogous DLT-hedged position cost of target option i, at time t, and (iii) $\dagger$ denotes the Moore–Penrose pseudo-inverse. The sixth restriction sets the combined time-t cost of the target and instrument positions to zero, so the portfolio requires no initial outlay: it is self-financing by construction. Every element of $\widehat{\boldsymbol{B}}_t$ and $\boldsymbol{e}_{i,t}$ is known at time t, so the weights are implementable ex ante. With 18 instruments and six restrictions, the system is underdetermined; among the weight vectors satisfying the six restrictions exactly, the Moore–Penrose solution selects the one with the smallest norm, which caps the gross size of the hedge positions.

A 0DTE option's 6-factor-neutral return is calculated as

$$r^{N}_{i,t+1} = r^{dh}_{i,t+1} + \boldsymbol{w}'_{i,t} \boldsymbol{r}^{H}_{i,t+1} \tag{17}$$

where $r^{dh}_{i,t+1}$ is time-$t+1$ DLT-hedged return of 0DTE option i, and $\boldsymbol{r}^{H}_{i,t+1}$ is the vector of DLT-hedged returns of the hedging instruments. We average the factor-neutral returns cross-sectionally and then over time. Table VI reports their summary statistics.

The delta-hedged returns of the target 0DTE options (in rows labeled "T") are significantly negative across all moneyness buckets, with Sharpe ratios between -1.0 and -1.35. After hedging, the magnitude of factor-neutral mean returns (rows labeled "N") shrinks substantially: for ATM options, from -0.61% to -0.13% (puts) and from -0.54% to -0.08% (calls), confirming that a large share of the returns on 0DTE is attributable to the six factors, even when the model is not at its best. Nonetheless, the residual factor-neutral returns remain

negative and statistically significant in all three put buckets, with Sharpe ratios between -0.27 and -0.45. Hence, part of the apparent 0DTE overpricing survives factor-neutral hedging. In calls, the residual returns are smaller: insignificant at ATM, and marginally significant, at the 10% level, at OTM and DOTM.

The intraday pattern of these factor-neutral returns is consistent with the seasonality of the pricing errors documented above: negative at the beginning of the day, positive early afternoon, and strongly negative in the last 30 minutes before expiration, with only the late-afternoon leg statistically significant (Figures IA.X and IA.XI).

The key question is whether these returns survive transaction costs. Figure 7 presents the cumulative returns of *shorting* the six-factor-neutral portfolio between 15:15 and 15:45 ET—i.e., selling 0DTE options and hedging with longer-dated instruments—under two scenarios. When the strategy is executed at mid-quote prices (green line), it generates significant positive cumulative returns across all three moneyness buckets over the sample period. The grey line serves as a reference: it is the profit of a long position in the ES futures over the same 30-minutes interval. Whereas the futures position leads to virtually no cumulative profit over our entire sample, the frictionless Sharpe ratio of our trading strategy is 1.4 in annualized terms (aggregating across moneyness levels).

When, instead, the 0DTE options are shorted at the bid price and the hedging instruments are bought (shorted) at the ask (bid) price—i.e., when the full quoted spread is paid on every leg—the cumulative return becomes substantially negative (purple line). Muravyev and Pearson (2020) show that conventional estimates of option trading costs overstate what sophisticated traders actually pay: accounting for execution timing, the effective half-spread on S&P 500 index options is approximately 29% of the quoted half-spread for algorithmic traders, and 58% when averaged over all market participants. In our setting, the factor-neutral strategy's cumulative return reaches zero when the effective half-spread is just 5% of the quoted half-spread (not in the figure). This is an order of magnitude below the most optimistic estimate from Muravyev and Pearson. The pricing errors of our model are therefore not exploitable by this exact-neutralization strategy under any realistic transaction-cost scenario. The claim is deliberately narrow: the weights in equation (16) enforce exact six-factor neutrality and are derived neglecting spreads, which are then charged on all 19 legs. We cannot exclude that a design trading off residual factor exposure against execution costs would fare better. Such a design, however, retains open factor positions: part of any Sharpe

ratio it earns is compensation for factor risk rather than gains from pricing errors, and the gap between the "T" and "N" rows in Table VI shows that this compensation is far from trivial. This places our result squarely in the tradition of Novy-Marx and Velikov (2016), who show that many documented anomalies in equity returns disappear once implementation costs are accounted for.

Whereas the economic magnitude of the pricing errors is modest, the frictionless strategy being profitable nonetheless hints toward possible improvements of the model. Results reported in the appendix (Figure IA.VIII), for instance, suggest that instrumenting betas to account for event risk (essentially jumps whose timing is known in advance, Johannes, Kaeck, and Seeger, 2023) could reduce pricing errors around macroeconomic announcements. In regression (15), the strong significance of prevailing open interest and, more broadly, of variables linked to funding costs suggest that allowing the model to condition on market makers' intraday inventory could prove fruitful. We leave these as avenues for future research.

6. Conclusion

Our study documents that the returns on ultra-short-maturity (USM) SPX options, and in particular zero-days-to-expiration (0DTE) contracts, are driven by a compact yet rich factor structure that combines powers of underlying returns with a decomposition of innovations in variance dynamics. The factor-model approach also enables recovering intraday time series of the factors' risk-neutral expectations and, from these, their risk premiums, thereby delivering a new perspective on the option market stochastic discount factor at high frequency. The time-series of Q-expectations are interesting objects of future study in their own right, as they embed real-time market perception of risk premiums and go beyond what is accessible via the static option replication approach.

Within the same framework, we examine the model's systematic pricing errors, which are concentrated in 0DTE options and exhibit pronounced intraday seasonality: options are underpriced at the open, overpriced around midday, and underpriced again near the close. These patterns are diagnostically rich—linked to option inventories and, on FOMC announcement days, consistent with the model's inability to accommodate event risk with known timing. However, a six-factor-neutral trading strategy designed to exploit the pricing errors does not survive realistic transaction costs: the residual pricing errors are too small to

reap with an exact factor-neutral hedge once spreads are paid. The frictionless profitability of the strategy nonetheless points to specific avenues for model improvement: instrumenting factor exposures with event-day indicators and conditioning on market makers' intraday inventory.

Our findings contribute to several areas of finance research. First, we extend the study of ultra-short-maturity options by documenting the economic sources of their returns and by quantifying how their risk compensation evolves throughout the day. Second, by estimating the Q-moments of return and volatility factors at high frequency, we show that conditional factor models can yield interpretable and robust insights into risk pricing in markets where exposures are complex and time-varying. Third, we provide a comprehensive diagnostic of where structurally-motivated option pricing models fall short at the ultra-short horizon, identifying event risk and inventory pressure as the two leading candidates for model enrichment. Taken together, these results offer an empirical framework to study the risk and pricing of 0DTE and other ultra-short-maturity options, and a foundation on which richer specifications can be built.

REFERENCES

Adams, Greg, Chukwuma Dim, Bjørn Eraker, Jean-Sébastien Fontaine, Chayawat Ornthanalai, and Grigory Vilkov, 2025, Do S&P 500 Options Increase Market Volatility? Evidence from 0DTEs, Working paper.

Aït-Sahalia, Yacine, Mustafa Karaman, and Loriano Mancini, 2020, The term structure of equity and variance risk premia, *Journal of Econometrics* In Press.

Almeida, Caio, Gustavo Freire, and Rodrigo Hizmeri, 2025, 0DTE Asset Pricing, *Proceedings of the EUROFIDAI-ESSEC Paris December Finance Meeting* .

Andersen, Torben G., and Tim Bollerslev, 1997, Intraday periodicity and volatility persistence in financial markets, *Journal of Empirical Finance* 4, 115–158.

Andersen, Torben G., Oleg Bondarenko, and Maria T Gonzalez-Perez, 2015, Exploring Return Dynamics via Corridor Implied Volatility, *Review of Financial Studies* 28, 2902–2945.

Andersen, Torben G., Nicola Fusari, and Viktor Todorov, 2015, The Risk Premia Embedded in Index Options, *Journal of Financial Economics* 117, 558–584.

Bakshi, Gurdip, Charles Cao, and Zhiwu Chen, 1997, Empirical Performance of Alternative Option Pricing Models, *The Journal of Finance* 52, 2003–2049.

Bakshi, Gurdip, and Nikunj Kapadia, 2003, Delta-Hedged Gains and the Negative Market Volatility Risk Premium, *Review of Financial Studies* 16, 527–566.

Bakshi, Gurdip, Nikunj Kapadia, and Dilip Madan, 2003, Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options, *Review of Financial Studies* 16, 101–143.

Bali, Turan G., and Scott Murray, 2013, Does Risk-Neutral Skewness Predict the Cross-Section of Equity Option Portfolio Returns?, *Journal of Financial and Quantitative Analysis* 48, 1145–1171.

Bandi, Federico M, Nicola Fusari, and Roberto Renò, 2026, 0dte option pricing, *Journal of Finance, Forthcoming* .

Bollerslev, Tim, and Viktor Todorov, 2011, Tails, fears, and risk premia, *The Journal of Finance* 66, 2165–2211.

Bollerslev, Tim, Viktor Todorov, and Lai Xu, 2015, Tail risk premia and return predictability, *Journal of Financial Economics* 118, 113–134.

Broadie, Mark, Mikhail Chernov, and Michael Johannes, 2007, Model Specification and Risk Premia: Evidence from Futures Options, *The Journal of Finance* 62, 1453–1490.

Broadie, Mark, Mikhail Chernov, and Michael Johannes, 2009, Understanding Index Option Returns, *Review of Financial Studies* 22, 4493–4529.

Büchner, Matthias, and Bryan Kelly, 2022, A Factor Model for Option Returns, *Journal of Financial Economics* 143, 1140–1161.

Carr, Peter, and Liuren Wu, 2020, Option Profit and Loss Attribution and Pricing: A New Framework, *Journal of Finance* 75, 2271–2316.

Chen, Hui, Yuhan Cheng, Yanchu Liu, and Ke Tang, 2023, Teaching Economics to the Machines, *Available at SSRN 6093305* NBER Working Paper 34713.

Chen, Hui, Scott Joslin, and Sophie Xiaoyan Ni, 2019, Demand for Crash Insurance, Intermediary Constraints, and Risk Premia in Financial Markets, *The Review of Financial Studies* 32, 228–265.

Christoffersen, Peter, Bruno Feunou, Yoontae Jeon, and Chayawat Ornthanalai, 2021, Time-Varying Crash Risk Embedded in Index Options: The Role of Stock Market Liquidity, *Review of Finance* 25, 1261–1298.

Christoffersen, Peter, Mathieu Fournier, and Kris Jacobs, 2018, The Factor Structure in Equity Options, *The Review of Financial Studies* 31, 595–637.

Christoffersen, Peter, Kris Jacobs, and Chayawat Ornthanalai, 2012, Dynamic Jump Intensities and Risk Premiums: Evidence from S&P500 Returns and Options, *Journal of Financial Economics* 106, 447–472.

Dennis, Patrick, and Stewart Mayhew, 2002, Risk-neutral skewness: Evidence from stock options, *Journal of Financial and Quantitative Analysis* 37, 471–493.

Dew-Becker, Ian, Stefano Giglio, Anh Le, and Marius Rodriguez, 2017, The Price of Variance Risk, *Journal of Financial Economics* 123, 225–250.

Dickerson, Alexander, Mathieu Fournier, Kris Jacobs, and Piotr Orłowski, 2025, Systematic variance risk everywhere in equity option markets, *HEC Montréal / Canadian Derivatives Institute Working Paper* .

Dickerson, Alexander, Philippe Mueller, and Cesare Robotti, 2023, Priced risk in corporate bonds, *Journal of Financial Economics* 150, 103707.

Duarte, Jefferson, Christopher S Jones, Mehdi Khorram, Haitao Mo, and Junbo L. Wang, 2026, Too Good to Be True: Look-ahead Bias in Empirical Options Research, *The Review of Financial Studies* Forthcoming.

Duarte, Jefferson, Christopher S. Jones, and Junbo L. Wang, 2024, Very noisy option prices and inference regarding the volatility risk premium, *The Journal of Finance* 79, 3581–3621.

Egloff, Daniel, Markus Leippold, and Liuren Wu, 2010, The term structure of variance swap rates and optimal variance swap investments, *Journal of Financial and Quantitative Analysis* .

Eraker, Bjørn, 2004, Do Stock Prices and Volatility Jump? Evidence From Spot and Option Prices, *Journal of Finance* 59, 1367–1403.

Fournier, Mathieu, Kris Jacobs, and Piotr Orłowski, 2024, Modeling conditional factor risk premia implied by index option returns, *The Journal of Finance* 79, 2289–2338.

Fu, Lei, Su Li, David K. Musto, and Neil D. Pearson, 2024, Trading Like There's No Tomorrow: Uncovering the S&P 500 0DTE Options Market Dynamics

Gagliardini, Patrick, Elisa Ossola, and Olivier Scaillet, 2016, Time-Varying Risk Premium in Large Cross-Sectional Equity Data Sets, *Econometrica* 84, 985–1046.

Gârleanu, Nicolae, Lasse Heje Pedersen, and Allen M. Poteshman, 2009, Demand-Based Option Pricing, *Review of Financial Studies* 22, 4259–4299.

Harvey, Campbell R, and Akhtar Siddique, 2000, Conditional skewness in asset pricing tests, *Journal of Finance* 55, 1263–1295.

Hazelkorn, Todd M., Tobias J. Moskowitz, and Kaushik Vasudevan, 2023, Beyond basis basics: Liquidity demand and deviations from the law of one price, *The Journal of Finance* 78, 301–345.

He, Zhiguo, Bryan Kelly, and Asaf Manela, 2017, Intermediary asset pricing: New evidence from many asset classes, *Journal of Financial Economics* 126, 1–35.

Horenstein, Alex R., Aurelio Vasquez, and Xiao Xiao, 2026, Common Factors in Equity Option Returns, *Review of Financial Studies* 39, 835–874.

Jacod, J, and P Protter, 2012, *Discretization of Processes*, volume 67 of *Stochastic Modelling and Applied Probability* (Springer-Verlag, Heidelberg).

Jagannathan, Ravi, and Zhenyu Wang, 1996, The conditional CAPM and the cross-section of expected returns, *Journal of Finance* 51, 3–53.

Jiang, George J, and Yisong S Tian, 2005, The Model-Free Implied Volatility and Its Information Content, *Review of Financial Studies* 18, 1305–1342.

Johannes, Michael Slater, Andreas Kaeck, and Norman Seeger, 2023, FOMC Announcement Event Risk, *Available at SSRN 4484011* .

Jones, Christopher S., 2006, A Nonlinear Factor Analysis of S&P 500 Index Option Returns, *Journal of Finance* 61, 2325–2363.

Jones, Christopher S., and Joshua Shemesh, 2018, Option Mispricing around Nontrading Periods, *The Journal of Finance* 73, 861–900.

Kozhan, Roman, Anthony Neuberger, and Paul Schneider, 2013, The Skew Risk Premium in the Equity Index Market, *Review of Financial Studies* 26, 2174–2203.

Kraus, Alan, and Robert H. Litzenberger, 1976, Skewness Preference and the Valuation of Risk Assets, *Journal of Finance* 31, 1085–1100.

Lipson, Marc L., Davide Tomio, and Jiang Zhang, 2023, A Real Cost of Free Trades: Retail Option Trading Increases the Volatility of Underlying Securities, Working paper, Darden School of Business.

Martin, Ian, 2017, What is the Expected Return on the Market?*, *The Quarterly Journal of Economics* 132, 367–433.

Muravyev, Dmitriy, and Xuechuan (Charles) Ni, 2020, Why do option returns change sign

from day to night?, *Journal of Financial Economics* 136, 219–238.

Muravyev, Dmitriy, and Neil D. Pearson, 2020, Options Trading Costs Are Lower than You Think, *The Review of Financial Studies* 33, 4973–5014.

Nagel, Stefan, and Kenneth J. Singleton, 2011, Estimation and Evaluation of Conditional Asset Pricing Models, *Journal of Finance* 66, 873–909.

Ni, Shangkun X., Neil D. Pearson, Allen M. Poteshman, and Joshua White, 2021, Does Option Trading Have a Pervasive Impact on Underlying Stock Prices?, *Review of Financial Studies* 34, 1952–1986.

Novy-Marx, Robert, and Mihail Velikov, 2016, A Taxonomy of Anomalies and Their Trading Costs, *Review of Financial Studies* 29, 104–147.

Orłowski, Piotr, Paul Schneider, and Fabio Trojani, 2024, On the nature of (jump) skewness risk premia, *Management Science* 70, 1154–1174.

Pan, Jun, 2002, The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study, *Journal of Financial Economics* 63, 3–50.

Savor, Pavel, and Mungo Wilson, 2013, How Much Do Investors Care About Macroeconomic Risk? Evidence from Scheduled Economic Announcements, *Journal of Financial and Quantitative Analysis* 48, 343–375.

Schneider, Paul, 2015, Generalized Risk Premia, *Journal of Financial Economics* 116, 487–504.

Siriwardane, Emil, Aditya Sunderam, and Jonathan Wallen, 2025, Segmented Arbitrage, *The Journal of Finance* 80, 2543–2590.

Todorov, Viktor, 2019, Nonparametric spot volatility from options, *The Annals of Applied Probability* 29, 3590–3636.

Vasquez, Aurelio, 2017, Equity Volatility Term Structures and the Cross Section of Option Returns, *Journal of Financial and Quantitative Analysis* 52, 2727–2754.

Vasquez, Aurelio, Diego Amaya, Neil D. Pearson, and Pedro Angel Garcia-Ares, 2025, 0dte index options and market volatility: How large is their impact?, *SSRN Electronic Journal* .

Vilkov, Grigory, 2023, 0DTE Trading Rules, *Available at SSRN 4641356* .

Figure 1. Time-series of Estimated Conditional Risk-neutral Expectations of Risk Factors

The figure presents daily time-series of the estimated conditional risk-neutral expectations of the risk factors; the top row shows estimates at 9:45 ET (the beginning of the first intraday return interval) and the bottom row at 15:15 ET (the beginning of the last). They are obtained from the cross-sectional regression (12) of the conditional intercepts on option exposures at each 30-minute timestamp, estimated jointly for puts and calls with type-specific bucket fixed effects. The annotation in the μ^{Q}_{GAM} panels reports the correlation between μ^{Q}_{GAM} and the contemporaneous short-maturity corridor implied variance. The ordinate axis labels are in annualized percentage points; the units are factor-specific, as in Table II. The sample period is from 2016/09 to 2023/12.

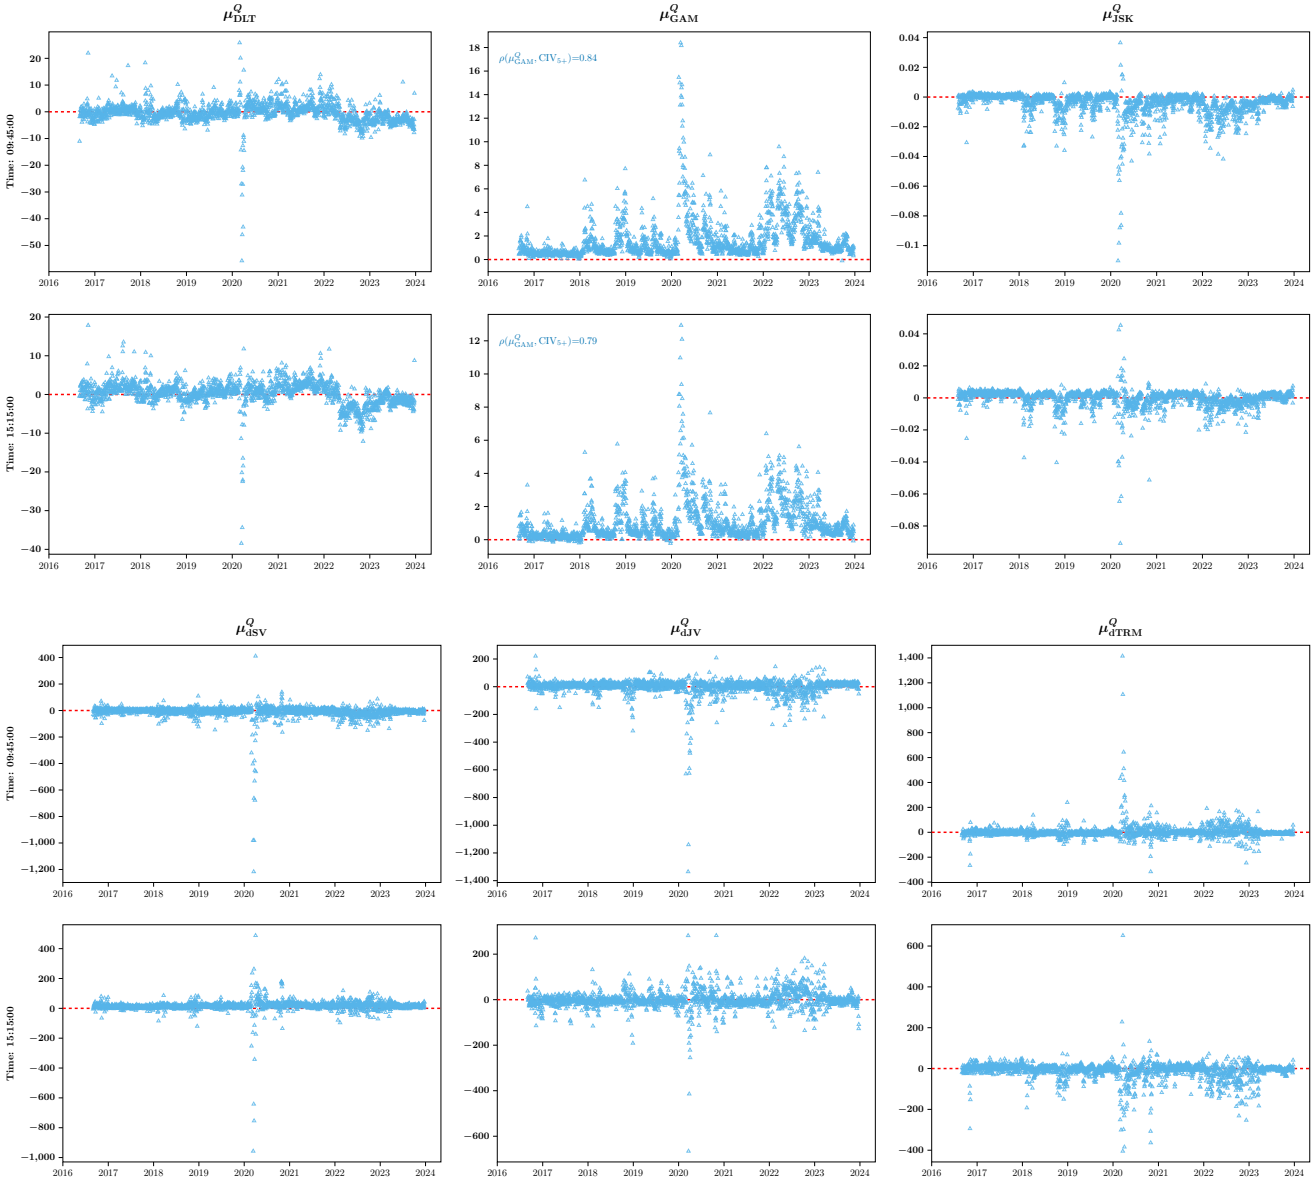


Figure 2. The Intraday Pattern of Estimated Conditional Risk-neutral Expectations of Risk Factors

The figure plots the intraday average conditional risk-neutral expectation, μ_t^Q, of each risk factor. The estimates are obtained from the cross-sectional regression (12) of the conditional intercepts on option exposures at each 30-minute timestamp, jointly for puts and calls with type-specific bucket fixed effects, and then averaged across days at each intraday timestamp. The ordinate axis labels are in annualized percentage points; the units are factor-specific, as in Table II. The sample period is from 2016/09 to 2023/12.

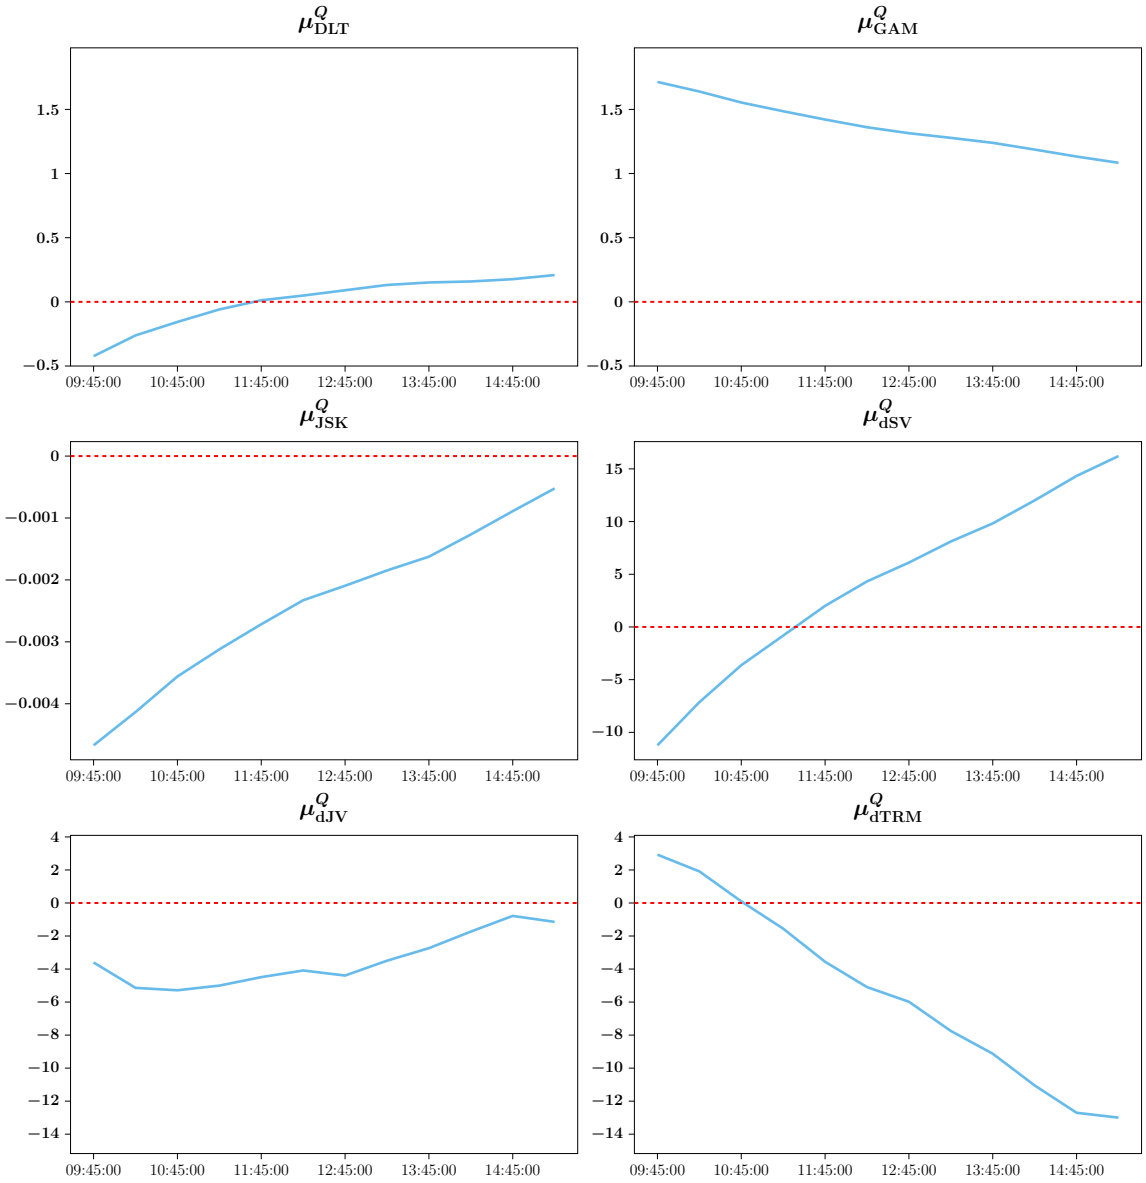


Figure 3. The Fitted Exposures across Relative Strikes

This figure plots the averages of the fitted exposures of puts and calls across relative strikes for 0, 1, 5, and 15DTE, respectively. For each factor, we evaluate the fitted exposure surface at each relative strike grid point, using the kernel-smoothed BMS implied volatility at that point as the implied-volatility characteristic. The fitted BMS implied volatility smile is obtained via kernel regression at each timestamp and maturity, where the grid is trimmed to the range of observed relative strikes $[\kappa_{t,\min}, \kappa_{t,\max}]$ at each timestamp to avoid extrapolation. The relative strike grids are from -2 to 1 for puts and from -0.5 to 1 for calls. The averages are computed by averaging the time-series values at each relative strike grid point. The sample period is from 2016/09 to 2023/12.

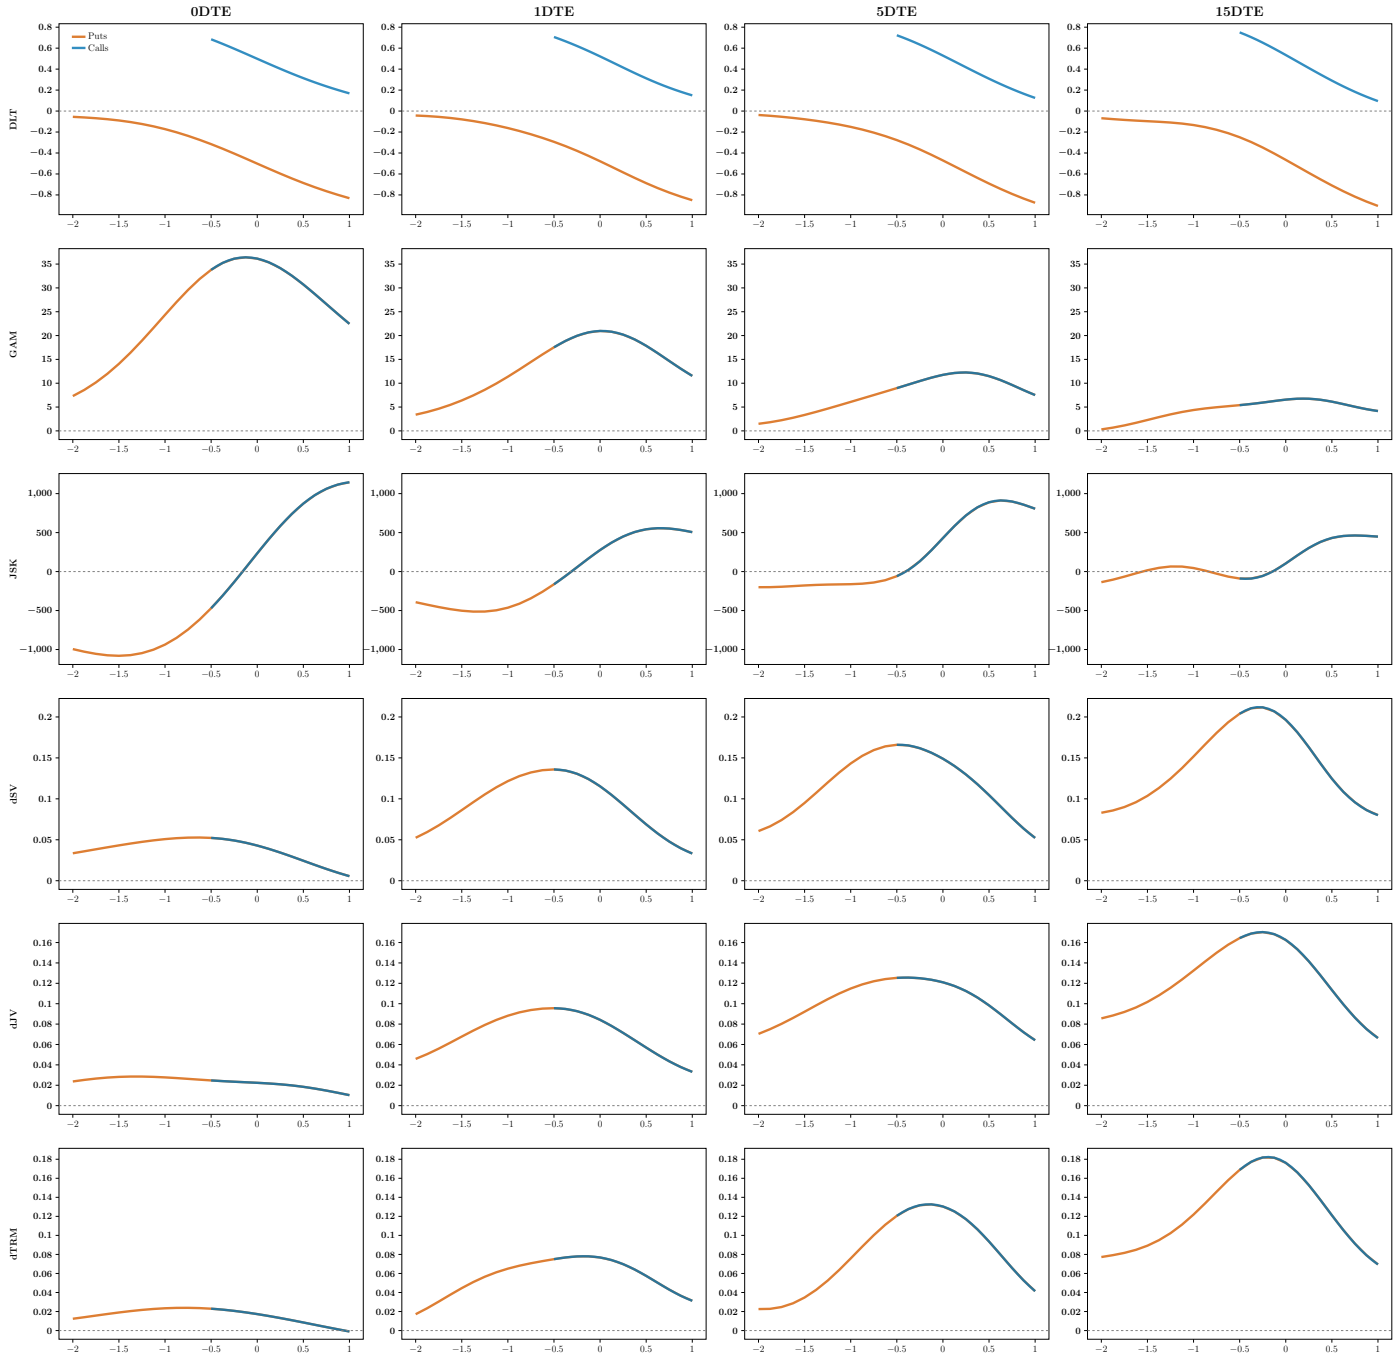


Figure 4. The Product of Exposures and Risk Premiums across Relative Strikes

This figure plots the time-series averages of the product of the fitted exposures and the realized risk premiums ($\widehat{\beta}_{i,t} \times \widehat{\lambda}_t$; solid lines), and the product of the averages of the fitted exposures and of the risk premiums ($\overline{\beta} \times \overline{\lambda}$; dashed lines), across relative strikes for 0, 1, 5, and 15DTE, respectively, for the five factors other than DLT. For each factor, we evaluate the fitted exposure surface at each relative strike grid point, using the kernel-smoothed BMS implied volatility at that point as the implied-volatility characteristic. The fitted BMS implied volatility smile is obtained via kernel regression at each timestamp and maturity, where the grid is trimmed to the range of observed relative strikes $[\kappa_{t,\min}, \kappa_{t,\max}]$ at each timestamp to avoid extrapolation. The relative strike grids are from -2 to 1 for puts. The realized risk premiums, $\widehat{\lambda}_t$, are computed as the difference between the factor realization and its conditional risk-neutral expectation, $\widehat{\lambda}_t = f_{t+1} - \widehat{E}_t^Q[f_{t+1}]$. The averages are computed by averaging the time-series values at each relative strike grid point. All the values are reported in percent per annum (multiplied by 12×252). The sample period is from 2016/09 to 2023/12.

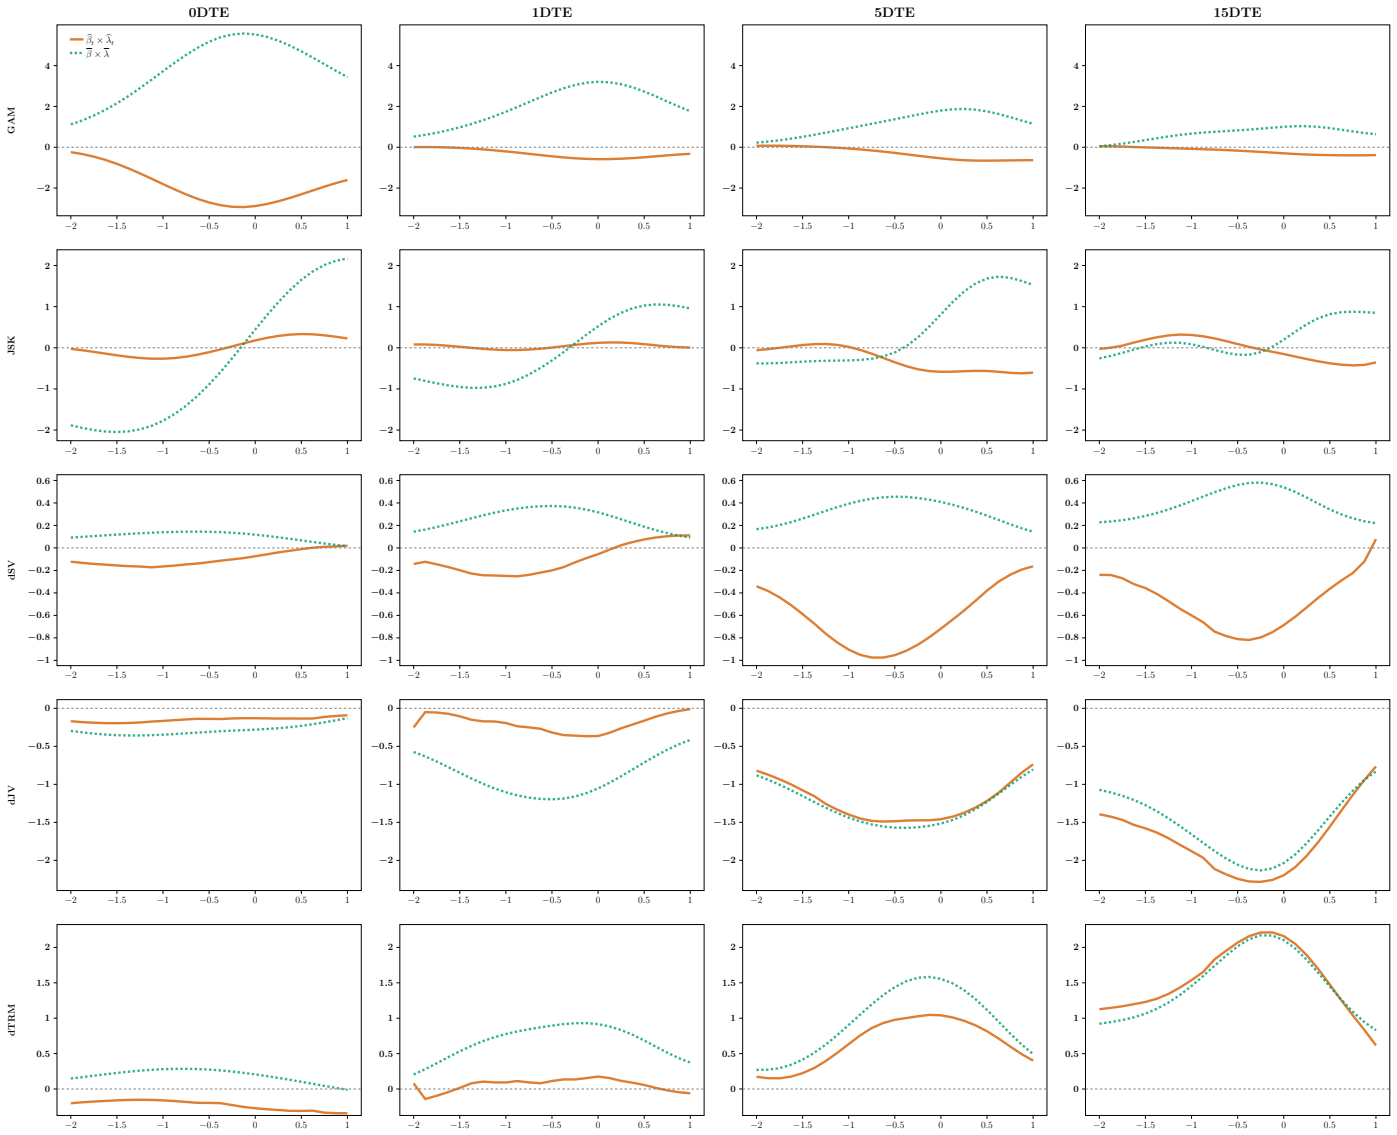


Figure 5. Time-series of 0DTE Estimated Alphas

The figure plots the time-series of 0DTE bucket-specific alphas, $\widehat{\alpha}_{b(i),t}$. They are obtained from the cross-sectional regression (14) of the conditional intercepts on option exposures at each 30-minute timestamp for puts and calls, respectively. Within each option type, the options are assigned to one of nine buckets based on maturity (0DTE, 1–3DTE, 4–15DTE) and moneyness (DOTM, OTM, ATM), for 18 type-specific buckets in total. This figure displays the estimated $\widehat{\alpha}_{b(i),t}$ for the 0DTE buckets, estimated from returns between 09:45–10:15 ET, 12:15–12:45 ET, and 15:15–15:45 ET, respectively (top, middle, and bottom row). The left (right) column reports results for puts (calls); within each panel, colors distinguish the moneyness buckets. The vertical red dashed line marks April 1, 2022; the new Tuesday and Thursday expirations began trading on April 18 and May 11, 2022, respectively. The vertical gray dashed lines indicate FOMC announcement days, including scheduled and emergency announcements. The ordinate axis labels are in annualized percentage points. The sample period is from 2016/09 to 2023/12.



Figure 6. The Intraday Pattern of Estimated Alphas

The figure plots the intraday average bucket-specific alphas, $\widehat{\alpha}_{b(i),t}$, at each intraday timestamp across days. The bucket-specific alphas are estimated from the cross-sectional regression (14) of the conditional intercepts on option exposures for puts and calls, at each 30-minute timestamp. Within each option type, the options are assigned to one of nine buckets based on maturity (0DTE, 1–3DTE, 4–15DTE) and moneyness (DOTM, OTM, ATM), for 18 type-specific buckets in total. Columns correspond to the maturity buckets; the top (bottom) row reports puts (calls); within each panel, colors distinguish the moneyness buckets. The ordinate axis labels are in annualized percentage points. The sample period is from 2016/09 to 2023/12.

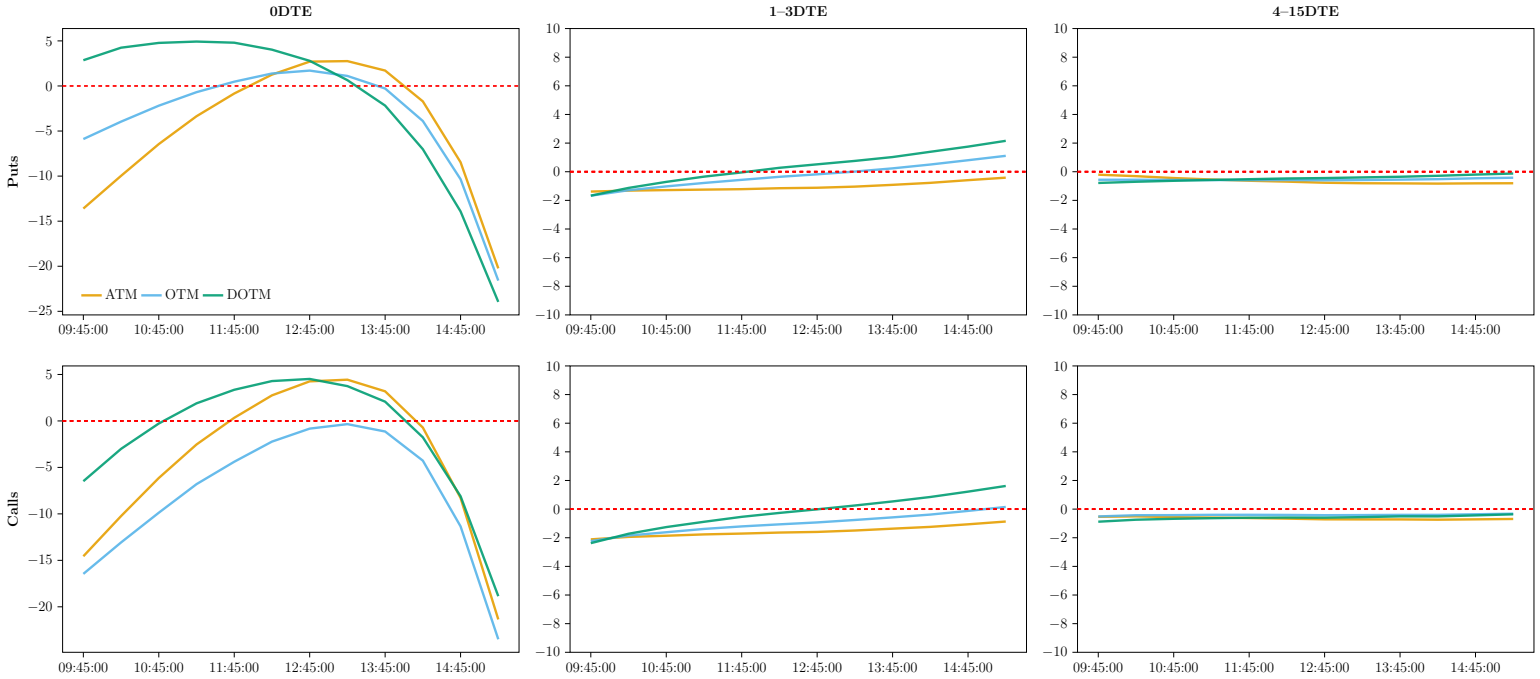


Figure 7. Cumulative Returns of Shorting Self-financing Six-factor-neutral Portfolio

This figure plots the time-series of the cumulative returns of shorting the self-financing six-factor-neutral portfolios—selling 0DTE options and hedging them with instrument option contracts—for positions held from 15:15 ET to 15:45 ET, using mid-quotes and bid-ask quotes, respectively. The hedging procedure is implemented for each individual 0DTE option contract, within each moneyness bucket (ATM, OTM, DOTM). The non-0DTE options are assigned to one of six buckets based on option maturities (1–3DTE, 4–15DTE) and moneyness (ATM, OTM, DOTM). The instrument contracts are selected, at each 30-minute timestamp, as the three contracts with the highest trading volume over the preceding 30 minutes in each of the six buckets (pooling puts and calls), yielding a total of 18 instrument contracts. The factor-neutral returns are averaged cross-sectionally at each timestamp. The green line represents the cumulative returns of shorting the six-factor-neutral strategy using mid-quotes of both the 0DTE options and the instrument options. The purple line represents the cumulative returns of shorting the six-factor-neutral strategy using bid and ask quotes, where the 0DTE options are shorted at the bid price, and the instrument options are bought (shorted) at the ask (bid) price. The gray line reports the cumulative return of a long position in ES futures initiated at the same intraday time and held for the same period as the option trade. The hedged portfolios are leveraged such that the unconditional (whole-sample) volatility of the hedged returns matches that of the corresponding ES futures returns: the factor-neutral return is multiplied by $\frac{\sigma_{ES}}{\sigma_{b,N}}$, where σ_{ES} is the unconditional volatility of ES futures returns, and $\sigma_{b,N}$ is the unconditional volatility of the risk-factor-neutral returns within each moneyness bucket b, computed separately for the mid-quote and bid-ask cases. This normalization uses full-sample volatilities and is therefore descriptive rather than implementable in real time. The cumulative returns are arithmetic running sums of the daily returns, rather than compounded portfolio wealth. The ordinate axis labels are in percent (multiplied by 100). The sample period is from 2016/09 to 2023/12.

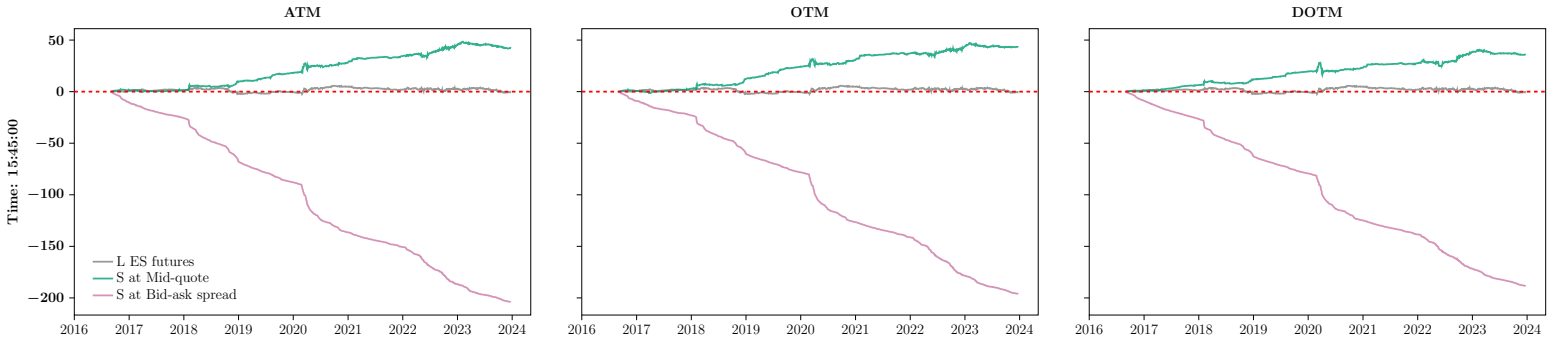


Table I
Summary Statistics of Intraday Delevered Excess Returns on S&P 500 Index Options

The table reports the summary statistics of intraday delevered excess returns on S&P 500 index options, before and after delta-hedging. The column labeled Moneyness indicates how options are grouped by their moneyness level. The column labeled DTE indicates how options are grouped by their remaining days to expiration. Further columns report, for unhedged and DLT-hedged returns respectively, the mean return (Mean), the standard deviation (SD), and the number of observations in each sub-sample (N). Means and standard deviations are reported in percent; accounting for 12 intraday periods per trading day and 252 trading days per year, means are annualized by the factor 12×252 and standard deviations by $\sqrt{12 \times 252}$. Significance is based on HAC standard errors for the mean of each series, with $\lceil T^{1/3} \rceil$ lags (Internet Appendix F). We delta-hedge the options with the S&P 500 E-Mini futures returns as in equation (2), and we estimate the option delta using our main specification of the model in Section 3. In the rare case an estimated delta has an incorrect sign (a positive delta for a put, or a negative delta for a call), we set it to zero. We measure moneyness as $\kappa_{t,\tau} = \frac{\ln(K/F_{t,\tau})}{\sqrt{\tau CIV_{t,\tau}}}$, and we supply all details about the calculations in Internet Appendix C. The option moneyness categories are: DOTM puts (calls) for $\kappa_{t,\tau} \leq -2.2$ ($\kappa_{t,\tau} > 1$), OTM puts (calls) for $-2.2 < \kappa_{t,\tau} \leq -1$ ($0.5 < \kappa_{t,\tau} \leq 1$) and ATM puts (calls) for $-1 < \kappa_{t,\tau} \leq 1$ ($-0.5 < \kappa_{t,\tau} \leq 0.5$). The sample period is from 2016/09 to 2023/12.

		Puts					Calls				
		Unhedged		DLT-hedged			Unhedged		DLT-hedged		
Moneyness	DTE	Mean	SD	Mean	SD	N	Mean	SD	Mean	SD	N
All	All	−4.71***	6.75	−1.57***	1.63	13,985,461	3.44***	5.38	−1.45***	1.89	6,798,057
	0	−9.01***	7.22	−6.45***	1.90	199,374	−4.09	5.88	−6.76***	1.93	126,628
	1–3	−6.31***	7.11	−2.10***	1.59	1,724,195	2.89**	5.78	−2.51***	1.81	910,685
	4–15	−4.41***	6.69	−1.41***	1.63	12,061,892	3.69***	5.30	−1.17***	1.91	5,760,744
ATM	All	−8.07***	9.23	−2.39***	2.05	7,265,721	5.17***	6.87	−0.91***	1.97	3,819,215
	0	−13.25**	9.83	−8.71***	2.37	104,201	−4.43	8.60	−8.54***	2.44	52,118
	1–3	−10.95***	9.73	−3.46***	1.97	893,277	4.88**	7.82	−2.41***	2.00	448,581
	4–15	−7.58***	9.14	−2.14***	2.06	6,268,243	5.36***	6.70	−0.59**	1.96	3,318,516
OTM	All	−1.52***	2.07	−0.90***	1.16	4,044,516	2.11***	2.81	−2.44***	2.05	1,959,969
	0	−5.42***	2.19	−4.80***	1.40	61,305	−3.94	3.97	−7.00***	2.02	26,255
	1–3	−1.99***	2.15	−0.94***	1.19	524,042	2.32*	3.35	−3.06***	1.98	234,071
	4–15	−1.38***	2.06	−0.83***	1.15	3,459,169	2.18***	2.70	−2.28***	2.06	1,699,643
DOTM	All	−0.42***	0.64	−0.33***	0.57	2,675,224	−0.51***	1.05	−1.58***	1.11	1,018,873
	0	−2.46***	0.68	−2.48***	0.61	33,868	−3.80***	1.45	−4.70***	1.04	48,255
	1–3	−0.15	0.73	−0.12	0.64	306,876	−0.46	1.22	−2.17***	1.14	228,033
	4–15	−0.43***	0.63	−0.32***	0.56	2,334,480	−0.31*	0.95	−1.19***	1.11	742,585

⋆ denotes statistically significant difference from zero at the 10% level, ⋆⋆ at the 5% level, and ⋆⋆⋆ at the 1% level.

Table II
Summary Statistics of Risk Factors under P and Q Measures

This table reports the sample averages of the risk factor realizations ($\overline{f}$), the sample averages of the factors' conditional risk-neutral expectations ($\overline{\mu}^Q$), and the risk premiums, i.e., differences between the two ($\widehat{\lambda} = \overline{f} - \overline{\mu}^Q$). The construction of the factors is described in Section 3. The conditional risk-neutral expectations, μ_t^Q, are estimated from cross-sectional regressions of option-specific conditional intercepts on estimated exposures at each 30-minute timestamp, jointly for puts and calls with type-, maturity-, and moneyness-specific bucket fixed effects, as specified in equation (12). Standard errors (in parentheses) account for serial dependence and for the sequential use of estimated inputs: heteroskedasticity- and autocorrelation-consistent (HAC) standard errors for $\overline{f}$; for $\overline{\mu}^Q$, the Fama–MacBeth time-series variance augmented with an adjustment for generated regressors, since the conditional risk-neutral expectations are estimated from estimated exposures; and for $\widehat{\lambda}$, the combination of the two, accounting for the covariance between the factor realizations and the estimated time series of Q-expectations. We provide the details in Internet Appendix F. The table reports these summary statistics for three sub-samples: the entire dataset (All), FOMC announcement days (FOMC), and the remaining days (Non-FOMC). All the values are reported in percent per annum (multiplied by 12×252); as indicated in the column header, the JSK entries are additionally multiplied by 100. The units are factor-specific: returns for DLT, squared returns for GAM, cubed returns for JSK, and changes in variance measures for dSV, dJV, and dTRM. The sample period is from 2016/09 to 2023/12.

Period	Qty	DLT	GAM	JSK×100	dSV	dJV	dTRM
All	$\overline{f}$	5.2	1.6***	0.05	5.7	−18.4**	7.6
		(4.4)	(0.1)	(0.1)	(7.5)	(9.2)	(5.9)
	$\overline{\mu}^Q$	0.01	1.4***	−0.2***	4.1	−3.5	−5.4
		(0.3)	(0.07)	(0.04)	(4.1)	(5.2)	(4.1)
	$\widehat{\lambda}$	5.2	0.2*	0.3**	1.5	−14.9	13.0*
		(4.5)	(0.10)	(0.1)	(8.5)	(10.1)	(6.9)
FOMC	$\overline{f}$	0.2	2.7***	-0.08	−54.7*	−0.7	45.2*
		(30.8)	(0.5)	(0.4)	(30.7)	(26.3)	(23.2)
	$\overline{\mu}^Q$	−0.5	2.0***	−0.6***	−12.0*	−5.4	0.5
		(0.5)	(0.2)	(0.1)	(6.8)	(9.3)	(7.8)
	$\widehat{\lambda}$	0.7	0.7*	0.6	−42.7	4.7	44.7*
		(30.8)	(0.4)	(0.4)	(30.6)	(27.7)	(23.3)
Non-FOMC	$\overline{f}$	5.4	1.5***	0.06	7.7	−19.0**	6.4
		(4.5)	(0.1)	(0.1)	(7.7)	(9.4)	(6.1)
	$\overline{\mu}^Q$	0.02	1.3***	−0.2***	4.7	−3.4	−5.6
		(0.3)	(0.07)	(0.04)	(4.1)	(5.1)	(4.1)
	$\widehat{\lambda}$	5.3	0.2*	0.3**	3.0	−15.6	12.0*
		(4.5)	(0.10)	(0.1)	(8.7)	(10.3)	(7.0)

$\star$ denotes statistically significant difference from zero at the 10% level, $\star\star$ at the 5% level, and $\star\star\star$ at the 1% level.

Table III
Model Explanatory Power for Delta-Hedged Delevered Option Returns

This table reports the incremental $R^2\%$ of the five risk factors for delta-hedged delevered excess option returns. For a given set of included factors, the $R^2\%$ is calculated as:

$$R^2\% = \left[1 - \frac{\sum_{i,t}\left(r^{dh}_{i,t+1} - \left(\widehat{c}_{i,t} + \sum_l \widehat{\beta}_{l,i,t} f_{l,t+1}\right)\right)^2}{\sum_{i,t}(r^{dh}_{i,t+1})^2}\right] \times 100,$$

where the sum runs over the included factors (DLT is never included, as the returns are delta-hedged), $\widehat{c}_{i,t}$ is an option's conditional intercept, and $\widehat{\beta}_{l,i,t}$ is the estimated option exposure. We give a breakdown of the results for puts and calls, respectively, and across option moneyness and maturity buckets. In each sub-panel, each column labeled with a factor name reports the increment in $R^2\%$ after adding that factor to a model that already includes the conditional intercept and the factors listed to its left; the increments therefore sum to the total. Because each specification is estimated on the full panel while the $R^2\%$ is evaluated within each subsample, increments can be negative. The column labeled "5F" reports the total $R^2\%$ of the five-factor model. We delta-hedge the options with the S&P 500 E-Mini futures returns as in equation (2), and we estimate the option delta using our main specification of the model in Section 3. In the rare case an estimated delta has an incorrect sign (a positive delta for a put, or a negative delta for a call), we set it to zero. We measure moneyness as $\kappa_{t,\tau} = \frac{\ln(K/F_{t,\tau})}{\sqrt{\tau}CIV_{t,\tau}}$, and we supply all details about the calculations in Internet Appendix C. The option moneyness categories are: DOTM puts (calls) for $\kappa_{t,\tau} \leq -2.2$ ($\kappa_{t,\tau} > 1$), OTM puts (calls) for $-2.2 < \kappa_{t,\tau} \leq -1$ ($0.5 < \kappa_{t,\tau} \leq 1$) and ATM puts (calls) for $-1 < \kappa_{t,\tau} \leq 1$ ($-0.5 < \kappa_{t,\tau} \leq 0.5$). The sample period is from 2016/09 to 2023/12.

		Puts						Calls					
Moneyness	DTE	GAM	JSK	dSV	dJV	dTRM	5F	GAM	JSK	dSV	dJV	dTRM	5F
All	All	10.4	19.2	28.5	20.1	6.0	84.3	11.7	24.7	30.6	12.0	6.0	85.0
	0	58.8	6.4	3.7	2.4	0.5	71.7	59.0	7.3	2.3	1.1	0.5	70.3
	1–3	17.5	13.6	28.6	18.8	1.1	79.4	17.4	20.1	31.5	10.7	1.1	80.8
	4–15	8.3	20.3	29.1	20.7	6.8	85.2	9.8	25.8	31.1	12.4	6.9	86.0
ATM	All	10.8	19.9	29.4	17.0	6.1	83.1	13.4	17.8	30.8	15.0	6.7	83.7
	0	61.8	4.9	3.6	2.2	0.5	73.0	66.7	3.8	1.7	0.9	0.5	73.5
	1–3	18.4	13.8	30.5	14.1	1.2	77.9	21.0	12.4	30.8	13.9	1.4	79.4
	4–15	8.6	21.0	29.8	17.7	6.9	84.0	11.0	18.9	31.5	15.5	7.7	84.6
OTM	All	9.3	15.7	25.0	35.5	5.5	90.9	10.4	35.0	30.5	7.2	5.2	88.2
	0	47.6	12.7	4.3	3.0	1.1	68.6	53.3	11.4	3.7	1.8	0.8	70.9
	1–3	15.3	12.2	21.3	36.8	0.8	86.4	13.2	30.4	33.9	5.9	0.8	84.2
	4–15	7.3	16.4	26.1	36.1	6.3	92.2	9.4	35.9	30.4	7.4	5.8	89.0
DOTM	All	2.7	18.7	21.3	34.8	5.2	82.7	0.6	39.4	29.3	7.4	2.9	79.7
	0	24.6	13.2	3.8	1.6	−1.4	41.9	25.8	20.1	3.2	0.9	0.4	50.4
	1–3	5.0	15.3	18.3	41.6	−0.4	79.8	8.4	35.7	28.5	5.8	0.4	78.9
	4–15	1.9	19.4	22.1	34.2	6.4	84.0	−3.4	41.7	31.1	8.4	3.9	81.7

Table IV
Summary Statistics of Systematic Pricing Errors

The table collects the summary statistics of systematic pricing errors estimated at the option maturity and moneyness bucket level, $\widehat{\alpha}_{b(i),t}$. The alphas are estimated as bucket-level fixed effects in cross-sectional regressions of the conditional intercepts on option exposures (as in equation (14)), at each 30-minute timestamp. Within each option type, the options are assigned to one of nine buckets based on maturity (0DTE, 1–3DTE, 4–15DTE) and moneyness (DOTM, OTM, ATM), for 18 type-specific buckets in total. The reported statistics include the mean ($\overline{\widehat{\alpha}}$), the mean of the absolute values ($\overline{|\widehat{\alpha}|}$), the standard deviation (SD), the first-order autocorrelation of the series of consecutive 30-minute observations ($\text{AC}^{30\text{min}}$), the 1-day autocorrelation ($\text{AC}^{1\text{day}}$), computed from the series of alphas estimated on successive trading days from returns between 15:15 ET and 15:45 ET, and the number of observations (N). Row labels $\alpha_{m,g}$ identify the maturity bucket, m (DTE), and the moneyness group, g (A, O, and D for ATM, OTM, and DOTM, respectively). Each alpha estimate is expressed in percent per annum (multiplied by 12×252); the reported means and standard deviations are the time-series average and dispersion of these annualized estimates. Significance is based on HAC standard errors for the mean of each bucket's alpha time series, with $\lceil T^{1/3} \rceil$ lags (Internet Appendix F). The sample period is from 2016/09 to 2023/12.

	Puts						Calls					
Factor	$\overline{\widehat{\alpha}}$	$\overline{\lvert\widehat{\alpha}\rvert}$	SD	$\text{AC}^{30\text{min}}$	$\text{AC}^{1\text{day}}$	N	$\overline{\widehat{\alpha}}$	$\overline{\lvert\widehat{\alpha}\rvert}$	SD	$\text{AC}^{30\text{min}}$	$\text{AC}^{1\text{day}}$	N
$\alpha_{0,A}$	$-4.63^{\star\star\star}$	10.85	16.05	0.92	0.32	15,493	$-3.94^{\star\star\star}$	12.06	18.28	0.93	0.29	14,990
$\alpha_{0,O}$	$-3.52^{\star\star\star}$	7.23	11.66	0.91	0.48	15,244	$-7.73^{\star\star\star}$	11.25	15.92	0.93	0.42	13,474
$\alpha_{0,D}$	$-1.53^{\star\star\star}$	8.34	11.83	0.93	0.53	14,698	$-1.51^{\star\star\star}$	10.22	16.02	0.94	0.68	15,550
$\alpha_{1\sim3,A}$	$-1.04^{\star\star\star}$	2.84	4.45	0.98	0.66	21,941	$-1.56^{\star\star\star}$	3.43	5.91	0.99	0.76	21,941
$\alpha_{1\sim3,O}$	$-0.27^{\star\star\star}$	2.00	3.28	0.97	0.58	21,941	$-1.00^{\star\star\star}$	2.76	5.42	0.99	0.80	21,941
$\alpha_{1\sim3,D}$	$0.33^{\star\star\star}$	2.36	4.08	0.98	0.61	21,655	-0.22	3.48	6.16	0.99	0.85	21,907
$\alpha_{4\sim15,A}$	$-0.64^{\star\star\star}$	2.22	3.54	0.97	0.66	21,953	$-0.65^{\star\star\star}$	2.37	4.70	0.99	0.80	21,953
$\alpha_{4\sim15,O}$	$-0.55^{\star\star\star}$	1.83	2.92	0.98	0.64	21,953	$-0.41^{\star\star\star}$	2.38	4.53	0.99	0.82	21,953
$\alpha_{4\sim15,D}$	$-0.46^{\star\star\star}$	1.72	2.99	0.99	0.71	21,953	$-0.59^{\star\star\star}$	2.71	4.36	0.99	0.88	21,772

$\star$ denotes statistically significant difference from zero at the 10% level, $\star\star$ at the 5% level, and $\star\star\star$ at the 1% level.

Table V
Pricing Errors and Intermediary Frictions

The table reports the coefficients and R^2 of the regressions of ATM option pricing errors measured from returns between 15:15 ET and 15:45 ET on variables related to frictions faced by financial intermediaries and on controls. The top panel reports the results for alphas, $\widehat{\alpha}_{b(i),t}$, while the bottom panel reports the results for absolute alphas, $|\widehat{\alpha}_{b(i),t}|$. The alphas are estimated as bucket-level fixed effects in cross-sectional regressions of the conditional intercepts on option exposures (as in equation (14)), at each 30-minute timestamp. *ESF* is the spread between the forward interest rate implied by the S&P 500 E-Mini futures and the three-month OIS rate, as in Siriwardane, Sunderam, and Wallen (2025) and Internet Appendix G, estimated using futures prices at the futures closing time, 16:15 ET, of the previous day. *HKM* is the level of the aggregate market equity ratio of primary dealers at the end of the previous day as defined by He, Kelly, and Manela (2017). ε_{HKM}, proposed therein, is the AR(1) innovation to the capital ratio on the current day, scaled by the lagged ratio. *CIV* is the corridor implied variance at 15:15 ET on the current day, computed from 0DTE options for the 0DTE regressions and averaged across the maturities in the bucket for the 1–3DTE and 4–15DTE regressions. *OI* is the open interest of all the puts (calls) within each bucket at the end of the previous day. *SP500* is the intraday return of the S&P 500 ES futures from 9:45 ET to 15:15 ET on the current day. For each combination of option maturity and type, the table reports the estimated coefficients, their Newey–West standard errors with $\lceil T^{1/3} \rceil$ lags (in parentheses), robust to the serial dependence of the alphas, and the adjusted R^2 (in %, in the row labeled R^2) in the first column, and the adjusted $R^2_{\text{univ.}}$ (in %) from a univariate regression on the given variable in the second column. Within each maturity–type regression, the dependent variable and the regressors are normalized by their respective sample standard deviations. The sample period is from 2016/09 to 2023/12.

	$\widehat{\alpha}_{b(i),t}$											
	0DTE				1–3DTE				4–15DTE			
	Puts		Calls		Puts		Calls		Puts		Calls	
	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$
ESF	−0.03	0.32	−0.04	0.51	−0.03	0.74	−0.11***	2.52	0.03	0.40	−0.06*	2.36
	(0.05)		(0.04)		(0.03)		(0.03)		(0.03)		(0.03)	
HKM	−0.05	0.33	−0.06	0.20	0.13***	11.06	0.22***	19.58	0.04	10.59	0.12***	19.58
	(0.05)		(0.04)		(0.04)		(0.03)		(0.04)		(0.04)	
ε_{HKM}	0.05	1.81	0.06*	0.91	0.05*	0.85	0.01	−0.05	0.02	0.91	−0.07*	0.03
	(0.03)		(0.03)		(0.03)		(0.02)		(0.02)		(0.04)	
CIV	−0.53***	34.31	−0.62***	39.44	−0.64***	52.00	−0.71***	62.43	−0.67***	51.34	−0.75***	65.06
	(0.10)		(0.08)		(0.11)		(0.05)		(0.07)		(0.04)	
OI	−0.13***	14.58	−0.08***	1.40	−0.12***	12.55	−0.05***	0.06	−0.09**	13.21	−0.02	5.13
	(0.03)		(0.02)		(0.03)		(0.02)		(0.04)		(0.02)	
SP500	0.06**	3.39	0.06**	2.50	0.10***	1.41	0.06**	0.15	0.10***	1.22	0.06	−0.05
	(0.03)		(0.03)		(0.02)		(0.03)		(0.03)		(0.05)	
Const.	−0.56*		−0.43*		−0.43		−0.95***		−0.01		−0.28	
	(0.31)		(0.26)		(0.28)		(0.20)		(0.28)		(0.24)	
R^2	37.34		41.10		56.50		67.94		53.60		66.82	

	$\lvert\widehat{\alpha}_{b(i),t}\rvert$											
	0DTE				1–3DTE				4–15DTE			
	Puts		Calls		Puts		Calls		Puts		Calls	
	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$
ESF	0.06	0.87	0.07**	1.12	0.03	0.80	0.09***	2.39	−0.02	0.57	0.04	2.41
	(0.04)		(0.04)		(0.04)		(0.03)		(0.03)		(0.04)	
HKM	0.02	0.98	0.03	0.74	0.10***	1.73	0.02	5.34	0.13***	3.85	0.11**	5.85
	(0.05)		(0.04)		(0.04)		(0.03)		(0.03)		(0.04)	
ε_{HKM}	−0.04	1.80	−0.05	0.87	−0.04	0.33	0.02	−0.04	−0.01	0.57	0.08	0.03
	(0.04)		(0.03)		(0.03)		(0.02)		(0.02)		(0.05)	
CIV	0.57***	40.84	0.66***	46.60	0.75***	51.24	0.80***	64.55	0.77***	49.72	0.84***	61.88
	(0.11)		(0.09)		(0.09)		(0.03)		(0.06)		(0.05)	
OI	0.14***	16.75	0.08***	1.48	0.00	6.06	0.00	−0.05	−0.01	9.61	0.05*	1.63
	(0.03)		(0.02)		(0.03)		(0.02)		(0.04)		(0.02)	
SP500	−0.07**	3.90	−0.07**	2.87	−0.06***	0.64	−0.04	0.05	−0.10***	0.84	−0.06	−0.04
	(0.03)		(0.03)		(0.02)		(0.03)		(0.03)		(0.05)	
Const.	0.88***		0.73***		−0.31		0.01		−0.60**		−0.82***	
	(0.32)		(0.28)		(0.25)		(0.19)		(0.26)		(0.27)	
R^2	44.13		48.36		53.05		65.48		52.19		63.94	

Table VI
Summary Statistics of Self-financing Six-factor-neutral Portfolio Returns

This table reports the summary statistics of the risk-factor-neutral returns of 0DTE options hedged by instrument option contracts using a self-financing strategy. The "Object" column indicates the return type. "T" (for target) denotes delevered DLT-hedged returns of 0DTE options without additional hedging; "H" denotes the return on the hedge portfolio of instrument option contracts, $\boldsymbol{w}'_{i,t}\boldsymbol{r}^H_{i,t+1}$, with the weights of equation (16); and "N" denotes the six-factor-neutral portfolio returns, the sum of the two. The hedging procedure described in Section 5.6 is implemented for each individual 0DTE put and call contract, within each moneyness bucket (ATM, OTM, DOTM), at each 30-minute timestamp. The factor-neutral returns of option i are averaged cross-sectionally and then over time. For example, the reported mean is calculated as $\frac{1}{T}\sum_{t=1}^{T}\frac{1}{n_t}\sum_{i=1}^{n_t} r^N_{i,t}$, where n_t is the number of contracts at timestamp t. All the statistics are annualized under a once-per-day 30-minute position: means and quartiles are multiplied by 252×100, standard deviations by $\sqrt{252}\times 100$, and Sharpe ratios by $\sqrt{252}$; the scaled quartiles are per-position statistics, not quantiles of annual returns. Significance is based on HAC standard errors for the mean of each return series, with $\lceil T^{1/3}\rceil$ lags (Internet Appendix F). The sample period is from 2016/09 to 2023/12.

		Puts						Calls					
Mny	Object	Mean	SD	SR	25th	50th	75th	Mean	SD	SR	25th	50th	75th
ATM	T	−0.61***	0.49	−1.25	−3.56	−1.47	0.80	−0.54***	0.54	−1.00	−3.89	−1.60	1.06
	H	0.48***	0.70	0.68	−1.98	1.28	4.39	0.46***	0.79	0.58	−2.31	1.40	4.94
	N	−0.13**	0.50	−0.27	−2.43	−0.07	2.30	−0.08	0.55	−0.15	−2.64	0.03	2.69
OTM	T	−0.35***	0.32	−1.11	−2.04	−0.94	0.22	−0.49***	0.44	−1.11	−3.21	−1.39	0.83
	H	0.21***	0.44	0.47	−1.36	0.75	2.80	0.38***	0.59	0.65	−1.70	1.06	3.53
	N	−0.14***	0.33	−0.43	−1.86	−0.05	1.74	−0.11*	0.44	−0.24	−1.96	−0.09	1.86
DOTM	T	−0.19***	0.15	−1.25	−0.90	−0.41	0.06	−0.30***	0.22	−1.35	−1.32	−0.56	0.15
	H	0.09***	0.23	0.40	−0.76	0.29	1.31	0.23***	0.35	0.66	−0.95	0.51	1.89
	N	−0.09***	0.20	−0.45	−1.07	−0.03	1.03	−0.06*	0.28	−0.23	−1.16	0.00	1.16

Internet Appendix

for

"The Factor Structure of 0DTE Option Returns"

Appendix A. Data

Our dataset comprises high-frequency records of quotes of S&P 500 Index options from September 2016 to December 2023. We obtain this data from the Chicago Board Options Exchange (CBOE)'s Option Intervals dataset, which includes option quotes aggregated to a 1-minute frequency alongside synchronized intraday records of the underlying S&P 500 index. We use the midpoint of the high-frequency quotes as the option price and then compute returns over thirty-minute intervals. Since the S&P 500 index is costly to replicate at a high frequency, we use the E-mini S&P 500 (ES) futures as the tradable hedging asset; these are the most actively traded futures on the S&P 500 index. We obtain the millisecond-frequency quotes for ES futures from Tickdata Inc. Finally, we obtain risk-free rates from the zero-yield curve of OptionMetrics. We assume that the risk-free rates remain constant across all timestamps within each day.

To study the performance of investments in short-term, continuously hedged options, we construct an investment period which starts 15 business days before expiration and ends at expiration. We avoid irregular option quotes and anomalies in bid-ask spreads shortly after the market opens by setting 9:45 a.m. ET as the time of the first trade, and 3:45 p.m. ET as the time of the last trade to demarcate the intraday and overnight periods. To reduce the impact of microstructure noise, we subsample these datasets at a 30-minute frequency.

Throughout this period, whenever an option contract satisfies the selection conditions stated below, it is considered to be held by an investor for the next 30-minute or overnight period. Overnight periods enter this tracking only to determine position continuity, entry, exit, and total losses; the returns underlying all reported results are the intraday 30-minute returns between 9:45 and 15:45 ET. We only consider PM-settled options. For each option position with 15 or fewer DTE the first 30-minute or overnight return included in our sample is the earliest interval in which the option has positive trading volume and a positive bid price at the outset, and is neither in-the-money (ITM)[1] nor extremely deep out-of-the-money (DOTM)[2] indicating real trading activity. Subsequently, we do not require the option to have

[1]The trading of short-term options is skewed toward OTM options (Andersen, Fusari, and Todorov, 2017). We keep options with $\kappa_{t,\tau} > -0.5$ for calls and $\kappa_{t,\tau} \leq 1$ for puts where $\kappa_{t,\tau}$ is the relative strike with maturity τ at time t as defined in Section C.

[2]We exclude the options in the 1% extreme left and right tails, due to their illiquidity and lack of meaningful informational content. The tail ratio and strike boundary are defined in Section B.

positive trading volume at the start of each 30-minute or overnight interval. Once the option becomes ITM or extreme DOTM, the investment position is assumed to be closed and the option is excluded from the sample. However, it may re-enter the sample if it satisfies the prior conditions again.

Finally, we track an option's value until expiration in the following manner. If the option has a positive bid price at the end of each trading period, we keep all 30-minute and overnight returns satisfying prior criteria, and calculate the final return at expiration using the (potentially zero) settlement payoff of the option. If the option's bid price reaches zero prior to expiration, we apply the following procedure: we track all trading periods where the start-of-period bid price is positive and the end-of-period bid price is zero; in the last such period before expiration, we consider the option to become valueless, i.e., we consider the investor writes the option off and realizes a -100% raw return on the option; in prior such periods, we use half of the ask price to calculate returns. Quoted bids are read with a 5-minute tolerance: at each 30-minute mark, we use the last positive bid observed in the preceding 5 minutes and record a zero quote if none exists; ES futures prices are matched at the selected timestamps.

Appendix B. Variance Proxies and Factor Construction

The implementation of our model requires proxies of market variance at all maturities. This proxy has two uses: first, it serves as a normalization factor for option moneyness (in Section C); second, we use its high-frequency evolution to decompose variance dynamics into three orthogonal components. Whereas the the VIX index calculated at selected maturities is an obvious candidate, Andersen, Bondarenko, and Gonzalez-Perez (2015) highlight sudden fluctuations in the high-frequency VIX series due to the cut-off rule for including options in the VIX calculation used by the CBOE. Following their work, we employ the square of the model-free corridor implied variance (CIV) index as the expected market variance measure. For option time to expiration τ, CIV is defined as

$$CIV = 2e^{r\tau}\left[\int_{\kappa_L}^{F_0} \frac{P_{t,\tau}(K)}{K^2}dK + \int_{F_0}^{\kappa_H} \frac{C_{t,\tau}(K)}{K^2}dK\right] \times \frac{252}{\tau} \tag{B1}$$

where $P_{t,\tau}(K)$ and $C_{t,\tau}(K)$ are put and call prices with strike K and time to expiration τ at time t respectively, and to compute it we must define the boundary strikes κ_L and κ_H.

To compute the CIV at each timestamp t, we first select all options with maturity τ. Then, we select all OTM options with positive bid quotes.[3] We determine the strike truncation

[3]We define OTM as $\frac{K}{F_t} \geq 1$ for calls and $\frac{K}{F_t} \leq 1$ for puts where K represents the strike and F_t is the forward price at time t.

level with the use of the following ratio statistic $R_{t,\tau}(K)$:

$$R_{t,\tau}(K) := \frac{P_{t,\tau}(K)}{P_{t,\tau}(K) + C_{t,\tau}(K)} \tag{B2}$$

which naturally runs from zero to one across the strike range for a fixed time to expiration. To maximize data retention, we set the constant 1% cut-off level; this ensures valid quotes spanning the corridor in each high-frequency interval. Then we set the lower integration boundary to $\kappa_L = R_{t,\tau}(K)^{-1}(0.01)$ and the upper boundary to $\kappa_H = R_{t,\tau}(K)^{-1}(0.99)$, which also serves as our criterion for identifying the extreme DOTM options, as described in prior sample selection conditions.

We then construct the variance dynamics factors of Section 3.2. We estimate spot variance (SV) following Cboe's methodology for the calculation of the SPOTVOL index, based on the nonparametric spot variance estimator proposed by Todorov (2019).[4] Cboe calculates initial estimates from options with the two shortest maturities exceeding 2DTE (labeled 2^+) and averages them to obtain the final value, SV_{2^+}; we define dSV := $SV_{2^+,t+1} - SV_{2^+,t}$. For the short-maturity risk-neutral variance, we repeat the procedure with the CIV estimator and denote the averaged estimator CIV_{2^+}; the dJV factor is then defined as dJV := $(CIV_{2^+,t+1} - SV_{2^+,t+1}) - (CIV_{2^+,t} - SV_{2^+,t})$. Finally, for the long-maturity risk-neutral variance, we use the CIV estimator calculated from options with the longest available maturity (14–15DTE), denoted CIV_{14^+}; the dTRM factor is defined as dTRM := $(CIV_{14^+,t+1} - CIV_{2^+,t+1}) - (CIV_{14^+,t} - CIV_{2^+,t})$.

Appendix C. Option Stochastic Maturity and Moneyness

These two quantities are key inputs to our modeling of option exposures. With our focus on intraday option return dynamics, we carefully describe their calculation in this section.

First, in the calculation of option maturity, we follow Bondarenko (2026) to account for the intraday variance profile. General linear maturity computation implicitly assumes that each minute contributes equally to the total variance an option is exposed to. This assumption is at odds with the observed intraday variance pattern. Figure IA.I plots average realized volatility by minute-of-day for E-mini S&P 500 futures over Monday to Thursday, spanning the full 24-hour trading cycle from the prior 18:00 ET, Global Trading Hours (GTH) open time, through the current day's 18:00 ET, post-market close. The volatility is low during GTH with only modest increases near the 08:30 ET U.S. macro release window, then jumps by roughly an order of magnitude at the open time of Regular Trading Hours (RTH) at 09:30 ET, traces the familiar U-shape, peaks again into close time at the 16:00 ET, and collapses near zero during the CME settlement halt (HALT, 16:15–16:30 ET) and technical break

[4] https://www.cboe.com/us/indices/dashboard/spotvol/

(BRK, 17:00–18:00 ET). For USM options, the assumption of linear maturity is problematic, especially for 0DTE options. Bondarenko (2026) shows that measuring maturity in calendar time rather than in units of variance accumulation accounts for a substantial portion of the apparent irregularity in the ultra-short implied volatility term structure.

To solve this problem, we adopt the stochastic clock of Bondarenko (2026), re-scaling each interval of clock time by its realized variance contribution, so that one unit of stochastic time corresponds to the average pricing-relevant variation in one RTH minute. We assign every minute of the week to one of six trading regimes on the CME E-mini S&P 500 futures market: GTH (18:00 to 09:30 ET), RTH (09:30 to 16:00 ET), Post-market (Post, 16:00–16:15 and 16:30–17:00 ET), HALT, the BRK (Monday to Thursday), and weekends and holidays (HOL, weekends start from 17:00 ET on Friday). We define minute-level weight $\lambda(m)$ that measures how much pricing time minute m of each day represents relative to one average RTH minute:

$$\lambda(m) = \frac{\bar{\sigma}^2_m}{\bar{\sigma}^2_{\text{RTH}}}, \qquad \bar{\sigma}^2_{\text{RTH}} = \frac{1}{390} \sum_{m' \in \text{RTH}} \bar{\sigma}^2_{m'}, \tag{C3}$$

where $\bar{\sigma}^2_m$ is the sample average of 1-minute realized variance across days at minute-of-day m. By construction, the mean weight across the 390 RTH minutes equals one. The cumulative stochastic time from a reference point t_0 to time t is

$$C(t) = \sum_{s=t_0}^{t} \lambda\big(m(s)\big), \tag{C4}$$

where $m(s)$ is the minute-of-day at calendar time s. $C(t)$ counts RTH-equivalent minutes: it advances quickly when prices are active and almost stalls when they are not. Figure IA.II plots the cumulative stochastic clock $C(t)$ alongside the calendar clock and an RTH clock, measuring the passage of time over a single Monday-through-Thursday trading day (18:00 the prior evening to 18:00 ET). The calendar clock rises linearly. The RTH clock, which counts only the 390 minutes between 09:30 and 16:00, is flat outside the RTH and rises linearly during the regular session. Neither captures the within-day variation in price activity. The stochastic clock advances slowly during GTH, accelerates during the open and close of regular trading period, and is nearly flat during the HALT and BRK regime.

To compute the option maturity using the stochastic clock, we need to convert the calendar year to stochastic time as the annualization factor. A calendar year contains 252 trading days, each with a GTH, RTH, HALT, and Post session. Of these trading days, 191 days have a standard 60-minute BRK period, and 61 days include pre-holiday gaps (17:00–18:00 ET of 52 Fridays and 9 pre-holiday days) classified as HOL. Adding the remaining 113 full 24-hour HOL, the total λ-weighted minutes across all regimes yields an annualization constant of $\Omega \approx 181{,}750$ RTH-equivalent minutes per year. We define the stochastic maturity of an

option with expiration T, observed at time t, as:

$$\tau = \frac{C(T)}{\Omega} = \frac{\sum_{s=t}^{T} \lambda\big(m(s)\big)}{\Omega}. \tag{C5}$$

To account for the fact that investors trade narrower strike ranges as options approach expiration, we adjust the logarithmic moneyness by incorporating volatility and define moneyness as:

$$\kappa_{t,\tau} = \frac{\ln\left(K/F_{t,\tau}\right)}{\sqrt{\tau\, CIV_{\,t,\tau}}} \tag{C6}$$

where CIV is the maturity-specific aggregate market variance proxy which captures the market conditions, described in Section B, $F_{t,\tau}$ is the forward price of the index computed from put-call parity using ATM option pairs, $F_{t,\tau} = K_{ATM} + e^{r\tau}(C_{t,\tau}(K_{ATM}) - P_{t,\tau}(K_{ATM}))$, at time t with time to expiration τ.

Figure IA.III plots the distribution of the moneyness for puts and calls across by maturity bucket (0DTE, 1-3 DTE, and 4-15 DTE). Considering the narrower moneyness ranges for calls, we divide moneyness into finer intervals for calls in our analysis. We define at-the-money (ATM) puts (calls) as $-1 < \kappa_{t,\tau} \leq 1$ ($-0.5 < \kappa_{t,\tau} \leq 0.5$), OTM puts (calls) as $-2.2 < \kappa_{t,\tau} \leq -1$ ($0.5 < \kappa_{t,\tau} \leq 1$), and DOTM puts (calls) as $\kappa_{t,\tau} \leq -2.2$ ($\kappa_{t,\tau} > 1$). The distribution shapes are very similar across maturities and stable across the intraday timestamps: for puts, density declines smoothly toward DOTM moneyness; for calls, density presents a steep drop beyond ATM and a short right tail.

Appendix D. Structural Model Derivations

This appendix develops the three-factor structural model supporting Section 3 of the main text: the affine mapping between model-free implied variance and the latent states, the structural content of the dJV and dTRM proxies, and the decomposition of the delevered option return that underlies the composite sensitivities $\widetilde{O}^*$.

Model. Under the risk-neutral measure, the dynamics of the S&P 500 index level S_t, the variance processes, and jumps are given by (restating equation (3) of the main text):

$$\begin{aligned}
\frac{\mathrm{d}S_t}{S_{t^-}} &= (r - \delta - \nu_t)\,\mathrm{d}t + \sqrt{(1-\rho^2)\,v_t}\,\mathrm{d}W_{St} + \rho\sqrt{v_t}\,\mathrm{d}W_{vt} + \left(e^J - 1\right)\mathrm{d}N_t \\
\mathrm{d}v_t &= k_v\,(m_t - v_t)\,\mathrm{d}t + \sigma_v v_t^{\delta_v}\,\mathrm{d}W_{vt} \\
\mathrm{d}m_t &= k_m\,(\theta_m - m_t)\,\mathrm{d}t + \sigma_m m_t^{\delta_m}\,\mathrm{d}W_{mt} \\
\mathrm{d}u_t &= k_u\,(\theta_u - u_t)\,\mathrm{d}t + \sigma_u u_t^{\delta_u}\,\mathrm{d}W_{ut}\,,
\end{aligned} \tag{D7}$$

with $J_t \sim N(\mu_J, \sigma_J^2)$, jump intensity $\lambda_t = \lambda_0 + \lambda_1 v_t + \lambda_2 u_t$, and jump compensator $\nu_t = \lambda_t(e^{\mu_J + \sigma_J^2/2} - 1)$.

The affine structure of model-free implied variance. In this model, the model-free implied variance (Jiang and Tian, 2005; Andersen et al., 2015) decomposes into integrated variance and a jump correction:

$$MFIV_{t,\tau} = -\frac{2}{\tau}\mathbb{E}_t^{\mathbb{Q}}\left[\log F_{t+\tau}/F_t\right] = \frac{1}{\tau}\,\mathbb{E}_t^{\mathbb{Q}}\left[\int_t^{t+\tau} v_s\,\mathrm{d}s\right] - \frac{2\left(\mu_J - (e^{\mu_J+\sigma_J^2/2}-1)\right)}{\tau}\,\mathbb{E}_t^{\mathbb{Q}}\left[\int_t^{t+\tau}\lambda_s\,\mathrm{d}s\right]. \tag{D8}$$

Because the drifts of v_t, m_t, and u_t are affine in the state, the conditional expectations of the integrals in equation (D8) are affine functions of (v_t, m_t, u_t), and so is λ_t; hence

$$MFIV_{t,\tau} = c_0 + c_v(\tau)v_t + c_m(\tau)m_t + c_u(\tau)u_t\,, \tag{D9}$$

where the loadings $c_x(\tau)$ depend only on time to maturity and the model parameters.

The dJV **proxy.** A Taylor expansion of $MFIV_{t,\tau} - v_t$ around $\tau = 0$ gives

$$MFIV_{t,\tau} - v_t = -2\left(\mu_J - (e^{\mu_J+\sigma_J^2/2}-1)\right)\lambda_t + O(\tau)\,,$$

and expanding the leading coefficient for small μ_J approximates it by the variance of jumps, ultimately giving

$$MFIV_{t,\tau} - v_t = \sigma_J^2\lambda_t + O(\mu_J) + O(\tau)\,.$$

At leading order, the gap between short-maturity implied variance and spot variance is thus proportional to the jump intensity, with the jump variance as the factor of proportionality; its change, dJV, captures the dynamics of expected jump variance.

The dTRM **proxy.** Define $\zeta_x := c_x(\tau_L) - c_x(\tau_S)$ for $x \in \{v, m, u\}$, the difference between the implied-variance loadings in equation (D9) at the long and short maturities. (We write ζ_x rather than δ_x to avoid a collision with the diffusion elasticities δ_v, δ_m, and δ_u in equation (D7).) Substituting equation (D9) into the definition of dTRM and using $\lambda_t = \lambda_0 + \lambda_1 v_t + \lambda_2 u_t$ to eliminate u_t, we can express dTRM in terms of $\mathrm{d}v_t$, $\mathrm{d}m_t$, and $\mathrm{d}\lambda_t$ as

$$\mathrm{dTRM}_t = \zeta_m\,\mathrm{d}m_t + \left(\zeta_v - \frac{\lambda_1}{\lambda_2}\,\zeta_u\right)\mathrm{d}v_t + \frac{\zeta_u}{\lambda_2}\,\mathrm{d}\lambda_t\,. \tag{D10}$$

When orthogonalized to dSV and dJV, which span $\mathrm{d}v_t$ and $\mathrm{d}\lambda_t$, dTRM is therefore identified by shocks to the long-term variance process m_t.

Dynamics of the delevered option return. In this setting, the dynamics of the delevered option return, $\mathrm{d}O/S_{t^-}$, can be stated in terms of the underlying return, the three variance span proxies, and underlying jumps:

$$\begin{aligned}\frac{\mathrm{d}O}{S} - r\frac{O_{t^-}}{S}\mathrm{d}t = \frac{\Theta_{t^-}}{S}\,\mathrm{d}t + O_{S,t^-}\frac{\mathrm{d}S_t}{S_{t^-}} + \frac{1}{S_{t^-}}\left(O_{v,t^-} - \frac{\lambda_1}{\lambda_2}O_{u,t^-} - \frac{\zeta_v - \frac{\lambda_1}{\lambda_2}\zeta_u}{\zeta_m}\,O_{m,t^-}\right)\mathrm{dSV} \\ + \frac{1}{\sigma_J^2\lambda_2 S_{t^-}}\left(O_{u,t^-} - \frac{\zeta_u}{\zeta_m}\,O_{m,t^-}\right)\,\mathrm{dJV} + \frac{O_{m,t^-}}{S_{t^-}\,\zeta_m}\,\mathrm{dTRM} \\ + \frac{1}{S_{t^-}}\left[O(S_{t^-}e^J) - O(S_{t^-}) - O_{S,t^-}\,S_{t^-}(e^J-1)\right]\mathrm{d}N_t\,, \end{aligned} \tag{D11}$$

where Θ_{t^-} collects all terms of order $\mathrm{d}t$ (the option time decay and Itô terms). Matching coefficients with the corresponding display in Section 3.1.2 of the main text identifies the composite sensitivities:

$$\widetilde{O}^v_{t^-} = O_{v,t^-} - \frac{\lambda_1}{\lambda_2}O_{u,t^-} - \frac{\zeta_v - \frac{\lambda_1}{\lambda_2}\zeta_u}{\zeta_m}\,O_{m,t^-}\,, \qquad \widetilde{O}^u_{t^-} = \frac{1}{\sigma_J^2\lambda_2}\left(O_{u,t^-} - \frac{\zeta_u}{\zeta_m}\,O_{m,t^-}\right)\,, \tag{D12}$$

$$\widetilde{O}^m_{t^-} = \frac{O_{m,t^-}}{\zeta_m}\,, \qquad \widetilde{O}^N_{t^-} = O(S_{t^-}e^J) - O(S_{t^-}) - O_{S,t^-}\,S_{t^-}(e^J-1)\,.$$

Each $\widetilde{O}^x$, $x \in \{v, u, m\}$, is thus not the bare partial derivative O_x, but it is dominated by it: the remaining terms are scalings and cross-corrections in the other partials.

The option exposures have clear interpretations. The exposure to the underlying return is the option's delta. The exposures to the variance span proxies are best read starting from dTRM: this is the only observable driven by m (central tendency) shocks, and the associated exposure measures the option's central-tendency sensitivity per unit of term-structure slope movement. The exposure to dJV is the option's jump-intensity sensitivity, net of the component of the factor m already captured by dTRM. Finally, the exposure to dSV is the pure spot-variance sensitivity net of its impact on dTRM and dJV.

Appendix E. Taylor-series Motivation

A complementary perspective on the factor structure in option returns can be obtained from a Taylor expansion of the option value dynamics:

$$\mathrm{d}O = O_S\,\mathrm{d}S + O_t\,\mathrm{d}t + O_V\,\mathrm{d}V + \frac{1}{2}O_{SS}\,(\mathrm{d}S)^2 + \frac{1}{6}O_{SSS}(dS)^3 + \ldots\,, \tag{E13}$$

where O_x, O_{xy}, and O_{xyz} denote derivatives of the option price with respect to pricing inputs x, y, and z up to third order; in our case the inputs being the underlying value, S, and the level of variance of the underlying, V. Equation (E13) accounts for the nonlinear impact of the underlying asset's dynamics and the first-order effect of variance changes on option prices. To obtain the delevered option return defined in equation (1), we multiply both sides of equation (E13) by $\frac{1}{S}$ and collect terms in such a manner that the expansion is in terms of the powers of the underlying return $\mathrm{d}S/S$ and changes in variance $\mathrm{d}V$:

$$\frac{\mathrm{d}O}{S} = O_S \frac{\mathrm{d}S}{S} + \frac{O_t}{S}\,\mathrm{d}t + \frac{O_V}{S}\,\mathrm{d}V + \frac{1}{2}S \times O_{SS}\left(\frac{\mathrm{d}S}{S}\right)^2 + \frac{1}{6}S^2 \times O_{SSS}\left(\frac{dS}{S}\right)^3 + \ldots \qquad \text{(E14)}$$

Equation (E14) states that an option delevered return's direct sensitivity to the underlying asset return is equal to its delta, while its sensitivity to the second and third powers of the underlying return are scaled versions of its gamma and third-order price sensitivity, respectively. For instance, in the Black–Merton–Scholes model, $O_{SS} = \frac{\phi(d_1(\kappa,\tau))}{S\sigma\sqrt{\tau}}$, where σ is the implied volatility and τ the option's maturity. In equation (E14), after scaling by S and accounting for the Taylor expansion, we obtain $\frac{\phi(d_1)}{2\sigma\sqrt{\tau}}$. One of the benefits of the delevered return representation is that the sensitivities in equation (E14) do not depend on the level of the underlying asset if one considers a fixed moneyness level $\kappa = K/S$. The option's time decay term O_t/S is accounted for by the conditional intercept of the option.[5]

Using the BMS pricing formula as a bijection between the space of options prices and that of implied volatilities, Carr and Wu (2020) introduce a decentralized option valuation framework where V in equation (E13) is replaced by each option's implied volatility, and where exposures are calculated using the BMS model. Our approach differs by constructing a set of factors that are common to all options and employing a nonparametric estimation of the scaled risk exposures.

Appendix F. Standard Errors for Risk-Premium Estimates

Three layers of estimation enter Table II: the first stage delivers exposures and conditional intercepts as in equation (7) (first-stage model in linear regression form), the period-by-period cross-sectional regressions (12) deliver the $\widehat{\mu}^Q_{l,t}$, and time-series averaging delivers $\overline{f}_l$, $\overline{\mu}^Q_l$, and $\widehat{\lambda}_l = \overline{f}_l - \overline{\mu}^Q_l$. The standard errors reflect all three layers.

For $\overline{f}_l$, we use heteroskedasticity- and autocorrelation-consistent (HAC) standard errors with a Bartlett kernel and $\lceil T^{1/3} \rceil$ lags, where T is the number of 30-minute periods in the sample. For $\overline{\mu}^Q_l$, the variance must combine the period-to-period dispersion of the cross-

[5]See the derivation of the expected return on an option in Broadie, Chernov, and Johannes (2009). In this derivation, the time decay is substituted out via the pricing PDE to accommodate the Q-expectations of the factors.

sectional estimates $\widehat{\mu}^{Q}_{l,t}$ with the sampling error carried over from the first stage, because both the regressand $\widehat{c}_{i,t}$ and the regressors $\widehat{\beta}_{l,i,t}$ in (12) are generated quantities.

Collect the pooled coefficients of equation (7) into $\boldsymbol{\theta} = (\boldsymbol{q}', \boldsymbol{b}_1', \ldots, \boldsymbol{b}_L')'$, with estimate $\widehat{\boldsymbol{\theta}}$ and covariance $\widehat{\boldsymbol{\Sigma}}_{\boldsymbol{\theta}}$, the latter a Driscoll–Kraay panel estimator (cross-sectional averaging within each 30-minute period, Newey–West across periods); because the first stage is estimated separately for each maturity, $\widehat{\boldsymbol{\Sigma}}_{\boldsymbol{\theta}}$ is block-diagonal across maturities. The conditional intercept and the exposures are linear in these coefficients,

$$\widehat{c}_{i,t} = \boldsymbol{\phi}_{i,t}'\widehat{\boldsymbol{q}}\,, \qquad \widehat{\beta}_{l,i,t} = \boldsymbol{\phi}_{i,t}'\widehat{\boldsymbol{b}}_l\,, \tag{F15}$$

so their derivatives with respect to $\widehat{\boldsymbol{\theta}}$ are the basis vector $\boldsymbol{\phi}_{i,t}$ itself. For each period t, stack the bucket indicators $\boldsymbol{e}_{b(i)}$ and the exposures into $\boldsymbol{z}_{i,t} = (\boldsymbol{e}_{b(i)}', \widehat{\boldsymbol{\beta}}_{i,t}')'$, with $\widehat{\boldsymbol{\beta}}_{i,t} = (\widehat{\beta}_{1,i,t}, \ldots, \widehat{\beta}_{L,i,t})'$, and write (12) as $\widehat{c}_{i,t} = \boldsymbol{z}_{i,t}'\boldsymbol{\xi}_t + \nu_{i,t}$ with $\boldsymbol{\xi}_t = (\boldsymbol{\alpha}_t', -\boldsymbol{\mu}_t^{Q\prime})'$. Weighting each contract by $\omega_{i,t} = 1/N_t$,

$$\widehat{\boldsymbol{\xi}}_t = \boldsymbol{Q}_t^{-1}\sum_i \omega_{i,t}\,\boldsymbol{z}_{i,t}\,\widehat{c}_{i,t}\,, \qquad \boldsymbol{Q}_t = \sum_i \omega_{i,t}\,\boldsymbol{z}_{i,t}\,\boldsymbol{z}_{i,t}'\,, \tag{F16}$$

and $\widehat{\boldsymbol{\mu}}_t^{Q} = -\boldsymbol{S}\,\widehat{\boldsymbol{\xi}}_t$, where $\boldsymbol{S}$ selects the exposure block; $\overline{\boldsymbol{\mu}}^{Q} = |\mathcal{T}_*|^{-1}\sum_{t\in\mathcal{T}_*}\widehat{\boldsymbol{\mu}}_t^{Q}$ averages over the periods $\mathcal{T}_*$ in which all L slopes are identified, and $\widehat{\nu}_{i,t} = \widehat{c}_{i,t} - \boldsymbol{z}_{i,t}'\widehat{\boldsymbol{\xi}}_t$.

Treating the generated quantities as data gives the Fama–MacBeth term, the HAC long-run variance of the period estimates,

$$\boldsymbol{V}_{\mathrm{FMB}} = \frac{1}{|\mathcal{T}_*|^2}\sum_{\ell=-L_\mu}^{L_\mu}\Big(1-\tfrac{|\ell|}{L_\mu+1}\Big)\sum_t\big(\widehat{\boldsymbol{\mu}}_t^{Q}-\overline{\boldsymbol{\mu}}^{Q}\big)\big(\widehat{\boldsymbol{\mu}}_{t-\ell}^{Q}-\overline{\boldsymbol{\mu}}^{Q}\big)'\,, \tag{F17}$$

a Bartlett kernel with $L_\mu = \lceil T^{1/3}\rceil$ lags. The first-stage error enters through the dependence of $\widehat{\boldsymbol{\mu}}_t^{Q}$ on $\widehat{\boldsymbol{\theta}}$. Differentiating (F16) splits this dependence into a part carried by the cross-sectional residual and a part that is not:

$$\boldsymbol{\Psi}_t = -\boldsymbol{S}\,\boldsymbol{Q}_t^{-1}\sum_i \omega_{i,t}\,\boldsymbol{z}_{i,t}\big(\boldsymbol{\Phi}_{i,t}^{c} - \widehat{\boldsymbol{\xi}}_t'\widetilde{\boldsymbol{\Phi}}_{i,t}^{\beta}\big)\,, \tag{F18}$$

$$\boldsymbol{H}_t = -\boldsymbol{S}\,\boldsymbol{Q}_t^{-1}\sum_i \omega_{i,t}\,\widehat{\nu}_{i,t}\,\widetilde{\boldsymbol{\Phi}}_{i,t}^{\beta}\,, \tag{F19}$$

where $\boldsymbol{\Phi}_{i,t}^{c} = \partial\widehat{c}_{i,t}/\partial\boldsymbol{\theta}'$ places $\boldsymbol{\phi}_{i,t}'$ in the $\boldsymbol{q}$-block, and $\widetilde{\boldsymbol{\Phi}}_{i,t}^{\beta} = \partial\boldsymbol{z}_{i,t}/\partial\boldsymbol{\theta}'$ places $\boldsymbol{\phi}_{i,t}'$ in each $\boldsymbol{b}_l$-block of the exposure rows and zero on the bucket rows. The residual-free sensitivities $\boldsymbol{\Psi}_t$ multiply the common first-stage error coherently and aggregate before being sandwiched, whereas the residual score $\boldsymbol{H}_t$ inherits the serial dependence of $\widehat{\nu}_{i,t}$ and is accumulated with

the same Bartlett kernel:

$$\boldsymbol{V}_{\text{fs}} = \frac{1}{|\mathcal{T}_*|^2}\Big(\sum_t \boldsymbol{\Psi}_t\Big)\,\widehat{\boldsymbol{\Sigma}}_{\boldsymbol{\theta}}\,\Big(\sum_t \boldsymbol{\Psi}_t\Big)', \tag{F20}$$

$$\boldsymbol{V}_{xe} = \frac{1}{|\mathcal{T}_*|^2}\sum_{\ell=-L_\mu}^{L_\mu}\Big(1-\tfrac{|\ell|}{L_\mu+1}\Big)\sum_t \boldsymbol{H}_t\,\widehat{\boldsymbol{\Sigma}}^{bb}_{\boldsymbol{\theta}}\,\boldsymbol{H}'_{t-\ell}\,, \tag{F21}$$

with $\widehat{\boldsymbol{\Sigma}}^{bb}_{\boldsymbol{\theta}}$ the exposure-coefficient block of $\widehat{\boldsymbol{\Sigma}}_{\boldsymbol{\theta}}$. The variance of $\overline{\mu}^Q_l$ is the l-th diagonal entry of

$$\boldsymbol{V}\big(\overline{\boldsymbol{\mu}}^Q\big) = \boldsymbol{V}_{\text{FMB}} + \boldsymbol{V}_{\text{fs}} + \boldsymbol{V}_{xe}\,; \tag{F22}$$

dropping the last two terms recovers the uncorrected Fama–MacBeth variance. Because $\widehat{c}_{i,t}$ and $\widehat{\boldsymbol{\beta}}_{i,t}$ are read from the same first stage, the full $\widehat{\boldsymbol{\Sigma}}_{\boldsymbol{\theta}}$—including the $\widehat{\boldsymbol{q}}$–$\widehat{\boldsymbol{b}}$ cross-block—enters $\boldsymbol{V}_{\text{fs}}$. For the risk premium, the variance combines the pieces,

$$\text{Var}\big(\widehat{\lambda}_l\big) = \text{Var}\big(\overline{f}_l\big) - 2\,\text{Cov}\big(\overline{f}_l,\,\overline{\mu}^Q_l\big) + \text{Var}\big(\overline{\mu}^Q_l\big)\,, \tag{F23}$$

where the first two terms are computed with the $\lceil T^{1/3}\rceil$-lag Bartlett kernel from the time series of $f_{l,t+1}$ and $\widehat{\mu}^Q_{l,t}$, and the last term, $\text{Var}(\overline{\mu}^Q_l)$, is the l-th diagonal entry of (F22), carrying the generated-regressor correction.

Inference elsewhere in the paper follows the same logic, applied series by series. The significance stars in the summary-statistics tables are based on HAC standard errors for the mean of each underlying time series, with $\lceil T^{1/3}\rceil$ lags. The regressions of pricing errors on economic-environment variables (Table V and its appendix counterparts) are estimated on daily data separately for each bucket and option type, with Newey–West standard errors with $\lceil T^{1/3}\rceil$ lags; the persistence of the alphas documented in Table IV is thereby accounted for in the reported significance levels.

Appendix G. Spot-futures Spread

To proxy for liquidity conditions in the underlying futures market, we follow Siriwardane, Sunderam, and Wallen (2025) and construct the spot-futures spread, defined as the difference between the forward rate implied by the futures and the 3-month OIS swap rate.

Let F_{t,τ_1} and F_{t,τ_2} denote, respectively, the nearby and first-deferred futures prices at date t, with maturities $\tau_1 < \tau_2$. Under no-arbitrage condition, the implied forward rate yields

$$f_t^{\tau_1,\tau_2} = \left(\frac{F_{t,\tau_2} + \mathbb{E}^Q_t[D_{t,\tau_2}]}{F_{t,\tau_1} + \mathbb{E}^Q_t[D_{t,\tau_1}]} - 1\right)\frac{1}{\tau_2 - \tau_1} \tag{G24}$$

where $\mathbb{E}t^{Q}[D_{t,\tau}]$ denotes the risk-neutral expectation of cumulative dividends over $[t, t+\tau]$. We implement this construction for S&P 500 E-mini (ES) and Russell 2000 E-mini (RUT) futures using the nearby and first-deferred contract close prices at 16:15 ET on each day. ES futures prices are obtained from TickData Inc., and RUT futures prices are obtained from Bloomberg. Following Siriwardane, Sunderam, and Wallen (2025), we proxy for $\mathbb{E}_t^Q[D_{t,\tau}]$ using realized dividends, obtained from Bloomberg. We compute the future values of dividends using the risk-free rate, measured by SOFR; missing observations are filled using LIBOR. Finally, the spot-futures spread is defined as the difference between the implied forward rate and the 3-month OIS swap rate:

$$ESF_t = f_t^{\tau_1,\tau_2} - OIS_t^{3M} \tag{G25}$$

where all the interest-rate series are obtained from Bloomberg. Figure IA.VI plots the *ESF* series.

Appendix H. Robustness of the Pricing-Error Regressions

Section 5.4 discusses the economic mechanisms behind the time-series variation in option pricing errors. In particular, Table V reports results for alphas and their magnitudes in ATM options with different maturities.

We here provide evidence that those results are robust along the moneyness dimension, as evidenced by Tables IA.VII (for OTM options), and IA.VIII (for DOTM options).

In Tables IA.IX to IA.XI, we also address a concern that the alphas could be mechanically related to the day's underlying return, $o2c$, via our estimation process. We split the *SP500* variable into two components: the return from market open until 14:45 ET ($SP500_{[09:45,14:45]}$) and from 14:45 to 15:15 ET ($SP500_{[14:45,15:15]}$). The return in the last 30 minutes of trading before alpha measurement is not significant in any of the regressions, except for certain DOTM put cases, where it's only significant at the 10% level. Negative market returns are associated with a decrease in alphas of 0DTE options of both types, and of longer-maturity puts, with the changes much larger in 0DTEs, and with the highest explanatory power in this bucket—up to 4%, vs. below 1% in longer maturities. In other words, options become more expensive when the market declines. The magnitude of 0DTE pricing errors also increases on down days, with patterns broadly similar to those for the level of the alphas.

Figure IA.I. The Intraday Volatility of the ES Futures Returns

The figure plots the average 1-minute realized volatility of the E-mini S&P 500 futures return across days, from the prior day 18:00 ET to the current day 18:00 ET, covering the full 24-hour trading cycle from Monday to Thursday. The vertical dashed lines mark the open and close time of different regimes: Global Trading Hours (GTH, Prior 18:00 to 09:30 ET), Regular Trading Hours (RTH, 09:30 to 16:00 ET), Post-market (Post, 16:00–16:15 and 16:30–17:00 ET), CME settlement halt (HALT, 16:15–16:30 ET), and the technical break (BRK, 17:00–18:00ET). The red horizontal line represents the average volatility during each regime. The minute-level volatility is computed using five-second ES futures returns, and then averaged across days for each minute-of-day. The volatility is expressed in annualized percentage points ($\sqrt{252 \times 23 \times 60}$). The sample period is from 2016/09 to 2023/12.

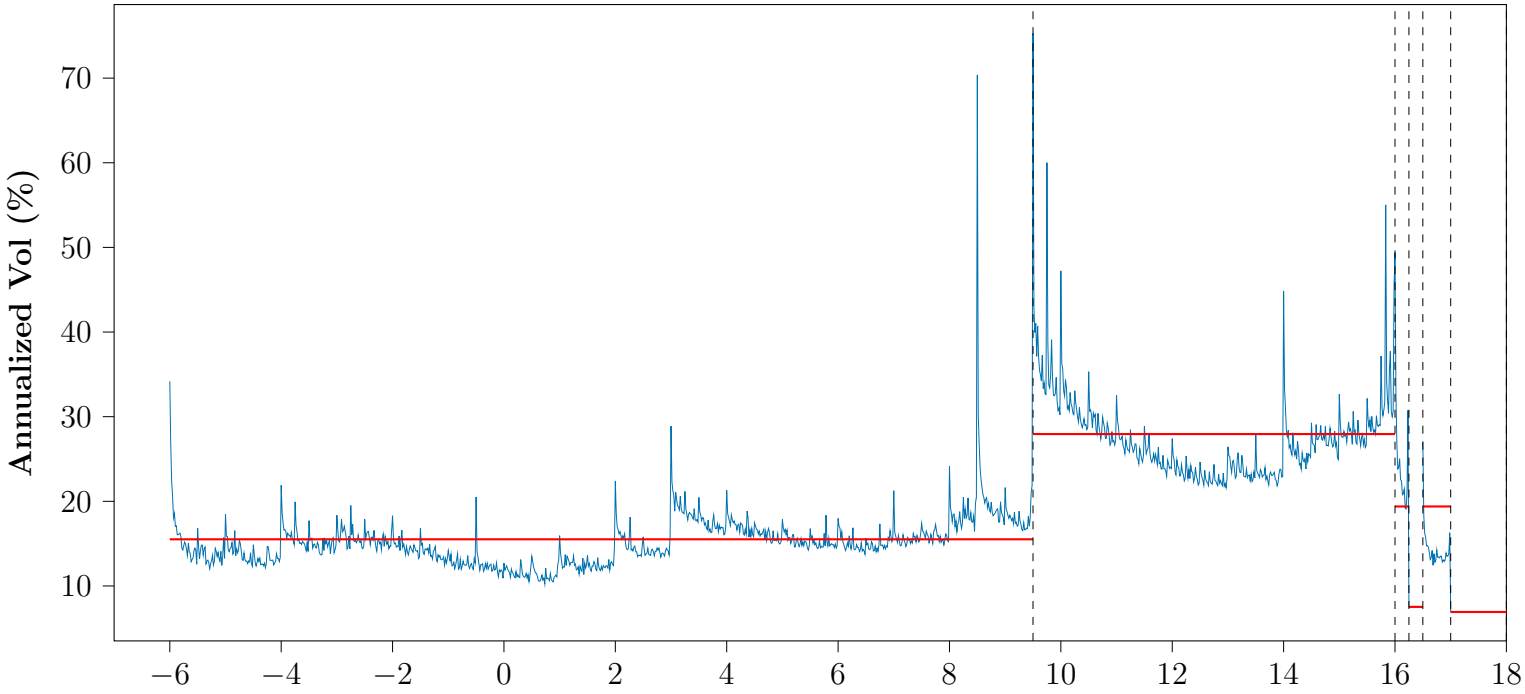


Figure IA.II. The Intraday Cumulative Clocks

The figure plots the cumulative clocks of the calendar time, regular trading hours (RTH) time, and the stochastic clock from the prior day 18:00 ET to the current day 18:00 ET, covering the full 24-hour trading cycle from Monday to Thursday. The calendar clock is the linear fraction of the 24 hours, while the RTH clock counts only the 6.5 hours between 09:30 and 16:00 ET. The stochastic clock is constructed by weighting each minute of the day by its average contribution to the total intraday variance, as described in Section C. The vertical dashed lines mark the open and close time of different regimes: Global Trading Hours (GTH, Prior 18:00 to 09:30 ET), Regular Trading Hours (RTH, 09:30 to 16:00 ET), Post-market (Post, 16:00–16:15 and 16:30–17:00 ET), CME settlement halt (HALT, 16:15–16:30 ET), and the technical break (BRK, 17:00–18:00ET). The sample period is from 2016/09 to 2023/12.

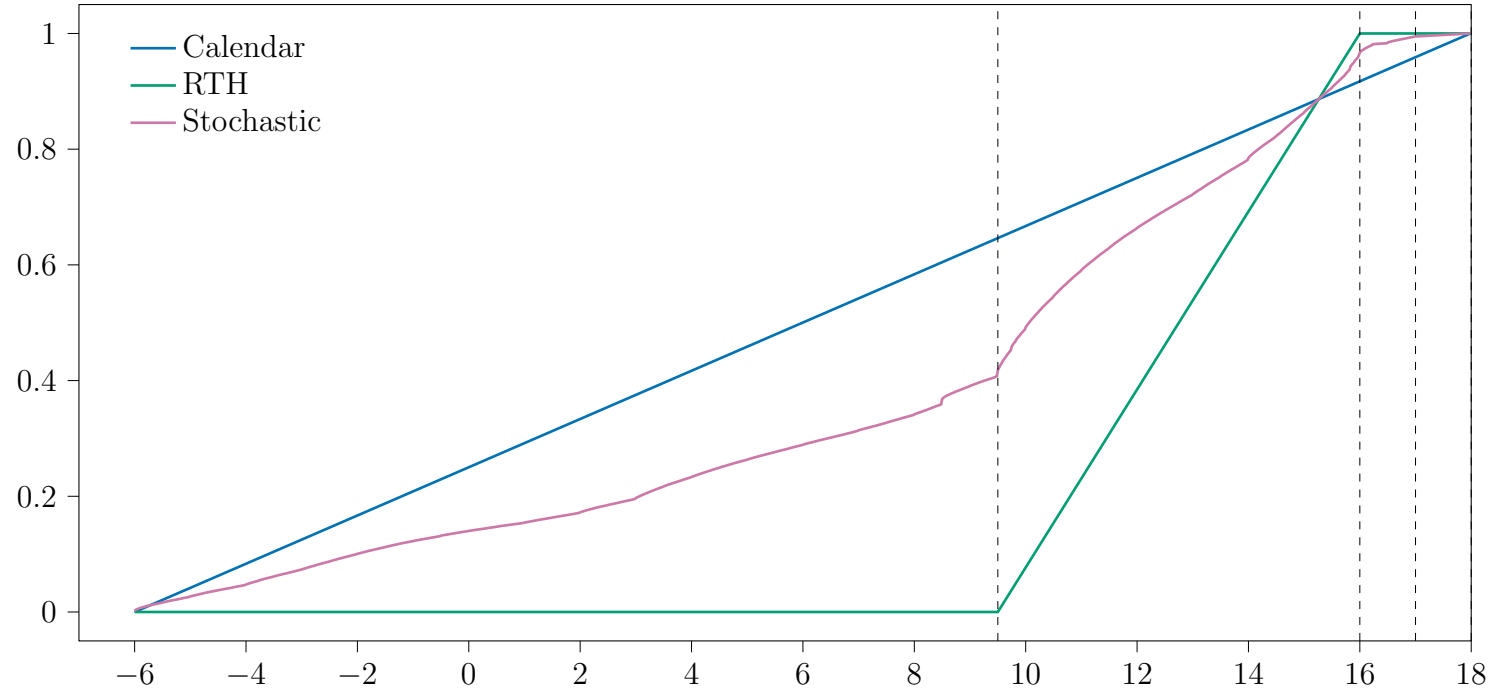


Figure IA.III. The Distribution Density of Moneyness

The figure displays the Gaussian kernel-estimated density of option moneyness, defined as $\kappa_{t,\tau} = \frac{\log(K/F)}{\sqrt{\tau}CIV_{t,\tau}}$, for our sample across three maturity groups (0DTE, 1–3DTE, and 4–15DTE) at three intraday timestamps. The left column shows the moneyness density for puts, while the right column shows the corresponding densities for calls. Each row corresponds to a different intraday time: 9:45, 12:15, and 15:15 ET, respectively. The sample period is from 2016/09 to 2023/12.

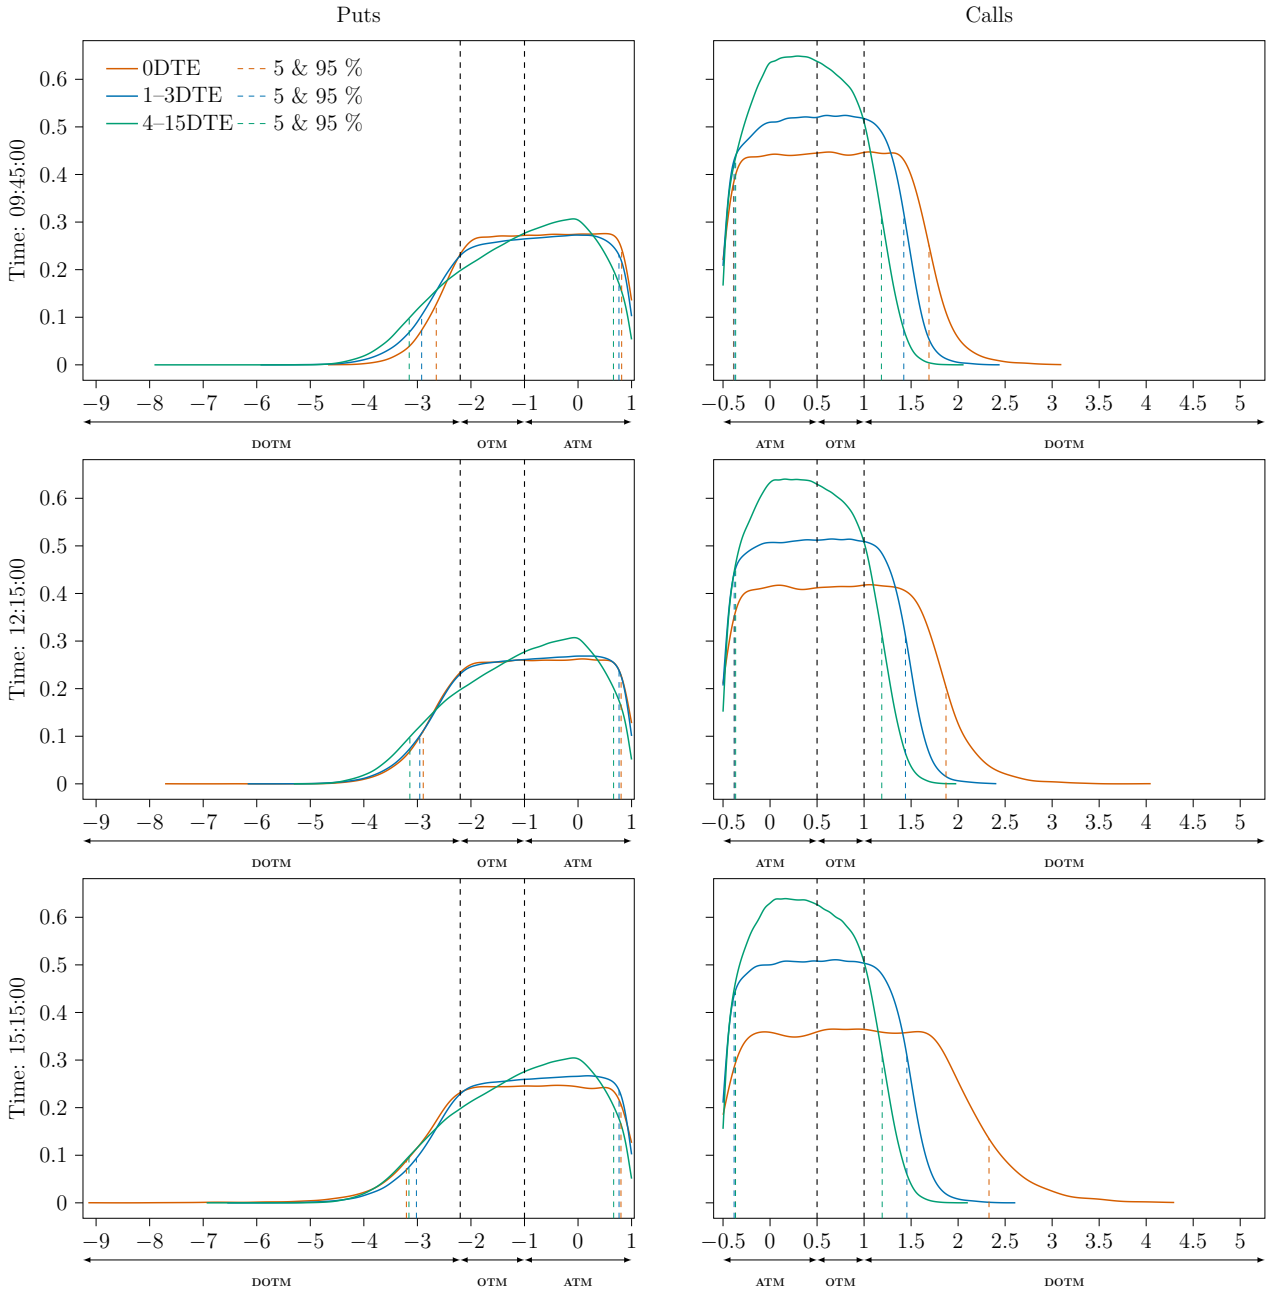


Figure IA.IV. The R^2 of the Factor Model across DTEs

The figure plots the R^2% of the factor model with the specification of Section 3, computed as in equation (8), for delevered unhedged and DLT-hedged option returns across DTEs. The bars report full-sample values, at the position labeled "A" on the horizontal axis. The solid blue bar reports $R_{6f}^{\mathrm{Uh}^2}$%, the fraction of the variation in delevered unhedged option returns explained by the six factors (DLT, GAM, JSK, dSV, dJV and dTRM) and the conditional intercept. The hatched blue bar (middle bar) represents $R_{\mathrm{DLT}}{}^2$%, the R^2 of the DLT factor alone, with the exposure obtained from the six-factor model. The orange bar reports $R_{5f}^{\mathrm{Dh}^2}$%, the fraction of the variation in DLT-hedged option returns explained by the five remaining factors (GAM, JSK, dSV, dJV and dTRM). The lines plot the same three statistics computed within each DTE subsample; the dashed line corresponds to $R_{\mathrm{DLT}}{}^2$%. The sample period is from 2016/09 to 2023/12.

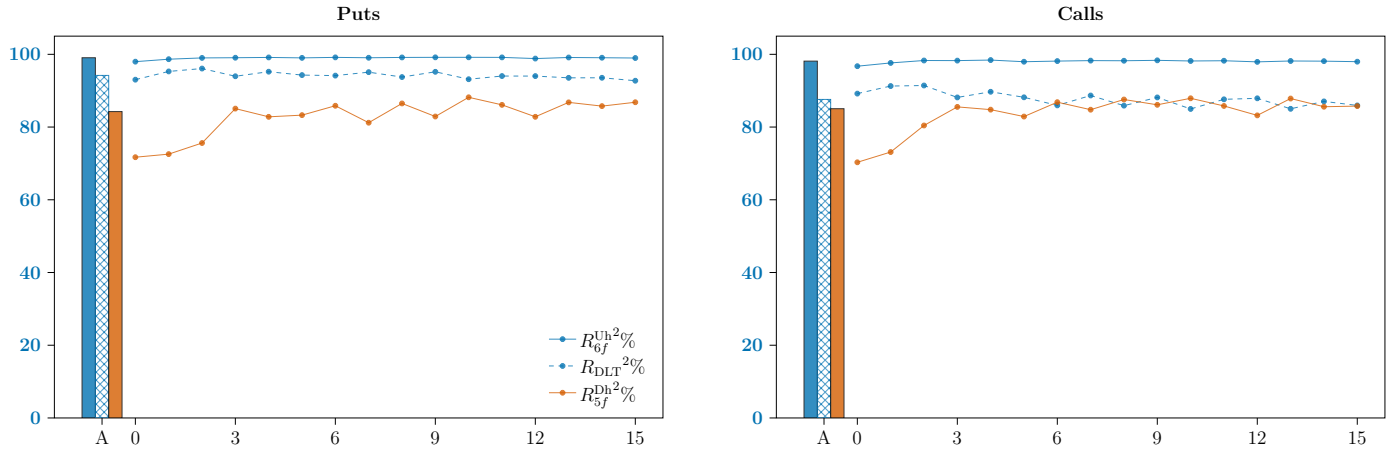


Figure IA.V. The Intraday Pattern of Estimated Conditional Risk-neutral Expectations of Risk Factors from Non-0DTE Options

The figure plots the intraday average conditional risk-neutral expectation, μ_t^Q, of each risk factor,estimated from non-0DTE puts and calls. They are obtained from the cross-sectional regression (12) of the conditional intercepts on option exposures at each 30-minute timestamp for puts and calls, excluding 0DTEs. The ordinate axis labels are in annualized percentage points. The sample period is from 2016/09 to 2023/12.

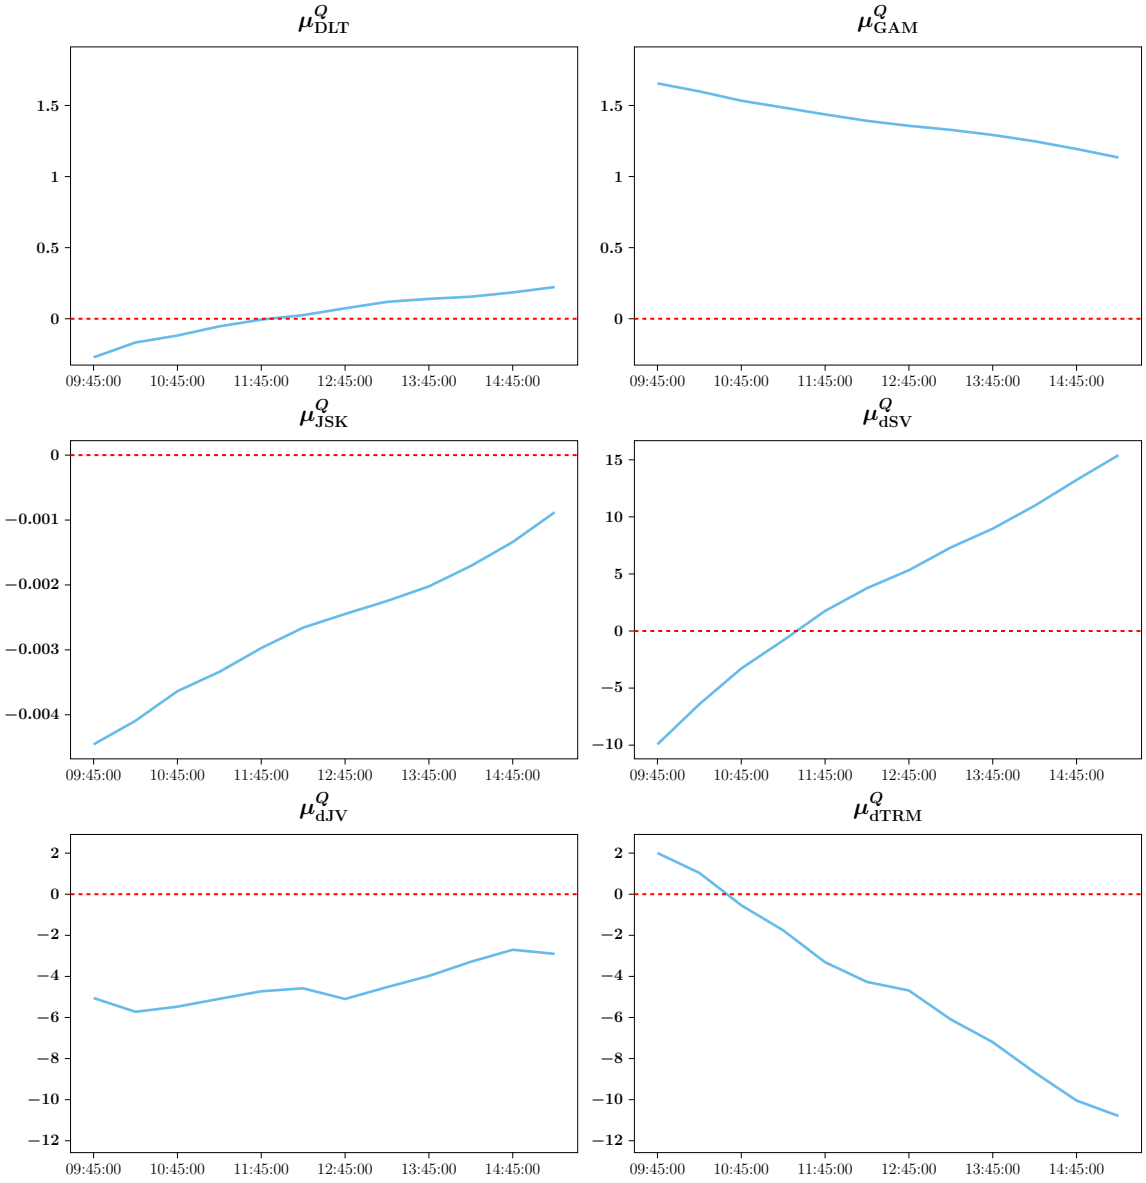


Figure IA.VI. Time-series of ESF (bps)

The figure plots the time-series of the daily spot-futures spread, *ESF*, in basis points (bps). The spot-futures spread is the spread between the forward interest rate implied by futures and the three-month OIS rate, as in Siriwardane, Sunderam, and Wallen (2025) and Internet Appendix G, estimated using futures prices at the futures closing time, 16:15 ET. The blue line represents the *ESF* estimated using S&P 500 E-Mini futures, while the orange line represents the *ESF* estimated using Russell 2000 E-Mini futures. The latter are only available from 2017/07. The sample period is from 2016/09 to 2023/12.

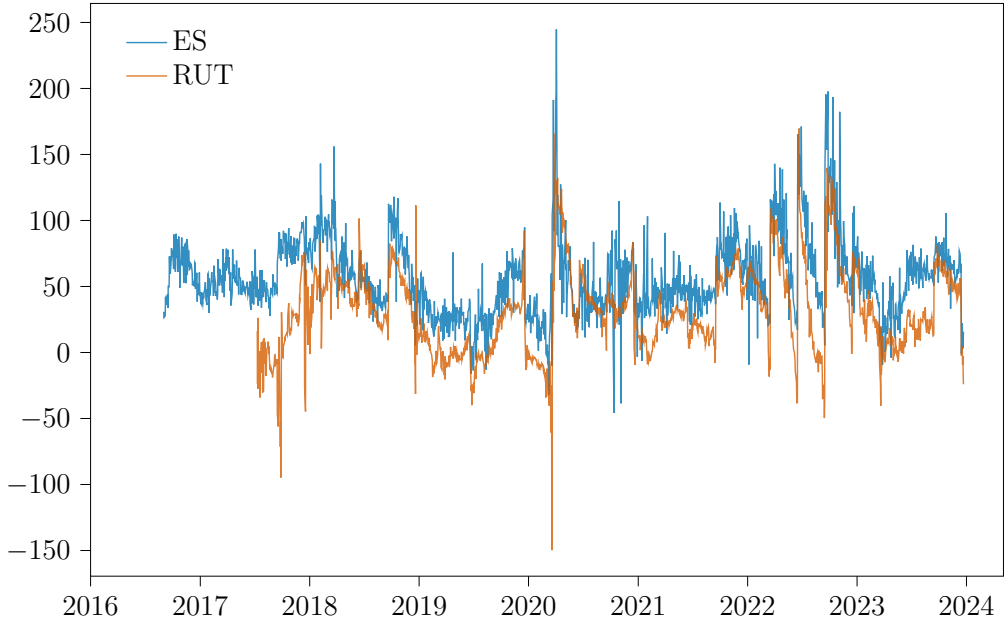


Figure IA.VII. Time-series of Non-0DTE Estimated Alphas

The figure plots the time-series of 1+DTE bucket-specific alphas, $\widehat{\alpha}_{b(i),t}$. They are obtained from the cross-sectional regression (14) of the conditional intercepts on option exposures at each 30-minute timestamp for puts and calls, respectively. Within each option type, the options are assigned to one of nine buckets based on maturity (0DTE, 1–3DTE, 4–15DTE) and moneyness (DOTM, OTM, ATM), for 18 type-specific buckets in total. This figure displays the estimated $\widehat{\alpha}_{b(i),t}$ for the non-0DTE buckets, estimated from returns between 09:45–10:15 ET, 12:15–12:45 ET, and 15:15–15:45 ET, respectively (top, middle, and bottom row); within each maturity block, the left (right) column reports puts (calls), and colors distinguish the moneyness buckets. The vertical red dashed line marks April 1, 2022; the new Tuesday and Thursday expirations began trading on April 18 and May 11, 2022, respectively. The vertical gray dashed lines indicate FOMC announcement days, including scheduled and emergency announcements. The ordinate axis labels are in annualized percentage points. The sample period is from 2016/09 to 2023/12.

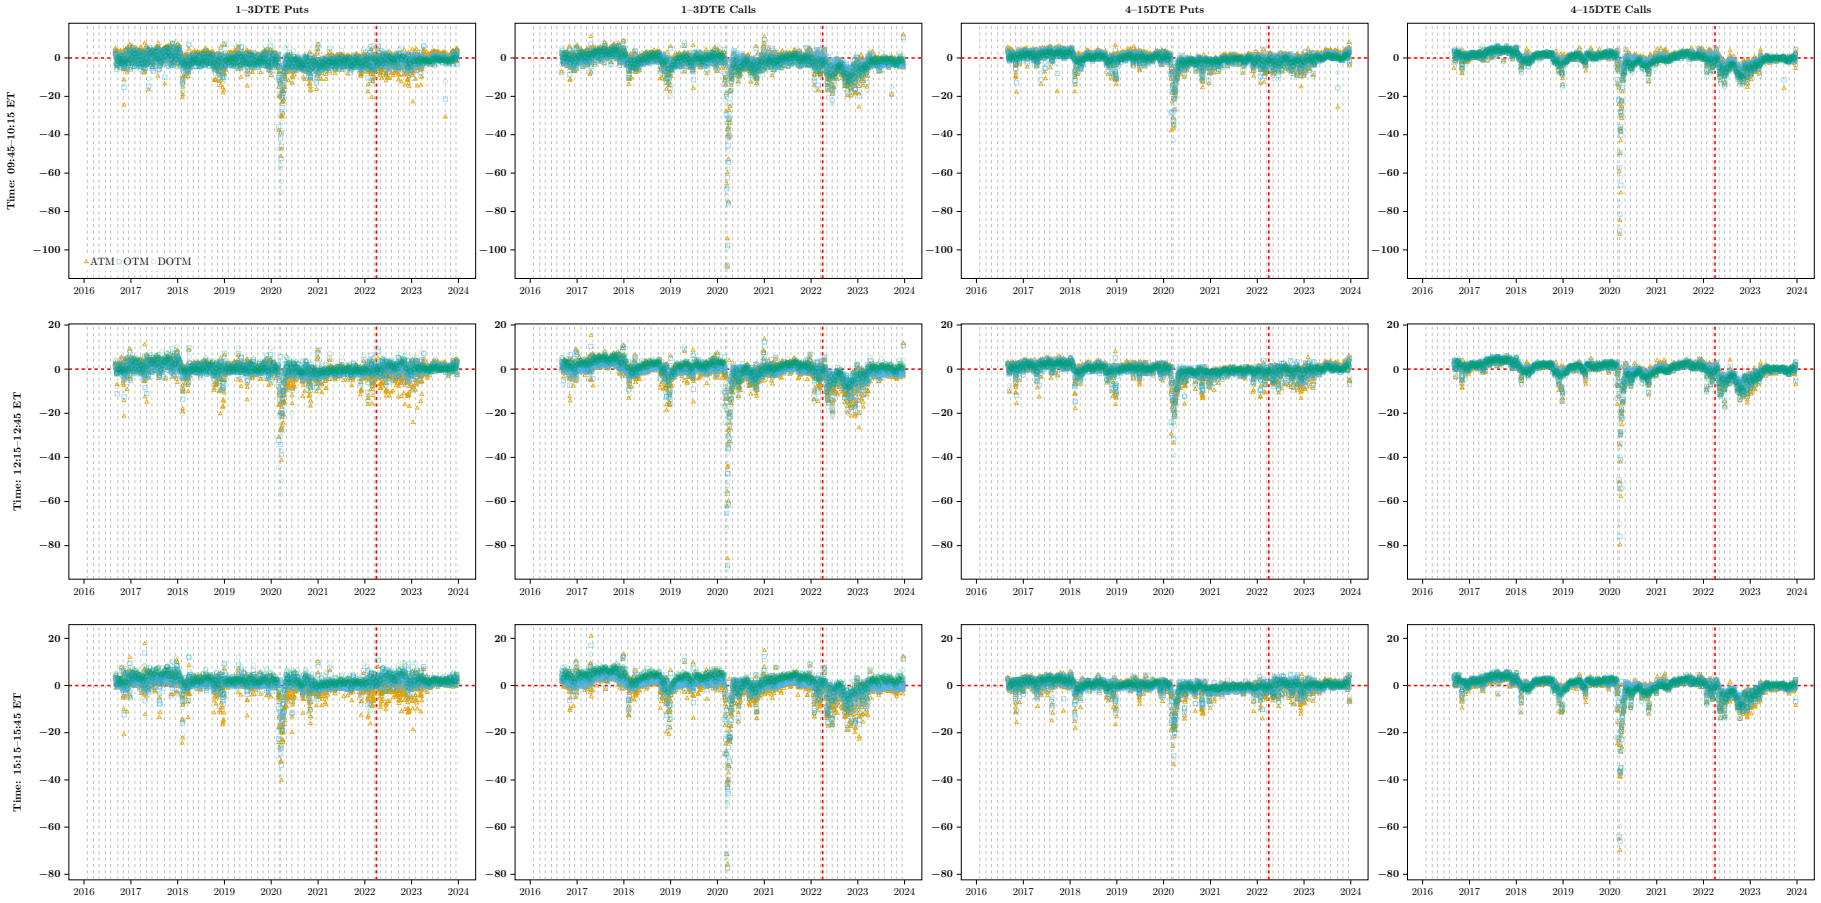


Figure IA.VIII. The Intraday Pattern of 0DTE Estimated Alphas on FOMC and Non-FOMC Days

The figure plots the intraday average 0DTE alphas, $\widehat{\alpha}_{b(i),t}$, at each intraday timestamp across FOMC and non-FOMC days. The bucket-specific alphas are estimated from the cross-sectional regression (14) of the conditional intercepts on option exposures for puts and calls, at each 30-minute timestamp. The options are assigned into 9 buckets based on option maturities (0DTE, 1–3DTE, 4–15DTE) and moneyness (DOTM, OTM, ATM). This figure plots the intraday average $\widehat{\alpha}_{b(i),t}$ for the 0DTE buckets only, computed across FOMC days (the second day of FOMC meetings) and non-FOMC days. The ordinate axis labels are in annualized percentage points. The sample period is from 2016/09 to 2023/12.

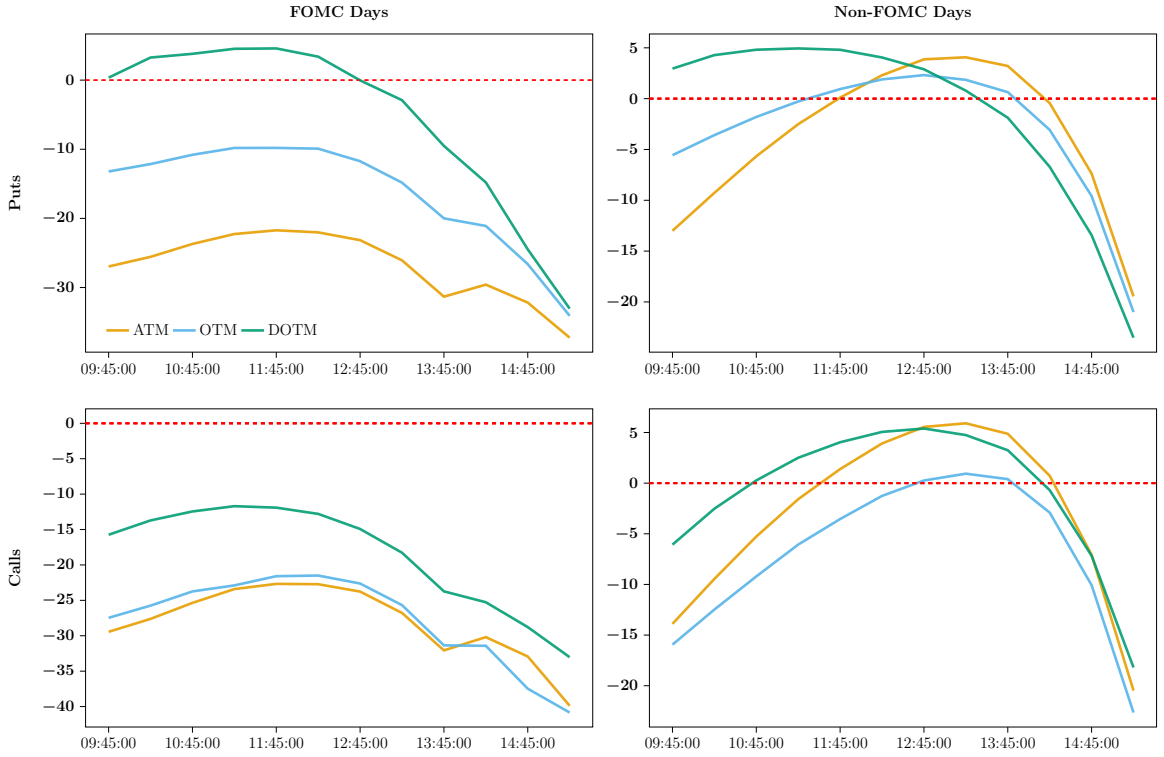


Figure IA.IX. The Intraday Pattern of 1–3DTE Estimated Alphas across Different Periods

The figure plots the intraday average 1–3DTE alphas, $\widehat{\alpha}_{b(i),t}$, at each intraday timestamp across different periods. The bucket-specific alphas are estimated from the cross-sectional regression (14) of the conditional intercepts on option exposures for puts and calls, at each 30-minute timestamp. The options are assigned into 9 buckets based on option maturities (0DTE, 1–3DTE, 4–15DTE) and moneyness (DOTM, OTM, ATM). This figure plots the intraday average $\widehat{\alpha}_{b(i),t}$ for the 1–3DTE buckets only, computed across four subsamples: pre-April 2022, post-April 2022, pre-April 2022 Tuesdays and Thursdays, and pre-April 2022 Mondays, Wednesdays and Fridays. All the subsamples include only normal trading days (non-FOMC days and excluding March 2020). The ordinate axis labels are in annualized percentage points. The sample period is from 2016/09 to 2023/12.

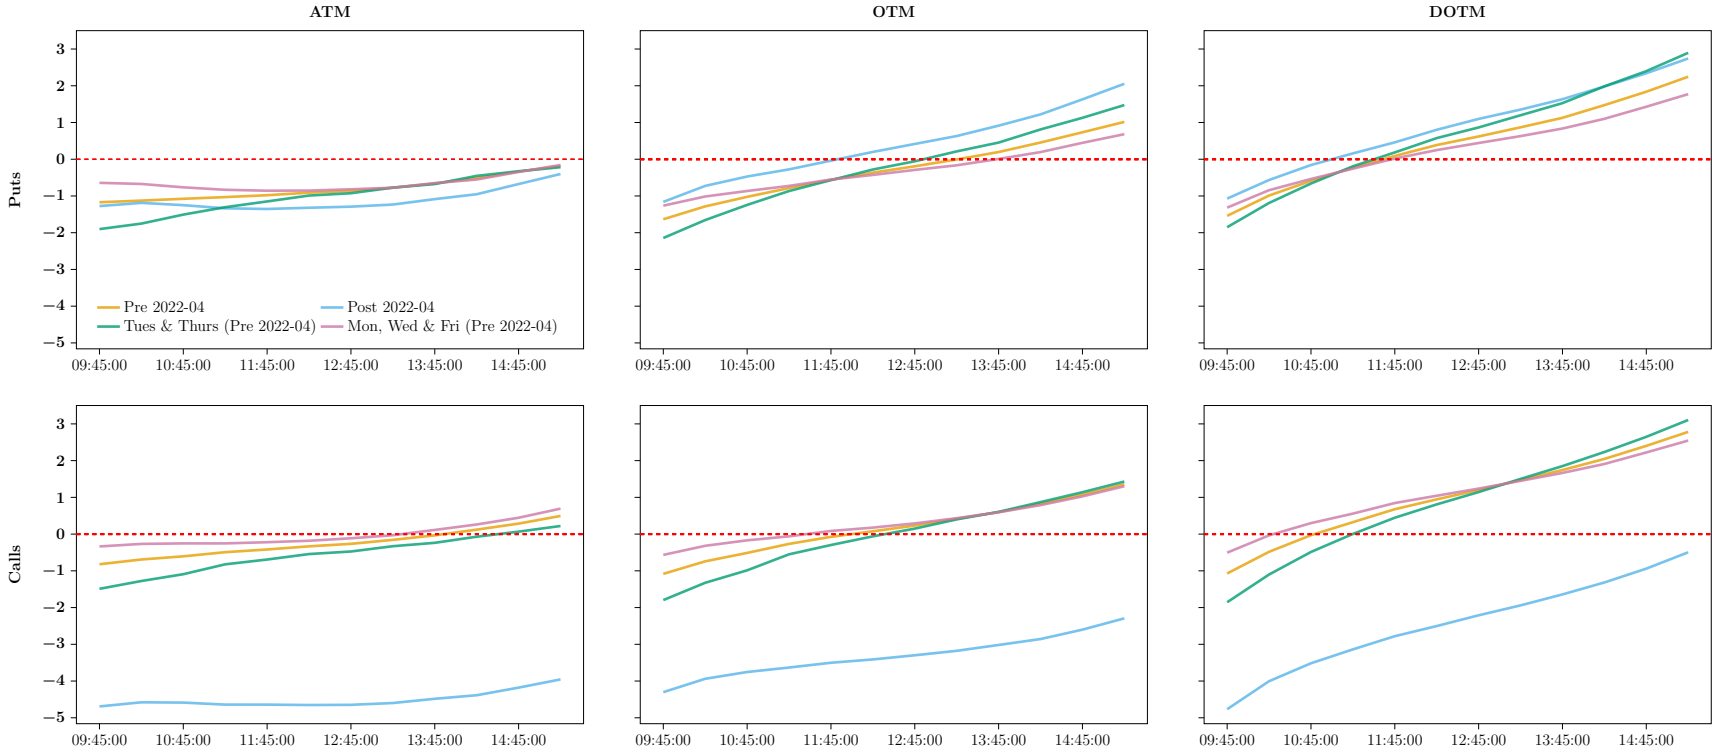


Figure IA.X. Intraday Patterns of Self-financing Six-factor-neutral Portfolio Returns

This figure plots the intraday patterns of the six-factor-neutral returns of the 0DTE options hedged by instrument option contracts across different periods using self-financing strategy. The hedging procedure is implemented foreach individual 0DTE put and call contract, within each moneyness bucket (ATM ,OTM, DOTM), at each 30-min timestamp. The non-zero options are assigned into 6 buckets based on option maturities (1–3 DTE, 4–15DTE) and moneyness (ATM, OTM, DOTM). The hedging instrument contracts are constructed, at each 30-min timestamp, by selecting top 3 traded contracts from each bucket of non-zero puts and calls, yielding a total of 18 instrument contracts. The factor-neutral returns are averaged cross-sectionally at each timestamp, and then averaged over days to obtain the intraday patterns. The mean return is then calculated by $\frac{1}{T}\sum_{t\in T}\frac{1}{n_t}\sum_{i\in n_t}R^{dh}_{N,i,t}$. The top row plots show the intraday patterns during the full sample period, while the middle and bottom rows plot the intraday patterns during the FOMC days and non-FOMC days, respectively. The shaded areas represent the 95% confidence intervals of the intraday means. The ordinate axis labels are in annualized percentage points (multiplied by 252×100). The sample period is from 2016/09 to 2023/12.

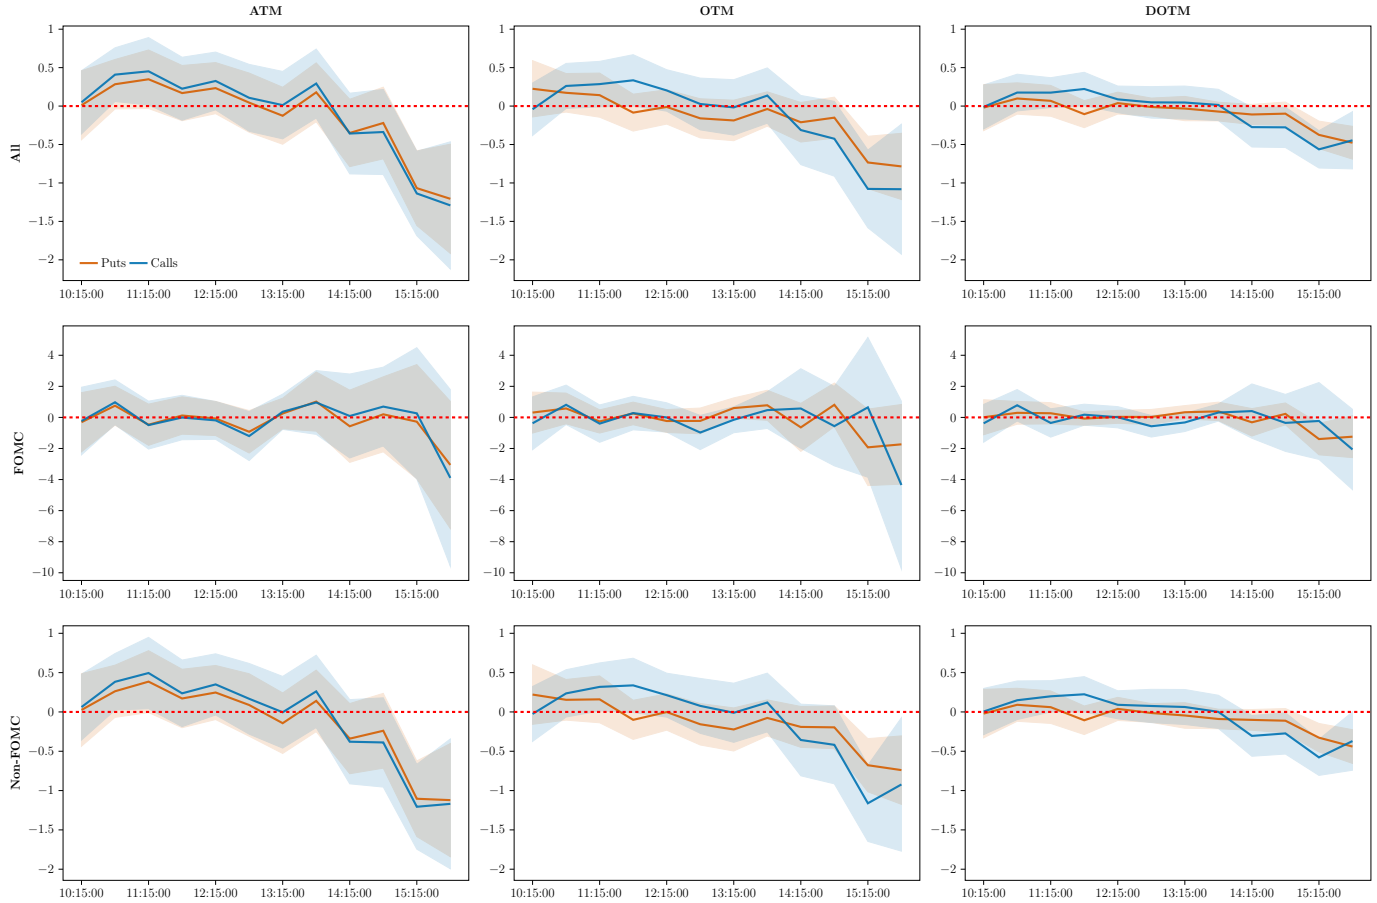


Figure IA.XI. Cumulative Returns of Self-financing Six-factor-neutral Portfolio

This figure plots the time-series of the cumulative six-factor-neutral returns of the 0DTE options hedged by instrument option contracts using self-financing strategy, at 10:15, 12:45 and 15:45 ET, respectively. The hedging procedure is implemented for each individual 0DTE put and call contract, within each moneyness bucket (ATM ,OTM, DOTM), at each 30-min timestamp. The non-zero options are assigned into 6 buckets based on option maturities (1–3 DTE, 4–15DTE) and moneyness (ATM, OTM, DOTM). The hedging instrument contracts are constructed, at each 30-min timestamp, by selecting top 3 traded contracts from each bucket of non-zero puts and calls, yielding a total of 18 instrument contracts. The factor-neutral returns are averaged cross-sectionally at each timestamp. We also report the cumulative return of a long position of ES futures initiated at the same intraday time and hold for the same period as the option trade. The hedging portfolios are leveraged such that the unconditional (whole sample) volatility of the hedged returns matches that of the corresponding ES futures returns. It means that the factor-neutral return is multiplied by $\frac{\sigma_{ES}}{\sigma_{b,N}}$, where σ_{ES} is the unconditional volatility of ES futures returns, and $\sigma_{b,N}$ is the unconditional volatility of the risk-factor-neutral returns within each moneyness bucket b. The cumulative returns are summed across previous days at each intraday timepoint. The ordinate axis labels are in percentage (multiplied by 100). The sample period is from 2016/09 to 2023/12.



Table IA.I
Summary Statistics of Intraday 0DTE DLT-hedged Delevered Returns

The table reports the summary statistics of 0DTE intraday delevered returns at each 30-minute interval. The number of observations is labeled as N. The table reports the average return (Mean), 25th percentile, median (50th), and 75th percentile, all reported in percent per annum (multiplied by 12×252). We delta-hedge the options with the S&P 500 E-Mini futures returns as in equation (2). The option delta, $\widehat{\beta}_{\mathrm{DLT},i,t}$, is estimated using the factor model described in Section 3. The sample period is from 2016/09 to 2023/12.

	Puts					Calls				
Time	N	Mean	25th	50th	75th	N	Mean	25th	50th	75th
10:15:00	22,987	−10.45***	−54.50	−18.65	11.37	14,319	−13.55***	−59.80	−19.33	14.64
10:45:00	21,413	−10.37***	−44.45	−15.91	6.57	13,408	−10.75***	−47.58	−16.06	8.06
11:15:00	20,126	−5.88***	−37.74	−13.30	7.68	12,570	−6.66***	−40.69	−13.64	9.11
11:45:00	19,029	3.42	−32.47	−10.08	9.69	11,932	0.19	−34.48	−11.69	8.00
12:15:00	18,101	−4.91***	−32.81	−10.87	5.59	11,381	−3.93*	−33.73	−12.04	6.95
12:45:00	17,107	−2.02	−29.19	−9.31	6.23	10,840	−0.84	−28.81	−9.53	7.58
13:15:00	16,197	−3.83**	−29.69	−8.66	6.85	10,330	−5.75***	−29.74	−9.27	7.79
13:45:00	15,400	0.25	−26.28	−7.60	8.37	9,801	−0.48	−27.42	−8.82	8.37
14:15:00	14,522	−10.70***	−36.25	−11.17	4.51	9,305	−8.36***	−38.41	−12.26	5.85
14:45:00	13,286	−6.38**	−41.13	−14.63	0.98	8,573	−10.06***	−40.93	−14.22	5.24
15:15:00	11,745	−6.97**	−44.14	−15.98	0.37	7,705	−1.22	−50.36	−14.74	8.28
15:45:00	9,461	−28.15***	−74.66	−24.48	−3.66	6,464	−22.43***	−76.74	−21.23	4.29

⋆ denotes statistically significant difference from zero at the 10% level, ⋆⋆ at the 5% level, and ⋆⋆⋆ at the 1% level.

Table IA.II
Persistence of Risk Factors and their Risk-Neutral Expectations

The table reports the autocorrelation coefficients of the risk factor realizations (f_{t+1}) and their risk-neutral expectations (μ_t^Q). The column labeled "30 minutes" reports the first-order autocorrelation of the series of consecutive 30-minute observations within each day, while the column labeled "1 day" reports the lag-one autocorrelation of the observations at 15:15 ET on successive trading days. The construction of the factors is described in Section 3. The conditional risk-neutral expectations, μ_t^Q, are estimated from cross-sectional regressions of option-specific conditional intercepts on estimated exposures at each 30-minute timestamp, jointly for puts and calls with type-specific bucket fixed effects, as specified in equation (12). The sample period is from 2016/09 to 2023/12.

Factor	Qty	30 min	1 day
DLT	f_{t+1}	−0.02	0.03
	μ_t^Q	0.97	0.69
GAM	f_{t+1}	0.35	0.36
	μ_t^Q	0.99	0.76
JSK	f_{t+1}	−0.16	0.10
	μ_t^Q	0.97	0.20
dSV	f_{t+1}	−0.08	0.07
	μ_t^Q	0.99	0.19
dJV	f_{t+1}	−0.06	0.23
	μ_t^Q	0.96	0.29
dTRM	f_{t+1}	−0.02	0.04
	μ_t^Q	0.97	0.31

Table IA.III
Correlation Matrix of Physical Risk Factors and Risk-neutral Expectations

The table presents the correlation matrices for risk factors under the physical values and conditional risk-neutral expectations. DLT is the market risk measured by the E-mini futures return. GAM denotes the gamma risk factor, computed by the squared DLT. JSK is the cubic return of the S&P 500 E-mini futures, defined as dS3 = DLT^3. The variance risk factors are decomposed to three components: (i) dSV denotes the change in spot variance proposed by Todorov (2019). (ii) dJV captures the change in the difference between average expected variance, CIV_{2_+DTE}, at the two shortest available maturities, satisfying DTE $>$ 2 and calendar maturity $>$ 4 and the corresponding spot variance, dJV = $(CIV_{2_+DTE,t+1} - \text{SV}_{t+1}) - (CIV_{2_+DTE,t} - \text{SV}_t)$. (iii) dTRM represents the shifts in the term structure of CIV_τ: dTRM = $(CIV_{15DTE,t+1} - CIV_{2_+DTE,t+1}) - (CIV_{15DTE,t} - CIV_{2_+DTE,t})$, where CIV_{15DTE} denotes the CIV with the longest available maturity. The aggregated variance risk factor, dVAR, is the sum of the three variance components: dVAR = dSV+dJV+dTRM. The conditional risk-neutral expectations, μ_t^Q, are estimated from the cross-sectional regression, as specified in equation (12), by regressing the six-factor model conditional intercepts on estimated exposures at each 30-minute timestamp, jointly for puts and calls with type-specific bucket fixed effects. The risk-neutral expectation of aggregated variance is also the sum of expectations of the three variance components: $\mu^Q_{\text{dVAR}} = \mu^Q_{\text{dSV}} + \mu^Q_{\text{dJV}} + \mu^Q_{\text{dTRM}}$. The sample period is from 2016/09 to 2023/12.

	DLT	GAM	JSK	dSV	dJV	dTRM	dVAR
DLT	1.00						
GAM	0.03	1.00					
JSK	0.39	0.35	1.00				
dSV	−0.56	−0.04	−0.47	1.00			
dJV	−0.33	−0.10	−0.34	0.31	1.00		
dTRM	0.37	0.09	0.42	−0.71	−0.80	1.00	
dVAR	−0.65	−0.07	−0.51	0.75	0.76	−0.75	1.00
	μ^Q_{DLT}	μ^Q_{GAM}	μ^Q_{JSK}	μ^Q_{dSV}	μ^Q_{dJV}	μ^Q_{dTRM}	μ^Q_{dVAR}
μ^Q_{DLT}	1.00						
μ^Q_{GAM}	−0.20	1.00					
μ^Q_{JSK}	−0.04	−0.81	1.00				
μ^Q_{dSV}	0.43	−0.50	0.38	1.00			
μ^Q_{dJV}	0.00	−0.43	0.14	0.39	1.00		
μ^Q_{dTRM}	−0.14	0.19	−0.07	−0.71	−0.74	1.00	
μ^Q_{dVAR}	0.30	−0.77	0.47	0.71	0.73	−0.50	1.00

Table IA.IV
Explanatory Power of DLT and Aggregated Six Factors for Unhedged Delevered Option Returns

This table reports the explanatory power of the S&P E-Mini futures returns (DLT) and the full model with six risk factors to delevered intraday returns of S&P 500 index options with maturity from 0DTE to 15DTE, estimated separately for each maturity. The R square of DLT is defined as

$$R^2_{\text{DLT}}\% = \left[1 - \frac{\sum_{i,t}\left(r_{i,t+1} - \widehat{\beta}_{\text{DLT},i,t}\text{DLT}_{t+1}\right)^2}{\sum_t (r_{i,t+1})^2}\right] \times 100$$

The exposure of DLT, $\widehat{\beta}_{\text{DLT},i,t}$, is estimated from the six-factor model. Any $\widehat{\beta}_{\text{DLT},i,t}$ with an incorrect sign—negative for a call or positive for a put—is set to zero. The R square of the six-factor model is defined as

$$R^2\% = \left[1 - \frac{\sum_{i,t}\left(r_{i,t+1} - \left(\widehat{c}_{i,t} + \sum_l \widehat{\beta}_{l,i,t} f_{l,t+1}\right)\right)^2}{\sum_{i,t}(r_{i,t+1})^2}\right] \times 100,$$

where $\widehat{c}_{i,t}$ is the factor model's conditional intercept for each option, i, and $\widehat{\beta}_{i,l,t}$ is the estimated option exposure to factor l. The risk factors l include: DLT, GAM, dSV, dJV, dTRM, JSK. The $R^2\%$ are reported for aggregates 0, 1–3, 4–15 DTEs and moneyness groups. The moneyness is defined as $\kappa_{t,\tau} = \frac{ln(K/F_{t,\tau})}{\text{CIV}_{t,\tau}\sqrt{\tau}}$, where $\text{CIV}_{t,\tau}$ is the corridor implied volatility with maturity τ. We divide options into six moneyness groups: DOTM puts (calls) for $\kappa_{t,\tau} \leq -2.2 (\kappa_{t,\tau} > 1)$, OTM puts (calls) for $-2.2 < \kappa_{t,\tau} \leq -1 (0.5 < \kappa_{t,\tau} \leq 1)$ and ATM puts (calls) for $-1 < \kappa_{t,\tau} \leq 1 (-0.5 < \kappa_{t,\tau} \leq 0.5)$. The sample period is from 2016/09 to 2023/12.

		Puts		Calls	
Moneyness	DTE	DLT	6F	DLT	6F
All	All	94.2	99.1	87.6	98.2
	0	93.0	98.0	89.2	96.8
	1–3	95.0	98.9	90.1	98.1
	4–15	94.1	99.1	87.1	98.2
ATM	All	95.0	99.2	91.7	98.7
	0	94.1	98.4	91.9	97.9
	1–3	95.9	99.1	93.5	98.7
	4–15	94.9	99.2	91.4	98.7
OTM	All	68.6	97.1	46.5	93.9
	0	59.0	87.1	74.0	92.7
	1–3	69.4	95.8	65.1	94.6
	4–15	68.7	97.6	41.6	93.8
DOTM	All	22.2	74.1	−13.2	77.5
	0	19.1	30.0	47.8	73.1
	1–3	22.1	69.1	13.1	81.6
	4–15	22.2	75.8	−35.7	76.1

Table IA.V
Explanatory Power of DLT and Aggregated Four Factors for Unhedged Delevered Option Returns

This table reports the explanatory power of the S&P E-Mini futures returns (DLT) and the full model with four risk factors to delevered intraday returns of S&P 500 index options with maturity from 0DTE to 15DTE, estimated separately for each maturity. The R square of DLT is defined as

$$R^2_{\text{DLT}}\% = \left[1 - \frac{\sum_{i,t}\left(r_{i,t+1} - \widehat{\beta}_{\text{DLT},i,t}\text{DLT}_{t+1}\right)^2}{\sum_t (r_{i,t+1})^2}\right] \times 100$$

The exposure of DLT, $\widehat{\beta}_{\text{DLT},i,t}$, is estimated from the four-factor model. Any $\widehat{\beta}_{\text{DLT},i,t}$ with an incorrect sign—negative for a call or positive for a put—is set to zero. The R square of the four-factor model is defined as

$$R^2\% = \left[1 - \frac{\sum_{i,t}\left(r_{i,t+1} - \left(\widehat{c}_{i,t} + \sum_l \widehat{\beta}_{l,i,t} f_{l,t+1}\right)\right)^2}{\sum_{i,t}(r_{i,t+1})^2}\right] \times 100,$$

where $\widehat{c}_{i,t}$ is the factor model's conditional intercept for each option, i, and $\widehat{\beta}_{i,l,t}$ is the estimated option exposure to factor l. The risk factors l include: DLT, GAM, dVAR, JSK, where dVAR is the change in aggregated variance risk factor, defined as dVAR $= CIV_{14^+,t+1} - CIV_{14^+,t}$ The $R^2\%$ are reported for aggregates 0, 1–3, 4–15 DTEs and moneyness groups. The moneyness is defined as $\kappa_{t,\tau} = \frac{ln(K/F_{t,\tau})}{\text{CIV}_{t,\tau}\sqrt{\tau}}$, where $\text{CIV}_{t,\tau}$ is the corridor implied volatility with maturity τ. We divide options into six moneyness groups: DOTM puts (calls) for $\kappa_{t,\tau} \leq -2.2 (\kappa_{t,\tau} > 1)$, OTM puts (calls) for $-2.2 < \kappa_{t,\tau} \leq -1 (0.5 < \kappa_{t,\tau} \leq 1)$ and ATM puts (calls) for $-1 < \kappa_{t,\tau} \leq 1 (-0.5 < \kappa_{t,\tau} \leq 0.5)$. The sample period is from 2016/09 to 2023/12.

		Puts		Calls	
Moneyness	DTE	DLT	4F	DLT	4F
All	All	94.4	98.8	89.0	97.7
	0	93.1	97.8	89.2	96.6
	1–3	95.0	98.4	91.0	97.2
	4–15	94.4	98.9	88.6	97.8
ATM	All	95.4	99.0	92.4	98.4
	0	94.2	98.3	92.0	97.8
	1–3	96.0	98.7	93.7	98.0
	4–15	95.3	99.0	92.2	98.5
OTM	All	65.9	95.7	55.0	92.0
	0	58.6	85.6	74.0	92.0
	1–3	65.4	91.8	69.9	92.1
	4–15	66.2	96.5	51.3	91.9
DOTM	All	18.3	68.1	9.0	69.8
	0	18.7	28.2	47.8	71.3
	1–3	14.6	57.6	28.4	75.2
	4–15	18.9	70.6	−6.8	66.9

Table IA.VI
Option Dislocations and Intermediary Frictions: Instrumental Variable Approach

This table repeats the analysis of Table V using an instrumental variable (IV) approach to address potential endogeneity concerns arising from the fact that time variation in frictions in the SPX option market and ES futures market (the primary hedging instrument) may be correlated. We instrument *ESF*, the spread between the forward interest rate implied by the S&P 500 E-Mini futures and the three-month OIS rate, using the corresponding spread for the Russell 2000 E-Mini futures. Since the latter are only available from 2017/07, the sample period is from 2017/07 to 2023/12.

	$\widehat{\alpha}_{b(i),t}$											
	0DTE				1–3DTE				4–15DTE			
	Puts		Calls		Puts		Calls		Puts		Calls	
	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$
ESF	−0.06	0.05	−0.05	−0.47	−0.09*	0.37	−0.13***	2.46	−0.01	−0.09	−0.02	2.35
	(0.06)		(0.05)		(0.05)		(0.04)		(0.05)		(0.06)	
HKM	−0.07	0.05	−0.08*	0.03	0.12***	10.23	0.22***	18.95	0.03	9.71	0.12***	18.98
	(0.05)		(0.04)		(0.04)		(0.03)		(0.04)		(0.04)	
ε_{HKM}	0.05	1.91	0.05*	0.92	0.06**	1.03	0.01	−0.06	0.03	0.92	−0.07*	0.05
	(0.03)		(0.03)		(0.03)		(0.02)		(0.02)		(0.04)	
CIV	−0.53***	34.57	−0.62***	39.75	−0.66***	54.31	−0.72***	63.49	−0.69***	53.62	−0.76***	65.10
	(0.10)		(0.08)		(0.10)		(0.05)		(0.07)		(0.04)	
OI	−0.12***	14.53	−0.08***	1.34	−0.11***	12.05	−0.03*	−0.06	−0.07*	12.70	−0.01	6.02
	(0.03)		(0.02)		(0.03)		(0.02)		(0.04)		(0.03)	
SP500	0.07**	3.62	0.07**	2.68	0.10***	1.55	0.07**	0.18	0.11***	1.18	0.06	−0.06
	(0.03)		(0.03)		(0.02)		(0.03)		(0.03)		(0.05)	
Const.	−0.39		−0.34		−0.33		−0.95***		0.11		−0.34	
	(0.30)		(0.25)		(0.27)		(0.19)		(0.28)		(0.23)	
R^2	37.69		41.62		58.18		68.90		55.46		66.84	
	$\lvert\widehat{\alpha}_{b(i),t}\rvert$											
	0DTE				1–3DTE				4–15DTE			
	Puts		Calls		Puts		Calls		Puts		Calls	
	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$
ESF	0.09*	0.50	0.07	−0.14	0.03	0.31	0.04	1.29	−0.07	0.26	−0.07	0.68
	(0.05)		(0.04)		(0.06)		(0.05)		(0.05)		(0.06)	
HKM	0.05	0.45	0.05	0.38	0.08**	2.17	0.01	5.58	0.12***	4.25	0.10**	5.87
	(0.05)		(0.04)		(0.04)		(0.03)		(0.04)		(0.04)	
ε_{HKM}	−0.04	2.00	−0.05*	0.97	−0.05	0.40	0.02	−0.03	−0.02	0.71	0.08*	0.05
	(0.04)		(0.03)		(0.03)		(0.02)		(0.02)		(0.05)	
CIV	0.58***	41.02	0.67***	46.88	0.77***	55.11	0.82***	66.51	0.79***	53.24	0.87***	63.01
	(0.11)		(0.08)		(0.10)		(0.03)		(0.07)		(0.04)	
OI	0.13***	16.74	0.08***	1.41	0.00	6.85	0.02	−0.06	0.00	10.60	0.06**	1.73
	(0.03)		(0.02)		(0.03)		(0.02)		(0.04)		(0.03)	
SP500	−0.08***	4.20	−0.08***	3.09	−0.07***	0.74	−0.04	0.04	−0.09***	0.87	−0.05	−0.06
	(0.03)		(0.03)		(0.03)		(0.03)		(0.03)		(0.05)	
Const.	0.71**		0.65**		−0.22		0.12		−0.50*		−0.65***	
	(0.32)		(0.28)		(0.25)		(0.19)		(0.26)		(0.23)	
R^2	44.40		48.83		56.82		67.22		55.41		63.69	

Table IA.VII
OTM Option Dislocations and Intermediary Frictions

The table reports the coefficients and R^2 of the regressions of OTM option pricing errors measured from returns between 15:15 ET and 15:45 ET on variables related to frictions faced by financial intermediaries and on controls. The top panel reports the results for alphas, $\widehat{\alpha}_{b(i),t}$, while the bottom panel reports the results for absolute alphas, $|\widehat{\alpha}_{b(i),t}|$. The alphas are estimated as bucket-level fixed effects in cross-sectional regressions of the conditional intercepts on option exposures (as in equation (14)), at each 30-minute timestamp. *SP500* is the intraday return of the S&P 500 ES futures from 9:45 ET to 15:15 ET on the current day. *HKM* is the level of the aggregate market equity ratio of primary dealers at the end of the previous day as defined by He, Kelly, and Manela (2017). ε_{HKM} proposed therein is the AR(1) innovation to the capital ratio, scaled by the lagged ratio, on the current day. *ESF* is the spread between the forward interest rate implied by the S&P 500 E-Mini futures and the three-month OIS rate, as in Siriwardane, Sunderam, and Wallen (2025) and Internet Appendix G, estimated using futures prices at the futures closing time, 16:15 ET, of the previous day. *OI* is the open interest of all the puts (calls) within each bucket at the end of the previous day. CIV is the corridor implied variance calculated from options with the given maturity at 15:15 ET on the current day. For each combination of option maturity and type, the table reports the estimated coefficients, their robust standard errors (in parentheses) and adjusted R^2 (in %) in the first column, and the $R^2_{\text{univ.}}$ (in %) from a univariate regression in the second column. All the variables are normalized using their full-sample standard deviations. The sample period is from 2016/09 to 2023/12.

	$\widehat{\alpha}_{b(i),t}$											
	0DTE				1–3DTE				4–15DTE			
	Puts		Calls		Puts		Calls		Puts		Calls	
	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$
ESF	−0.04	0.69	−0.08**	1.10	0.06	−0.05	−0.08***	1.96	0.07**	0.02	−0.06*	2.42
	(0.04)		(0.04)		(0.04)		(0.02)		(0.03)		(0.03)	
HKM	−0.02	1.29	−0.02	0.94	−0.08*	1.96	0.17***	17.38	0.00	8.23	0.13***	22.11
	(0.05)		(0.04)		(0.04)		(0.03)		(0.04)		(0.03)	
ε_{HKM}	0.03	1.57	0.11***	2.00	0.01	−0.03	−0.02	0.02	0.00	0.35	−0.07*	0.01
	(0.03)		(0.04)		(0.03)		(0.01)		(0.02)		(0.04)	
CIV	−0.65***	50.36	−0.67***	48.48	−0.69***	42.11	−0.78***	70.79	−0.70***	47.17	−0.77***	68.85
	(0.12)		(0.10)		(0.04)		(0.03)		(0.04)		(0.04)	
OI	−0.17***	13.97	−0.09***	3.44	0.03	0.55	−0.04**	1.36	−0.04	0.08	0.02	0.23
	(0.03)		(0.02)		(0.03)		(0.02)		(0.03)		(0.02)	
SP500	0.08***	3.24	0.04*	3.07	0.04*	0.14	0.03	−0.03	0.09***	0.64	0.07	−0.03
	(0.03)		(0.03)		(0.02)		(0.02)		(0.02)		(0.05)	
Const.	−0.96***		−0.70***		1.02***		−0.47**		0.29		−0.32	
	(0.34)		(0.27)		(0.29)		(0.19)		(0.30)		(0.22)	
R^2	53.92		51.02		43.17		73.98		48.49		71.12	

	$\lvert\widehat{\alpha}_{b(i),t}\rvert$											
	0DTE				1–3DTE				4–15DTE			
	Puts		Calls		Puts		Calls		Puts		Calls	
	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$
ESF	0.05	0.77	0.09**	1.34	0.05*	1.03	0.04	1.26	−0.01	0.94	0.03	2.12
	(0.04)		(0.04)		(0.03)		(0.03)		(0.03)		(0.04)	
HKM	0.01	1.64	0.01	1.28	0.14***	0.78	0.14***	1.78	0.23***	1.83	0.18***	3.52
	(0.05)		(0.04)		(0.03)		(0.03)		(0.03)		(0.04)	
ε_{HKM}	−0.02	1.37	−0.10**	1.89	−0.02	−0.05	0.03**	0.05	0.01	0.08	0.09*	0.05
	(0.03)		(0.04)		(0.03)		(0.02)		(0.02)		(0.05)	
CIV	0.67***	54.33	0.70***	52.75	0.73***	45.77	0.86***	67.13	0.81***	49.91	0.85***	60.03
	(0.12)		(0.10)		(0.03)		(0.04)		(0.04)		(0.06)	
OI	0.17***	14.93	0.08***	3.48	−0.06**	0.34	−0.02	0.25	−0.05**	−0.04	0.01	0.42
	(0.03)		(0.02)		(0.02)		(0.02)		(0.02)		(0.02)	
SP500	−0.08***	3.30	−0.05*	3.23	0.02	−0.04	−0.01	−0.05	−0.05	0.07	−0.06	−0.03
	(0.03)		(0.03)		(0.02)		(0.03)		(0.03)		(0.05)	
Const.	1.09***		0.85***		−0.35*		−0.73***		−1.09***		−1.16***	
	(0.34)		(0.28)		(0.20)		(0.22)		(0.23)		(0.29)	
R^2	57.94		55.27		48.01		69.23		53.79		63.46	

Table IA.VIII
DOTM Option Dislocations and Intermediary Frictions

The table reports the coefficients and R^2 of the regressions of DOTM option pricing errors measured from returns between 15:15 ET and 15:45 ET on variables related to frictions faced by financial intermediaries and on controls. The top panel reports the results for alphas, $\widehat{\alpha}_{b(i),t}$, while the bottom panel reports the results for absolute alphas, $|\widehat{\alpha}_{b(i),t}|$. The alphas are estimated as bucket-level fixed effects in cross-sectional regressions of the conditional intercepts on option exposures (as in equation (14)), at each 30-minute timestamp. *SP500* is the intraday return of the S&P 500 ES futures from 9:45 ET to 15:15 ET on the current day. *HKM* is the level of the aggregate market equity ratio of primary dealers at the end of the previous day as defined by He, Kelly, and Manela (2017). ε_{HKM} proposed therein is the AR(1) innovation to the capital ratio, scaled by the lagged ratio, on the current day. *ESF* is the spread between the forward interest rate implied by the S&P 500 E-Mini futures and the three-month OIS rate, as in Siriwardane, Sunderam, and Wallen (2025) and Internet Appendix G, estimated using futures prices at the futures closing time, 16:15 ET, of the previous day. *OI* is the open interest of all the puts (calls) within each bucket at the end of the previous day. CIV is the corridor implied variance calculated from options with the given maturity at 15:15 ET on the current day. For each combination of option maturity and type, the table reports the estimated coefficients, their robust standard errors (in parentheses) and adjusted R^2 (in %) in the first column, and the $R^2_{\text{univ.}}$ (in %) from a univariate regression in the second column. All the variables are normalized using their full-sample standard deviations. The sample period is from 2016/09 to 2023/12.

	$\widehat{\alpha}_{b(i),t}$											
	0DTE				1–3DTE				4–15DTE			
	Puts		Calls		Puts		Calls		Puts		Calls	
	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$
ESF	−0.05	0.50	−0.10***	2.24	0.06	−0.04	−0.09***	2.35	0.05*	0.26	−0.08***	3.95
	(0.04)		(0.03)		(0.04)		(0.02)		(0.03)		(0.03)	
HKM	0.11**	4.99	0.20***	11.75	−0.01	3.87	0.20***	21.77	0.01	11.63	0.24***	29.05
	(0.05)		(0.05)		(0.04)		(0.02)		(0.04)		(0.03)	
ε_{HKM}	0.09**	0.18	0.12***	2.83	0.01	−0.06	0.00	−0.05	−0.01	0.34	−0.02	0.10
	(0.04)		(0.04)		(0.03)		(0.01)		(0.01)		(0.01)	
CIV	−0.71***	56.49	−0.65***	52.58	−0.71***	48.88	−0.81***	79.10	−0.80***	62.64	−0.73***	70.34
	(0.16)		(0.14)		(0.04)		(0.05)		(0.04)		(0.04)	
OI	−0.13***	4.55	−0.07***	1.51	−0.08***	0.51	−0.02	0.01	−0.02	3.75	−0.12***	0.19
	(0.02)		(0.02)		(0.02)		(0.01)		(0.03)		(0.02)	
SP500	0.04*	2.02	0.07***	4.20	0.04*	0.12	0.05**	0.03	0.11***	0.81	0.04*	−0.05
	(0.02)		(0.02)		(0.02)		(0.02)		(0.03)		(0.03)	
Const.	−2.45***		−1.84***		1.00***		−0.45***		0.38		−0.76***	
	(0.38)		(0.33)		(0.30)		(0.17)		(0.28)		(0.22)	
R^2	60.16		60.19		49.97		83.58		63.94		75.80	

	$\lvert\widehat{\alpha}_{b(i),t}\rvert$											
	0DTE				1–3DTE				4–15DTE			
	Puts		Calls		Puts		Calls		Puts		Calls	
	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$
ESF	0.05	0.50	0.11***	2.47	0.03	0.64	−0.01	0.35	−0.01	1.19	0.04	3.42
	(0.04)		(0.03)		(0.03)		(0.04)		(0.03)		(0.04)	
HKM	−0.11**	4.99	−0.18***	10.40	0.24***	−0.05	0.26***	0.03	0.23***	0.93	0.13***	5.32
	(0.05)		(0.04)		(0.04)		(0.04)		(0.04)		(0.04)	
ε_{HKM}	−0.09**	0.18	−0.12***	2.88	−0.01	0.10	0.05***	0.24	0.04***	0.00	0.04**	0.04
	(0.04)		(0.04)		(0.03)		(0.02)		(0.01)		(0.02)	
CIV	0.71***	56.49	0.68***	55.10	0.70***	39.47	0.87***	61.13	0.88***	56.49	0.83***	60.13
	(0.16)		(0.14)		(0.03)		(0.05)		(0.03)		(0.05)	
OI	0.13***	4.55	0.06***	1.31	−0.09***	0.15	0.00	0.01	0.09***	−0.04	0.06**	0.22
	(0.02)		(0.02)		(0.02)		(0.02)		(0.03)		(0.03)	
SP500	−0.04*	2.02	−0.07***	4.28	0.05**	0.19	−0.01	−0.05	−0.05	0.07	−0.03	−0.05
	(0.02)		(0.02)		(0.02)		(0.02)		(0.03)		(0.03)	
Const.	2.45***		1.75***		−0.74***		−1.15***		−1.49***		−0.80***	
	(0.38)		(0.31)		(0.27)		(0.26)		(0.22)		(0.28)	
R^2	60.16		61.82		44.74		67.18		63.56		62.52	

Table IA.IX
ATM Option Dislocations and Intraday ES Futures Returns Decomposition

The table reports the coefficients and R^2 of the regressions of ATM option pricing errors measured from returns between 15:15 ET and 15:45 ET on the decomposition of intraday ES futures return at daily frequency across option types and maturity buckets. The top panel reports the results for alphas, $\alpha_{b(i),t}$, while the bottom panel reports the results for absolute alphas, $|\alpha_{b(i),t}|$. The alphas are estimated as bucket-level fixed effects in cross-sectional regressions of the conditional intercepts on option exposures (as in equation (14)), at each 30-minute timestamp. $SP500_{t,[09:45,14:45]}$ and $SP500_{t,[14:45,15:15]}$ denote the intraday returns of the S&P 500 ES futures from 9:45 ET to 14:45 ET and from 14:45 ET to 15:15 ET, respectively, on the current day. For each combination of option maturity and type, the table reports the estimated coefficients, their robust standard errors (in parentheses) and adjusted R^2 (in %) in the first column, and the $R^2_{\text{univ.}}$ (in %) from a univariate regression in the second column. All the variables are normalized using their full-sample standard deviations. The sample period is from 2016/09 to 2023/12.

	$\widehat{\alpha}_{b(i),t}$											
	0DTE				1–3DTE				4–15DTE			
	Puts		Calls		Puts		Calls		Puts		Calls	
	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$
$SP500_{[09:45,14:45]}$	0.15***	2.35	0.13***	1.64	0.10	0.90	0.03	2.93	0.10	0.83	−0.01	−0.04
	(0.04)		(0.04)		(0.06)		(0.06)		(0.06)		(0.06)	
$SP500_{[14:45,15:15]}$	0.12**	1.50	0.12**	1.32	0.09	0.63	0.07	0.42	0.07	0.45	0.06	0.31
	(0.06)		(0.05)		(0.06)		(0.06)		(0.05)		(0.06)	
Const.	−1.29***		−1.24***		−0.10*		−0.15**		−0.25***		−0.15**	
	(0.06)		(0.06)		(0.06)		(0.07)		(0.06)		(0.07)	
R^2	3.78		2.89		1.58		0.43		1.32		0.27	
	$\lvert\widehat{\alpha}_{b(i),t}\rvert$											
	0DTE				1–3DTE				4–15DTE			
	Puts		Calls		Puts		Calls		Puts		Calls	
	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$
$SP500_{[09:45,14:45]}$	−0.17***	2.76	−0.14***	1.97	−0.06	0.27	−0.01	−0.04	−0.08	0.50	0.01	−0.04
	(0.04)		(0.03)		(0.07)		(0.06)		(0.06)		(0.06)	
$SP500_{[14:45,15:15]}$	−0.13**	1.61	−0.12*	1.32	−0.09	0.74	−0.07	0.40	−0.07	0.47	−0.06	0.33
	(0.06)		(0.06)		(0.07)		(0.07)		(0.06)		(0.07)	
Const.	1.50***		1.44***		0.76***		0.72***		0.73***		0.61***	
	(0.06)		(0.06)		(0.06)		(0.07)		(0.06)		(0.07)	
R^2	4.29		3.21		1.04		0.36		1.01		0.28	

Table IA.X
OTM Option Dislocations and Intraday ES Futures Returns Decomposition

The table reports the coefficients and R^2 of the regressions of OTM option pricing errors measured from returns between 15:15 ET and 15:45 ET on the decomposition of intraday ES futures return at daily frequency across option types and maturity buckets. The top panel reports the results for alphas, $\alpha_{b(i),t}$, while the bottom panel reports the results for absolute alphas, $|\alpha_{b(i),t}|$. The alphas are estimated as bucket-level fixed effects in cross-sectional regressions of the conditional intercepts on option exposures (as in equation (14)), at each 30-minute timestamp. $SP500_{t,[09:45,14:45]}$ and $SP500_{t,[14:45,15:15]}$ denote the intraday returns of the S&P 500 ES futures from 9:45 ET to 14:45 ET and from 14:45 ET to 15:15 ET, respectively, on the current day. For each combination of option maturity and type, the table reports the estimated coefficients, their robust standard errors (in parentheses) and adjusted R^2 (in %) in the first column, and the $R^2_{\text{univ.}}$ (in %) from a univariate regression in the second column. All the variables are normalized using their full-sample standard deviations. The sample period is from 2016/09 to 2023/12.

	$\widehat{\alpha}_{b(i),t}$											
	0DTE				1–3DTE				4–15DTE			
	Puts		Calls		Puts		Calls		Puts		Calls	
	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$
$SP500_{[09:45,14:45]}$	0.15***	2.15	0.14***	2.04	0.01	−0.04	−0.01	−0.05	0.06	0.29	0.00	−0.05
	(0.04)		(0.04)		(0.07)		(0.07)		(0.06)		(0.06)	
$SP500_{[14:45,15:15]}$	0.13**	1.69	0.12**	1.53	0.11	1.05	0.07	0.51	0.09	0.66	0.07	0.38
	(0.06)		(0.06)		(0.08)		(0.07)		(0.06)		(0.06)	
Const.	−1.64***		−1.34***		0.37***		0.03		−0.16**		−0.08	
	(0.07)		(0.07)		(0.06)		(0.07)		(0.06)		(0.07)	
R^2	3.75		3.49		1.01		0.46		0.98		0.33	
	$\lvert\widehat{\alpha}_{b(i),t}\rvert$											
	0DTE				1–3DTE				4–15DTE			
	Puts		Calls		Puts		Calls		Puts		Calls	
	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$
$SP500_{[09:45,14:45]}$	−0.15***	2.20	−0.15***	2.20	0.04	0.13	0.02	−0.01	−0.01	−0.04	0.01	−0.05
	(0.04)		(0.04)		(0.07)		(0.07)		(0.07)		(0.07)	
$SP500_{[14:45,15:15]}$	−0.13**	1.68	−0.12*	1.50	−0.09	0.79	−0.07	0.47	−0.08	0.54	−0.07	0.42
	(0.06)		(0.07)		(0.08)		(0.08)		(0.07)		(0.07)	
Const.	1.74***		1.45***		0.89***		0.64***		0.79***		0.65***	
	(0.07)		(0.07)		(0.06)		(0.07)		(0.06)		(0.07)	
R^2	3.80		3.60		0.90		0.45		0.51		0.37	

Table IA.XI
DOTM Option Dislocations and Intraday ES Futures Returns Decomposition

The table reports the coefficients and R^2 of the regressions of DOTM option pricing errors measured from returns between 15:15 ET and 15:45 ET on the decomposition of intraday ES futures return at daily frequency across option types and maturity buckets. The top panel reports the results for alphas, $\alpha_{b(i),t}$, while the bottom panel reports the results for absolute alphas, $|\alpha_{b(i),t}|$. The alphas are estimated as bucket-level fixed effects in cross-sectional regressions of the conditional intercepts on option exposures (as in equation (14)), at each 30-minute timestamp. $SP500_{t,[09:45,14:45]}$ and $SP500_{t,[14:45,15:15]}$ denote the intraday returns of the S&P 500 ES futures from 9:45 ET to 14:45 ET and from 14:45 ET to 15:15 ET, respectively, on the current day. For each combination of option maturity and type, the table reports the estimated coefficients, their robust standard errors (in parentheses) and adjusted R^2 (in %) in the first column, and the $R^2_{\text{univ.}}$ (in %) from a univariate regression in the second column. All the variables are normalized using their full-sample standard deviations. The sample period is from 2016/09 to 2023/12.

	$\widehat{\alpha}_{b(i),t}$											
	0DTE				1–3DTE				4–15DTE			
	Puts		Calls		Puts		Calls		Puts		Calls	
	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$
$SP500_{[09:45,14:45]}$	0.13***	1.43	0.18***	3.18	0.02	−0.02	0.01	−0.05	0.07	0.40	−0.01	−0.04
	(0.04)		(0.04)		(0.09)		(0.07)		(0.07)		(0.07)	
$SP500_{[14:45,15:15]}$	0.09	0.66	0.11**	1.31	0.07	0.47	0.07	0.45	0.09	0.71	0.06	0.31
	(0.06)		(0.06)		(0.06)		(0.07)		(0.06)		(0.07)	
Const.	−2.30***		−1.06***		0.62***		0.28***		−0.05		−0.09	
	(0.07)		(0.07)		(0.06)		(0.07)		(0.07)		(0.08)	
R^2	2.15		4.40		0.46		0.41		1.14		0.26	
	$\lvert\widehat{\alpha}_{b(i),t}\rvert$											
	0DTE				1–3DTE				4–15DTE			
	Puts		Calls		Puts		Calls		Puts		Calls	
	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$
$SP500_{[09:45,14:45]}$	−0.13***	1.43	−0.18***	3.23	0.08	0.65	0.03	0.07	−0.01	−0.05	0.02	0.02
	(0.04)		(0.04)		(0.07)		(0.08)		(0.07)		(0.08)	
$SP500_{[14:45,15:15]}$	−0.09	0.66	−0.12**	1.37	−0.09	0.83	−0.08	0.60	−0.09	0.71	−0.07	0.40
	(0.06)		(0.06)		(0.07)		(0.07)		(0.07)		(0.08)	
Const.	2.30***		1.14***		1.05***		0.89***		0.73***		0.81***	
	(0.07)		(0.07)		(0.06)		(0.07)		(0.07)		(0.07)	
R^2	2.15		4.51		1.45		0.65		0.67		0.41	

Table IA.XII
OTM Option Dislocations and Intermediary Frictions: Instrumental Variable Approach

This table repeats the analysis of Table IA.VII using an instrumental variable (IV) approach to address potential endogeneity concerns arising from the fact that time variation in frictions in the SPX option market and ES futures market (the primary hedging instrument) may be correlated. We instrument *ESF*, the spread between the forward interest rate implied by the S&P 500 E-Mini futures and the three-month OIS rate, using the corresponding spread for the Russell 2000 E-Mini futures. Since the latter are only available from 2017/07, the sample period is from 2017/07 to 2023/12.

	$\widehat{\alpha}_{b(i),t}$											
	0DTE				1–3DTE				4–15DTE			
	Puts		Calls		Puts		Calls		Puts		Calls	
	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$
ESF	−0.10**	0.55	−0.10**	0.34	0.06	−0.39	−0.08**	1.97	0.04	−0.17	−0.03	2.40
	(0.05)		(0.05)		(0.06)		(0.04)		(0.05)		(0.06)	
HKM	−0.04	0.84	−0.03	0.79	−0.10**	1.53	0.17***	16.78	−0.04	6.87	0.13***	21.26
	(0.05)		(0.04)		(0.05)		(0.03)		(0.05)		(0.04)	
ε_{HKM}	0.03	1.55	0.10***	2.04	0.02	−0.02	−0.02	0.02	0.00	0.30	−0.07*	0.04
	(0.03)		(0.04)		(0.03)		(0.01)		(0.02)		(0.04)	
CIV	−0.65***	51.06	−0.67***	48.59	−0.72***	44.88	−0.79***	72.16	−0.72***	48.49	−0.77***	68.60
	(0.11)		(0.10)		(0.04)		(0.02)		(0.04)		(0.04)	
OI	−0.16***	13.54	−0.08***	3.37	0.06***	0.17	−0.04*	1.11	−0.01	−0.03	0.02	0.14
	(0.03)		(0.02)		(0.02)		(0.02)		(0.03)		(0.02)	
SP500	0.09***	3.39	0.05*	3.22	0.03	0.13	0.04	−0.03	0.09***	0.56	0.07	−0.04
	(0.03)		(0.03)		(0.02)		(0.03)		(0.03)		(0.05)	
Const.	−0.79**		−0.62**		1.09***		−0.49***		0.46		−0.37*	
	(0.33)		(0.27)		(0.30)		(0.18)		(0.30)		(0.21)	
R^2	54.20		51.08		46.65		75.33		49.74		70.90	

	$\lvert\widehat{\alpha}_{b(i),t}\rvert$											
	0DTE				1–3DTE				4–15DTE			
	Puts		Calls		Puts		Calls		Puts		Calls	
	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$
ESF	0.10**	0.56	0.10**	0.35	−0.03	−1.30	−0.04	−1.66	−0.10**	−0.38	−0.10	−0.26
	(0.05)		(0.05)		(0.04)		(0.05)		(0.04)		(0.07)	
HKM	0.03	1.08	0.02	1.10	0.12***	1.34	0.13***	2.10	0.22***	2.31	0.18***	3.75
	(0.05)		(0.04)		(0.04)		(0.03)		(0.04)		(0.04)	
ε_{HKM}	−0.02	1.43	−0.10**	1.94	−0.02	−0.06	0.03**	0.09	0.01	0.11	0.09*	0.06
	(0.03)		(0.04)		(0.03)		(0.02)		(0.02)		(0.05)	
CIV	0.68***	54.89	0.70***	52.97	0.76***	51.09	0.88***	70.27	0.84***	53.55	0.88***	62.06
	(0.12)		(0.10)		(0.03)		(0.04)		(0.05)		(0.05)	
OI	0.16***	14.54	0.08***	3.41	−0.06**	0.36	0.00	0.39	−0.04*	−0.02	0.01	0.60
	(0.03)		(0.02)		(0.02)		(0.02)		(0.02)		(0.02)	
SP500	−0.09***	3.46	−0.05*	3.40	0.02	−0.05	−0.01	−0.06	−0.05	0.08	−0.06	−0.04
	(0.03)		(0.03)		(0.02)		(0.03)		(0.03)		(0.05)	
Const.	0.91***		0.79***		−0.14		−0.56***		−0.91***		−0.95***	
	(0.34)		(0.28)		(0.21)		(0.20)		(0.24)		(0.24)	
R^2	58.14		55.48		52.20		71.25		56.04		63.45	

Table IA.XIII
DOTM Option Dislocations and Intermediary Frictions: Instrumental Variable Approach

This table repeats the analysis of Table IA.VIII using an instrumental variable (IV) approach to address potential endogeneity concerns arising from the fact that time variation in frictions in the SPX option market and ES futures market (the primary hedging instrument) may be correlated. We instrument *ESF*, the spread between the forward interest rate implied by the S&P 500 E-Mini futures and the three-month OIS rate, using the corresponding spread for the Russell 2000 E-Mini futures. Since the latter are only available from 2017/07, the sample period is from 2017/07 to 2023/12.

	$\widehat{\alpha}_{b(i),t}$											
	0DTE				1–3DTE				4–15DTE			
	Puts		Calls		Puts		Calls		Puts		Calls	
	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$
ESF	−0.11**	−0.89	−0.18***	2.17	0.04	−0.05	−0.13***	2.36	0.00	−0.15	−0.09*	3.57
	(0.05)		(0.05)		(0.06)		(0.03)		(0.04)		(0.05)	
HKM	0.11**	4.23	0.18***	9.90	−0.03	2.93	0.20***	20.92	−0.01	9.98	0.24***	28.56
	(0.05)		(0.05)		(0.05)		(0.02)		(0.04)		(0.03)	
ε_{HKM}	0.09**	0.14	0.12***	2.92	0.01	−0.06	0.00	−0.05	−0.01	0.29	−0.02*	0.05
	(0.04)		(0.04)		(0.04)		(0.01)		(0.01)		(0.01)	
CIV	−0.72***	57.82	−0.66***	53.65	−0.72***	50.32	−0.82***	79.83	−0.81***	63.75	−0.72***	70.19
	(0.16)		(0.13)		(0.04)		(0.05)		(0.04)		(0.04)	
OI	−0.13***	4.34	−0.06***	1.08	−0.07***	0.47	−0.01	0.04	−0.02	3.27	−0.13***	0.94
	(0.02)		(0.02)		(0.02)		(0.01)		(0.03)		(0.02)	
SP500	0.04*	2.06	0.08***	4.36	0.04	0.11	0.05**	0.04	0.11***	0.73	0.04	−0.06
	(0.03)		(0.02)		(0.03)		(0.02)		(0.03)		(0.03)	
Const.	−2.27***		−1.61***		1.14***		−0.43***		0.55**		−0.79***	
	(0.37)		(0.33)		(0.29)		(0.16)		(0.26)		(0.21)	
R^2	60.88		59.82		51.51		84.25		64.83		76.38	
	$\lvert\widehat{\alpha}_{b(i),t}\rvert$											
	0DTE				1–3DTE				4–15DTE			
	Puts		Calls		Puts		Calls		Puts		Calls	
	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$	Coef.	$R^2_{\text{univ.}}$
ESF	0.11**	−0.89	0.18***	2.52	−0.07	−2.88	−0.11**	−4.09	−0.13***	−1.38	−0.09	1.68
	(0.05)		(0.05)		(0.05)		(0.05)		(0.04)		(0.07)	
HKM	−0.11**	4.23	−0.16***	8.76	0.22***	0.08	0.24***	0.16	0.21***	1.44	0.13***	5.65
	(0.05)		(0.04)		(0.05)		(0.04)		(0.04)		(0.05)	
ε_{HKM}	−0.09**	0.14	−0.12***	2.99	−0.01	0.14	0.05***	0.26	0.03**	0.01	0.04**	0.02
	(0.04)		(0.04)		(0.04)		(0.02)		(0.01)		(0.02)	
CIV	0.72***	57.82	0.68***	55.82	0.74***	44.52	0.89***	65.06	0.92***	61.03	0.87***	62.59
	(0.16)		(0.13)		(0.03)		(0.05)		(0.04)		(0.04)	
OI	0.13***	4.34	0.05***	0.95	−0.09***	0.22	0.00	0.01	0.10***	−0.03	0.07**	0.51
	(0.02)		(0.02)		(0.02)		(0.02)		(0.03)		(0.03)	
SP500	−0.04*	2.06	−0.07***	4.42	0.05**	0.17	0.00	−0.06	−0.05	0.10	−0.03	−0.06
	(0.03)		(0.02)		(0.02)		(0.02)		(0.03)		(0.03)	
Const.	2.27***		1.55***		−0.45*		−0.93***		−1.20***		−0.64**	
	(0.37)		(0.31)		(0.25)		(0.24)		(0.22)		(0.25)	
R^2	60.88		61.57		47.48		69.31		65.79		63.07	

Table IA.XIV
The First-stage Results of IV Regressions

This table reports the first-stage regression results for the instrumental variable (IV) approach used in Tables IA.VI, IA.XII, and IA.XIII. Since the puts and calls with different maturities are not available on the exactly same dates, we report the first-stage regression results separately for each option type and maturity group. We instrument *ESF*, the spread between the forward interest rate implied by the S&P 500 E-Mini futures and the three-month OIS rate, using the corresponding spread for the Russell 2000 E-Mini futures. Since the latter are only available from 2017/07, the sample period is from 2017/07 to 2023/12.

	0DTE		1–3DTE		4–15DTE	
	R^2	*Fstats.*	R^2	*Fstats.*	R^2	*Fstats.*
P	57.4	298.2	53.1	203.8	53.3	200.6
C	58.3	310.4	53.2	209.8	53.2	211.5

REFERENCES

Andersen, Torben G., Oleg Bondarenko, and Maria T Gonzalez-Perez, 2015, Exploring Return Dynamics via Corridor Implied Volatility, *Review of Financial Studies* 28, 2902–2945.

Andersen, Torben G., Nicola Fusari, and Viktor Todorov, 2017, Short-Term Market Risks Implied by Weekly Options, *The Journal of Finance* 72, 1335–1386.

Bondarenko, Oleg, 2026, The Clock Matters: Time Measurement and the Ultra-Short Volatility Term Structure, *Available at SSRN 6705119* .

Broadie, Mark, Mikhail Chernov, and Michael Johannes, 2009, Understanding Index Option Returns, *Review of Financial Studies* 22, 4493–4529.

Carr, Peter, and Liuren Wu, 2020, Option Profit and Loss Attribution and Pricing: A New Framework, *Journal of Finance* 75, 2271–2316.

He, Zhiguo, Bryan Kelly, and Asaf Manela, 2017, Intermediary asset pricing: New evidence from many asset classes, *Journal of Financial Economics* 126, 1–35.

Jiang, George J, and Yisong S Tian, 2005, The Model-Free Implied Volatility and Its Information Content, *Review of Financial Studies* 18, 1305–1342.

Siriwardane, Emil, Aditya Sunderam, and Jonathan Wallen, 2025, Segmented Arbitrage, *The Journal of Finance* 80, 2543–2590.

Todorov, Viktor, 2019, Nonparametric spot volatility from options, *The Annals of Applied Probability* 29, 3590–3636.