

0DTE Asset Pricing*

Caio Almeida^{†a}, Gustavo Freire^{‡b}, and Rodrigo Hizmeri^{§c}

^aPrinceton University

^bErasmus University Rotterdam, Tinbergen Institute & ERIM

^cUniversity of Liverpool

This draft: October 18, 2024

First draft: January 20, 2024

Abstract

We document asset pricing implications of the new zero days-to-expiration (0DTE) options, which today account for half of the total S&P 500 option volume. We show that: (i) most of the intra-day equity premium is attributable to market returns between -5% and 0%; (ii) investors demand a high compensation to bear variance risk over the day, which is mainly due to compensation for upside risk; (iii) the variance risk premium predicts intra-day market returns, with a negative relation that is driven by the upside risk premium; and (iv) 0DTE options violate stochastic dominance restrictions, where exploiting this relative mispricing is highly profitable. Our findings contrast with evidence from longer horizons and are consistent with a nonmonotonic pricing kernel that is especially high for positive market returns.

Keywords: zero days-to-expiration (0DTE) options, equity premium, variance risk premium, pricing kernel, option returns.

JEL Classification: G11, G12, G13.

*We would like to thank Oleg Bondarenko (discussant), Stefanos Delikouras (discussant), Elise Gourier (discussant), Hyung Joo Kim, Alex Kostakis, Tobias Sichert and participants at the Cancun Derivatives and Asset Pricing Conference 2024, the TSE Financial Econometrics Conference 2024, the 16th Annual SoFiE Meeting, the Liverpool Workshop in Option Markets, the 24th Brazilian Finance Meeting and the “Frontiers of Option Pricing” session at the Econometric Society Summer Meeting 2024 for useful comments and suggestions.

[†]calmeida@princeton.edu, Department of Economics, Princeton University.

[‡]freire@ese.eur.nl, Erasmus School of Economics - Erasmus University Rotterdam, Tinbergen Institute and Erasmus Research Institute of Management (ERIM).

[§]r.hizmeri@liverpool.ac.uk, University of Liverpool Management School, University of Liverpool.

1 Introduction

We investigate the asset pricing implications of a new, relatively unexplored market: zero days-to-expiration (0DTE) S&P 500 options, i.e., options on the market index expiring by the end of the same day. These are weekly options that were listed at least a week before. Since May 2022, weeklies are listed every trading day by the Chicago Board Options Exchange (CBOE), resulting in the daily availability of 0DTE options. While one-month options were among the most traded contracts up to ten years ago, which partially justified the focus on these options by the literature, the landscape of the option market has changed dramatically more recently. Today, the daily volume of 0DTEs accounts for around half of total S&P 500 option volume, being thus by far the most traded maturity. The tremendous growth of these ultra short-maturity options has made them a trending topic in financial media outlets, trader forums and social media.

Our interest in 0DTE options is justified not only by the fact that they are now the most traded options in the market, but also because they contain new, valuable information about investors' risk preferences and risk premia over intra-daily horizons. Our goal is to extract this information by exploring, from different perspectives, how 0DTE option prices relate to the time series of high-frequency market returns. Our analysis provides new asset pricing stylized facts for ultra short-horizons and highlights how they are remarkably different from the extant empirical evidence of longer horizons.

We focus on two main questions. First, what is the compensation required by investors for bearing market risks over the day? The answer to this question is directly related to the intra-day pricing kernel implied by 0DTEs and its shape as a function of market returns.¹ We document a number of patterns in the equity and variance risk premia that are aligned with a nonmonotonic pricing kernel that is especially high for positive

¹Starting with [Jackwerth \(2000\)](#), [Aït-Sahalia and Lo \(2000\)](#) and [Rosenberg and Engle \(2002\)](#), a large literature estimates this pricing kernel for longer horizons and documents a nonmonotonic shape, which is puzzling under a representative investor framework. See [Cuesdeanu and Jackwerth \(2018\)](#) for a survey.

returns. Second, are 0DTE option prices consistent with risk-averse investors? We find that the majority of the options violate stochastic dominance bounds, where exploiting these violations is highly profitable. This suggests that 0DTEs are mostly “mispriced”.²

We start by analyzing average returns of 0DTE call and put options across strikes. These are informative about investors’ preferences over intra-day market return states, i.e., about the shape of the pricing kernel projected onto market returns over the options horizon. Calls experience low returns overall, which decrease with the strike and eventually get highly negative. Following the rationale of [Bakshi, Madan, and Panayotov \(2010\)](#), this is evidence that the pricing kernel is increasing for some region of positive market returns, violating monotonicity. Put returns are also negative, especially out-of-the-money (OTM), indicating that the pricing kernel is a decreasing but steep function of market returns in the negative return region, consistent with aversion to downside risk.

Such risk preferences have strong implications for the equity premium over intra-daily horizons. Applying the decomposition of [Beason and Schreindorfer \(2022\)](#) to 0DTE options and realized market returns, we find that most of the intra-day equity premium stems from compensation to market returns between -5% and 0% . In contrast, positive return states have a negative contribution. This happens when the pricing kernel is nonmonotonic and marginal utility is high for positive market returns. In this case, investors would be willing to pay a premium to hold Arrow-Debreu securities paying in those states, i.e., they are willing to give up part of the compensation for equity risk.

Implications are also substantial for the variance risk premium.³ We first document that the average returns of at-the-money (ATM) delta-hedged calls and straddles are significantly negative. Since these strategies essentially represent long positions in volatility,

²We emphasize that “mispriced” here means only that we cannot reconcile 0DTE prices with information from the underlying market under risk-averse preferences. This does not mean that 0DTEs violate no-arbitrage conditions (which imply much wider bounds that 0DTEs in general satisfy) or that it is not possible to price 0DTEs with an option pricing model (see, e.g., [Bandi, Fusari, and Renò, 2023](#)).

³As [Bollerslev, Tauchen, and Zhou \(2009\)](#), we define the variance risk premium as the difference between the risk-neutral and the physical expected variance of the market return.

this means that investors are willing to pay a high premium to be protected against variance risk over the day (Bakshi and Kapadia, 2003; Coval and Shumway, 2001). This is confirmed when we compute a direct estimate of the variance risk premium, similarly to Bollerslev et al. (2009), but for the ultra-short horizons associated with 0DTE options. The (annualized) average variance risk premium over the day can be up to four times larger than what is usually observed for the one-month horizon.

To disentangle the compensation demanded by investors to bear variation risk in positive versus negative market returns, we respectively compute the “good” and “bad” components of the total variance risk premium, as in Kilic and Shaliastovich (2019), for intra-daily horizons. Strikingly, while at the one-month horizon the “good” variance risk premium is negative and the “bad” is highly positive, we find that the intra-day upside risk premium is positive and often larger than the premium for downside risk. This, again, would be consistent with a pricing kernel that is exceptionally high for positive market returns, such that investors demand compensation for these states of high marginal utility.

We further investigate whether 0DTE options contain predictive information for excess market returns over the day. We consider predictive regressions using the variance risk premium, risk-neutral moments (Bakshi, Kapadia, and Madan, 2003) and the equity premium lower bound of Martin (2017). From these variables, only the variance risk premium significantly predicts returns, but with a negative coefficient, which is at odds with the existing evidence for longer horizons (Bollerslev et al., 2009). By substituting the total variance risk premium with its “good” and “bad” components, we find that only the “good” component helps predict market returns, with a strong and statistically significant negative sign. In other words, the compensation for upside risk drives the result for the total variance risk premium. This negative relation is, once more, aligned with high marginal utility in states of positive market returns: the higher the pricing kernel in this region, the higher is the compensation for positive return variation risk and

the more negative is the contribution of positive return states to the equity premium.

The findings above provide indirect evidence that the pricing kernel as a function of market returns is nonmonotonic and, in particular, high for positive returns. To confirm this evidence, we directly estimate the intra-day pricing kernel implied by 0DTE options and high-frequency market returns. We document that, regardless of the time of the day, on average there is a pronounced and statistically significant hump in the pricing kernel around returns close to zero. That is, in the region where returns are more likely to occur, marginal utility is higher for positive returns than for negative returns, which is aligned with our results on option returns and market risk premia. Outside this region, we observe less pronounced, but still statistically significant nonmonotonicities for positive returns, while the pricing kernel is mostly monotonically decreasing for negative returns.

Our evidence shows that 0DTE option prices can only be jointly reconciled with the physical distribution of market returns under a pricing kernel displaying pronounced nonmonotonicities. However, if we entertain the possibility that the 0DTE option market is segmented, there could still be a different risk-averse trader in the market index that is marginal in each option. To test for that, we compute, for each option, price bounds from the physical distribution consistent with all pricing kernels that are monotonically decreasing in market returns ([Ritchken, 1985](#)). Over our sample, only around 35% of 0DTE options (and 7% of the ATM options) satisfy these bounds, which is again in stark contrast with evidence from longer horizons ([Almeida and Freire, 2022](#)). As a result, 0DTE options are mostly “mispriced”, in the sense that they do not reflect the risks implied by the time series of intra-day market returns under risk-averse preferences.

Importantly, the mispricing we document is economically significant. A trading strategy purchasing (writing) the delta-hedged ATM option if it is cheap (expensive) according to risk-averse investors – i.e., if its price represents a lower (upper) bound violation – is highly profitable, even after transaction costs. In particular, the Sharpe ratio of this

strategy is about ten times as large as that obtained by always writing the delta-hedged ATM option, that is, exploiting the variance risk premium. The mispricing is more pronounced, in terms of higher Sharpe ratios of the trading strategy, when option liquidity is low, market uncertainty is low and attention to the 0DTE market (as measured by Google trends) is low.

The remainder of the paper is organized as follows. After a brief discussion of the related literature, Section 2 describes the theoretical framework behind our analysis. Section 3 presents the data and implementation details of the methods we use. Section 4 contains our empirical analysis. Section 5 reports robustness results with respect to different subsamples, monetary policy announcements, using E-mini S&P 500 futures data instead of the cash index, and alternative pricing kernel estimates. Section 6 concludes the paper. Appendix A collects the figures and tables of the paper.

1.1 Related literature

Our paper mainly relates to three strands of the literature. The first strand consists of an increasing number of papers studying the new 0DTE option market from different lenses. Brogaard, Han, and Won (2023) show that a higher fraction of 0DTE option trading increases the volatility of the underlying asset, while Dim, Eraker, and Vilkov (2024) and Adams, Fontaine, and Ornathanalai (2024) provide contrary evidence based on net open interest measures. Beckmeyer, Branger, and Gayda (2023) document that 0DTE options are popular among retail traders, even though these investors mainly experience losses in this market.⁴ Bandi et al. (2023) present a novel option pricing formula designed for 0DTEs and investigate how leverage and volatility-of-volatility affect instantaneous risk premia. Vilkov (2023) explores the performance of 0DTE option trading strategies. Focusing on 1DTE options instead, Johannes, Kaeck, Seeger, and Shah (2024) analyze

⁴While 0DTE options account for more than 75% of retail trading in S&P 500 options, the vast majority of 0DTE S&P 500 trading (around 94%) is still attributable to institutional investors.

option returns around macroeconomic announcements. [Chong and Todorov \(2024\)](#) use 0DTEs to show that there is no segmentation between the equity and option markets based on restrictions for how short-horizon volatility should behave.⁵ We contribute by recovering investors' risk preferences implied by 0DTEs, analyzing their implications for intra-day equity and variance risk premia, and investigating whether these options are mispriced relative to the underlying asset in the stochastic dominance sense.

The second strand recovers information about investors' expectations and risk preferences from relatively long-maturity options. [Jackwerth \(2000\)](#), [Ait-Sahalia and Lo \(2000\)](#) and [Rosenberg and Engle \(2002\)](#) estimate the projection of the pricing kernel onto market return states. [Almeida and Freire \(2022\)](#) find that S&P 500 option prices satisfy bounds consistent with risk-averse investors, where the preferences of the marginal agent vary across the options. [Beason and Schreindorfer \(2022\)](#) decompose the equity premium into different parts of the return state space. [Bollerslev et al. \(2009\)](#) show that the variance risk premium is generally positive and helps predict market returns. [Bollerslev, Todorov, and Xu \(2015\)](#) and [Andersen, Fusari, and Todorov \(2015, 2017\)](#) document the special role of compensation for jump risk in determining market risk premia. Using information from the new 0DTE market, which is the most relevant option market today, we provide novel asset pricing stylized facts for intra-daily horizons that are strikingly different from those obtained for longer horizons. In particular, we study the intra-day pricing kernel implied by 0DTEs.⁶ As a methodological contribution, we discuss how all these elements are connected through the shape of the pricing kernel as a function of market returns.

The third literature strand investigates return predictability over relatively short horizons. [Gao, Han, Zhengzi Li, and Zhou \(2018\)](#) and [Baltussen, Da, Lammers, and Martens](#)

⁵As an implication of their result, the aggregate pricing kernel that reconciles 0DTE option prices with high-frequency market returns exists, which is the object of our study.

⁶Relatedly, [Aletti and Bollerslev \(2024\)](#) study intra-day realizations of a pricing kernel obtained from high-frequency returns of factors constructed from a monthly conditioning set of variables. In contrast, the pricing kernel we analyze is a function of market returns that is forward-looking and conditional on the investors' information set at the time of the day that the 0DTE options are observed.

(2021) document intra-day momentum patterns across different markets, relating them to infrequent portfolio rebalancing and hedging demand, respectively. [Aït-Sahalia, Fan, Xue, and Zhou \(2022\)](#) study the predictability of ultra high-frequency stock returns using machine learning methods. [Aleti, Bollerslev, and Siggaard \(2023\)](#) predict intra-day market returns with high-frequency cross-sectional returns of the factor zoo. [Almeida, Ardison, Freire, Garcia, and Orłowski \(2023\)](#), [Almeida, Freire, Garcia, and Hizmeri \(2023\)](#) and [Alexiou, Bevilacqua, and Hizmeri \(2023\)](#) use high-frequency market returns, cross-sectional stock returns and option returns, respectively, to estimate volatility and tail risk measures and predict risk premia over daily horizons. We show that 0DTEs contain useful predictive information about intra-day risk premia over the options horizon. More specifically, the variance risk premium negatively predicts excess market returns over the day, which is driven by a strong negative relation between the premium for positive return variation risk and future market returns.

2 Theoretical background

In this section, we present the theoretical background behind our analysis of the asset pricing implications of 0DTE options. We first describe what is the pricing kernel implied by options. Then, we discuss its relation with expected option returns, market risk premia and option price bounds. Finally, we explain how this framework will be used to study the 0DTE option market.

2.1 The pricing kernel

In the absence of arbitrage, the current price P_t of any asset is given by the expectation of the future asset payoff X_T at time $T = t + \tau$ multiplied by the pricing kernel $m_{t,T}$:

$$P_t = \mathbb{E}^{\mathbb{P}}[m_{t,T} X_T | \mathcal{F}_t] \equiv \int X_T(s) m_{t,T}(s) \pi_{t,T}^{\mathbb{P}}(s) ds, \quad (1)$$

where s represents the state of the economy, \mathcal{F}_t is the information available to investors at time t and $\pi_{t,T}^{\mathbb{P}}(s)$ is the probability density function (PDF) under the physical measure \mathbb{P}_t . The pricing kernel distorts the physical measure as to reflect investors' compensation for risk, such that one can take simple expectations to calculate the price of any asset.

More specifically, given the almost sure positivity of $m_{t,T}$ under no-arbitrage, the pricing kernel induces a change of measure from the physical measure \mathbb{P}_t to the risk-neutral measure \mathbb{Q}_t . Given a risk-free rate R_f from t to T , this can be seen by noting that $\mathbb{E}_t[m_{t,T}] = 1/R_f$, and dividing and multiplying (1) by $\mathbb{E}_t[m_{t,T}]$:

$$P_t = \frac{1}{R_f} \int X_T(s) \frac{m_{t,T}(s)}{\mathbb{E}_t[m_{t,T}]} \pi_{t,T}^{\mathbb{P}}(s) ds = \frac{1}{R_f} \int X_T(s) \pi_{t,T}^{\mathbb{Q}}(s) ds \equiv \frac{1}{R_f} \mathbb{E}^{\mathbb{Q}}[X_T | \mathcal{F}_t], \quad (2)$$

where $\pi_{t,T}^{\mathbb{Q}}(s)$ is the PDF under the risk-neutral measure \mathbb{Q}_t . This PDF is often called the state-price density, as it defines the (forward) prices of Arrow-Debreu securities paying one dollar at time T if state of nature s is realized, and zero elsewhere. In contrast, $\pi_{t,T}^{\mathbb{P}}(s)$ can be interpreted as the expected payoff of an Arrow-Debreu security for state s .

From (1) and (2), it becomes evident that the pricing kernel is the ratio of discounted risk-neutral probabilities and physical probabilities:

$$m_{t,T}(s) = \frac{1}{R_f} \frac{\pi_{t,T}^{\mathbb{Q}}(s)}{\pi_{t,T}^{\mathbb{P}}(s)}. \quad (3)$$

The economy-wide pricing kernel above depends on the realization of the state s of the economy. However, there is no consensus among researchers on which are the relevant state variables to consider from a modeling perspective.

As an alternative, a large strand of the literature, starting with [Jackwerth \(2000\)](#), [Ait-](#)

Sahalia and Lo (2000) and Rosenberg and Engle (2002), has proposed to focus instead on the projection of the pricing kernel onto states $R_{t,T}$ of the market return:

$$m_{t,T}(R_{t,T}) = \frac{1}{R_f} \frac{\pi_{t,T}^{\mathbb{Q}}(R_{t,T})}{\pi_{t,T}^{\mathbb{P}}(R_{t,T})}. \quad (4)$$

The main advantage is that this projection can be estimated using S&P 500 options and the time series of market returns. On the one hand, the seminal result of Breeden and Litzenberger (1978) allows to recover from option prices across different strikes the risk-neutral distribution of underlying returns over the maturity τ of the options, $\pi_{t,T}^{\mathbb{Q}}(R_{t,T})$. On the other hand, historical market returns are informative about the physical distribution $\pi_{t,T}^{\mathbb{P}}(R_{t,T})$.⁷ Importantly, $m_{t,T}(R_{t,T})$ has the same pricing implications as the economy-wide pricing kernel $m_{t,T}(s)$ for assets with payoffs that depend only on $R_{t,T}$.

The focus on the projection of the pricing kernel onto S&P 500 returns is also justified by the general interest in learning about investors' risk preferences towards the market index and associated equity and variance risk premia. In particular, if one assumes that the market index is equal to the aggregate wealth and a representative agent exists, $m_{t,T}(R_{t,T})$ is the marginal utility of this agent. Under this interpretation, the pricing kernel should be monotonically decreasing in market returns if the representative agent is risk-averse. However, the literature provides extensive evidence for monthly or longer horizons that $m_{t,T}(R_{t,T})$ is usually a nonmonotonic (generally U-shaped) function of market returns instead, characterizing the pricing kernel puzzle (Cuesdeanu and Jackwerth, 2018).⁸ In the next subsections, we discuss how various objects of interest in our analysis relate to the pricing kernel and its shape.

⁷There is a potential mismatch of conditioning information sets when estimating $\pi_{t,T}^{\mathbb{Q}}(R_{t,T})$ with option prices, that are forward-looking, and $\pi_{t,T}^{\mathbb{P}}(R_{t,T})$ with historical returns, that are backward-looking (see, e.g., Linn, Shive, and Shumway, 2018). We discuss how we handle that empirically in Section 3.4.

⁸More recently, Almeida and Freire (2023) show that if one interprets $m_{t,T}(R_{t,T})$ as representing the preferences of a marginal agent in the option market instead of a representative investor, a nonmonotonic shape is not puzzling but rather reflects the risk exposures from the marginal agent's options positions.

2.2 Expected option returns

The shape of the pricing kernel has direct implications for expected option returns. The expected return of a call option can be defined as below:⁹

$$\mu_t^c(S_t, R_{t,T}, K) = \frac{\mathbb{E}_t[\max(S_t R_{t,T} - K, 0)]}{\mathbb{E}_t[\max(S_t R_{t,T} - K, 0) m_{t,T}(R_{t,T})]} - 1, \quad (5)$$

where S_t is the market index at t and K is the option strike price. The expected return of a put option is analogously defined for the payoff $\max(K - S_t R_{t,T}, 0)$. The numerator in (5) is the expected payoff of the option under the physical measure, while the denominator is the expected payoff under the risk-neutral measure, i.e., the option price.

Coval and Shumway (2001) show that if $m_{t,T}(R_{t,T})$ is monotonically decreasing, calls (puts) have expected returns that are positive (negative) and increase with the strike price. Intuitively, a monotonically decreasing $m_{t,T}(R_{t,T})$ shifts probability mass towards states where the call (put) is less (more) valuable. Therefore, as the strike increases, the call price decreases by more than the expected payoff, increasing the expected return. Conversely, as the strike decreases, the put price decreases by less than the expected payoff, decreasing the expected return. If, instead, the pricing kernel is a U-shaped function of returns (i.e., increasing for some region of positive returns), Bakshi et al. (2010) show that expected call returns are decreasing in the strike and negative beyond a strike threshold. This is because $m_{t,T}(R_{t,T})$ shifts probability mass towards states where the call option is more valuable, such that the call price decreases by less than the expected payoff as the strike increases, decreasing the expected return.¹⁰

⁹To see that, note that $\mathbb{E}_t\left[\frac{\max(S_t R_{t,T} - K, 0)}{\mathbb{E}_t[\max(S_t R_{t,T} - K, 0) m_{t,T}(R_{t,T})]}\right] = \frac{\mathbb{E}_t[\max(S_t R_{t,T} - K, 0)]}{\mathbb{E}_t[\max(S_t R_{t,T} - K, 0) m_{t,T}(R_{t,T})]}$.

¹⁰Since a U-shaped pricing kernel is declining in the region of negative market returns, the implications for expected put returns are the same as under a monotonically decreasing shape.

2.3 Equity premium

The shape of $m_{t,T}(R_{t,T})$, or, equivalently, how $\pi_{t,T}^{\mathbb{Q}}(R_{t,T})$ relates to $\pi_{t,T}^{\mathbb{P}}(R_{t,T})$, is informative about which return states are compensated via the equity premium. To see that, first note that the conditional equity premium can be written as:

$$\mathbb{E}_t[R_{t,T}] - R_f = \int_0^\infty R_{t,T}[\pi_{t,T}^{\mathbb{P}}(R_{t,T}) - \pi_{t,T}^{\mathbb{Q}}(R_{t,T})]dR_{t,T}. \quad (6)$$

To decompose the unconditional equity premium, [Beason and Schreindorfer \(2022\)](#) take the unconditional expectation of (6) and consider net market returns $\tilde{R} = R - 1$ to define:

$$EP(x) = \frac{\int_{-1}^x \tilde{R} [\pi^{\mathbb{P}}(\tilde{R}) - \pi^{\mathbb{Q}}(\tilde{R})]d\tilde{R}}{\int_{-1}^\infty \tilde{R} [\pi^{\mathbb{P}}(\tilde{R}) - \pi^{\mathbb{Q}}(\tilde{R})]d\tilde{R}}, \quad (7)$$

where $\pi^{\mathbb{P}}(\tilde{R}) = \mathbb{E}[\pi_{t,T}^{\mathbb{P}}(\tilde{R}_{t,T})]$ and $\pi^{\mathbb{Q}}(\tilde{R}) = \mathbb{E}[\pi_{t,T}^{\mathbb{Q}}(\tilde{R}_{t,T})]$. $EP(x)$ measures the fraction of the average equity premium that is associated with market returns below x .

Returns around zero contribute only marginally to the equity premium (as $\tilde{R} \approx 0$), such that $EP(x)$ is expected to be flat around those states. Outside this region, the shape of the $EP(x)$ function will depend on the shape of the (average) pricing kernel. Under a monotonically decreasing pricing kernel, every state \tilde{R} contributes positively to the equity premium, i.e., $EP(x)$ is always increasing. To see that, note that in the negative return region, $\tilde{R} < 0$ and the pricing kernel is above 1, which means that $\pi^{\mathbb{P}}(\tilde{R}) - \pi^{\mathbb{Q}}(\tilde{R}) < 0$, such that $\tilde{R} [\pi^{\mathbb{P}}(\tilde{R}) - \pi^{\mathbb{Q}}(\tilde{R})] > 0$ and $EP(x)$ is increasing. Analogously, in the positive return region, $\tilde{R} > 0$ and the pricing kernel is below 1, which means that $\pi^{\mathbb{P}}(\tilde{R}) - \pi^{\mathbb{Q}}(\tilde{R}) > 0$, such that $\tilde{R} [\pi^{\mathbb{P}}(\tilde{R}) - \pi^{\mathbb{Q}}(\tilde{R})] > 0$ and $EP(x)$ is again increasing. Now, if we consider instead a U-shaped pricing kernel where $\pi^{\mathbb{Q}}(\tilde{R})$ is above $\pi^{\mathbb{P}}(\tilde{R})$ (i.e., pricing kernel is above 1) for some positive return $\tilde{R} > 0$, then $\tilde{R} [\pi^{\mathbb{P}}(\tilde{R}) - \pi^{\mathbb{Q}}(\tilde{R})] < 0$ and $EP(x)$ is decreasing, i.e., such positive returns contribute negatively to the equity premium.

We give a new economic interpretation for these relations, which is as follows. For

each state \tilde{R} , consider the asset that pays \tilde{R} in this state and zero otherwise, i.e., the asset defined by buying \tilde{R} units of the Arrow-Debreu security of state \tilde{R} . Then, $\tilde{R} [\pi^{\mathbb{P}}(\tilde{R}) - \pi^{\mathbb{Q}}(\tilde{R})]$ is the expected payoff minus the price of this asset. If the pricing kernel is monotonically decreasing, these assets have a low (high) payoff when the pricing kernel is high (low), such that they are speculative assets with a positive expected return, i.e., $\tilde{R} [\pi^{\mathbb{P}}(\tilde{R}) - \pi^{\mathbb{Q}}(\tilde{R})] > 0$. In other words, investors would require compensation for holding any of these assets, such that all states contribute positively to the equity premium. In contrast, if there is a U-shape where the pricing kernel is high for a region of positive returns, the assets in this region will have a high payoff when the pricing kernel is high, such that they are hedging assets with a negative expected return, i.e., $\tilde{R} [\pi^{\mathbb{P}}(\tilde{R}) - \pi^{\mathbb{Q}}(\tilde{R})] < 0$. That is, investors would be willing to give up compensation to hedge against these states, such that they contribute negatively to the equity premium.

2.4 Variance risk premium

The pricing kernel projection onto market returns is also informative about the magnitude of the variance risk premium. Defined as the difference between the risk-neutral and physical expected variance of the market return over horizon τ , it reflects the compensation investors require for bearing variance risk (Bollerslev et al., 2009). Baele, Driessen, Ebert, Londono, and Spalt (2019) demonstrate that the variance risk premium is closely related to expected option returns. In particular, it can be written as a weighted average of expected returns of put and call options across strikes, with negative weights. Consequently, the variance risk premium is higher under a U-shaped pricing kernel, where both call and put expected returns are negative, than under a monotonically decreasing pricing kernel. Intuitively, this premium reflects compensation for extreme negative and positive return states, such that this compensation is higher when $m_{t,T}(R_{t,T})$ is U-shaped.

It is also possible to decompose the total variance risk premium into the specific com-

compensation for variation risk in positive and negative market returns (Kilic and Shaliastovich, 2019). Analogously to above, the risk premium for positive (negative) return variation is a weighted average of expected call (put) returns, with negative weights. If the pricing kernel is monotonically decreasing (U-shaped), expected call returns are positive (mostly negative) and the positive return variation premium is negative (positive).¹¹ That is, if marginal utility is high for positive market returns, investors would be willing to pay a premium to be protected against large positive returns. On the other hand, the more negative expected put returns are, the steeper is the pricing kernel for negative returns and the higher is the compensation for negative return variation.

The expected returns of specific option strategies contain further information about the variance risk premium. Bakshi and Kapadia (2003) and Coval and Shumway (2001) show, respectively, that negative delta-hedged option returns and negative straddle returns reflect a positive variance risk premium. These strategies profit from (and are a hedge against) increases in market volatility, such that a negative expected return indicates that investors are willing to pay a premium to be protected against variance risk.

2.5 Option price bounds

The shape of $m_{t,T}(R_{t,T})$ tells us how option prices are jointly reconciled with the physical distribution implied by market returns. A related, but alternative way of investigating this relation is by comparing each observed option price with option price bounds consistent with the physical distribution $\pi_{t,T}^{\mathbb{P}}(R_{t,T})$ and specific risk preferences. Of particular interest for us are the second-order stochastic dominance (SSD) bounds (Levy, 1985; Perrakis and Ryan, 1984; Ritchken, 1985), which provide the maximum and minimum price for a given option consistent with risk-averse investors trading in the

¹¹More precisely, under a U-shaped pricing kernel, expected call returns are only negative beyond a strike that depends on how pronounced the U-shape is. Therefore, if the pricing kernel is only mildly U-shaped, it can still be the case that expected call returns are mostly positive and the positive return variation premium is negative.

underlying asset and the risk-free rate. In other words, these bounds allow for market segmentation by giving all option prices compatible with the set of pricing kernels that are monotonically decreasing in the market returns.

A violation of the SSD lower (upper) bound by a given option means that any risk-averse investor can improve expected utility by taking a long (short) position in the option, or, equivalently, that the option dominates (is dominated by) the underlying asset by second-order stochastic dominance. That is, a violation would mean that there is no marginal investor in the index, risk-free rate and the option with a monotonically decreasing pricing kernel (i.e., satisfying risk aversion). This option is often regarded as “mispriced” as its price cannot be reconciled with the physical distribution under reasonable risk preferences.

It is important to note that SSD bounds offer related, but complementary insights relative to the pricing kernel in (4). If $m_{t,T}(R_{t,T})$ is monotonically decreasing, then this pricing kernel is part of the SSD admissible set and the option prices will be inside the SSD bounds. In contrast, if $m_{t,T}(R_{t,T})$ is nonmonotonic, this means that no unique monotonically decreasing pricing kernel prices all options, but it is still possible that different monotonically decreasing pricing kernels price the different options in the cross-section. Conversely, if option prices satisfy SSD bounds, this does not mean that $m_{t,T}(R_{t,T})$ is monotonically decreasing, whereas if they violate the bounds, this implies that $m_{t,T}(R_{t,T})$ is nonmonotonic. In fact, while empirical evidence favors a U-shaped pricing kernel projection at the one-month horizon, [Almeida and Freire \(2022\)](#) show that S&P 500 option prices generally satisfy SSD bounds, where options with different moneyness require heterogeneous investors who differ in their assessment of tail risk to be priced.¹²

¹²This evidence differs from that by [Constantinides, Jackwerth, and Perrakis \(2009\)](#), who find substantial violations of SSD bounds in the S&P 500 option market. The main reason for this difference is that the conditional physical distribution they estimate keeps the volatility constant over long periods of time, during which volatility varies considerably. In contrast, [Almeida and Freire \(2022\)](#) adjust volatility daily in the estimation of the conditional physical distribution.

2.6 Empirical strategy

The theoretical framework described above provides insights into investors' risk preferences over the horizon τ corresponding to the maturity of the options used. Most of the literature applying these methods has focused on the one-month horizon or longer, partially motivated by the liquidity of the associated options. However, as previously described, the option market has changed dramatically over the last few years. Now, the most traded contracts are 0DTE options, which give investors the opportunity to hedge against or make leveraged bets on specific market movements over the day. As such, 0DTEs are a valuable source of information about risk premia and compensation for risk at intra-daily horizons. We aim to extract and analyze this information.

We will consider different times of the day for t , while T will always be the market close, which is usually at 16:00. More specifically, our analysis will be based on the cross-section of 0DTEs at 10:00, 10:30, 11:00, 11:30, 12:00, 12:30, 13:00, 13:30 and 14:00, such that the horizon/option maturity τ will be 6, 5.5, 5, 4.5, 4, 3.5, 3, 2.5 and 2 hours, respectively. In the next section, we show that this range of times of the day is where the relative option bid-ask spread is reasonably stable at its minimum over the day. For each of those times of the day, we will estimate the pricing kernel in (4), calculate option returns of different strategies (from t to T), decompose the corresponding equity premium, estimate the variance risk premium and compute SSD bounds for each option.

3 Data description and implementation details

3.1 Data

We obtain the intra-day S&P 500 option data from CBOE, which includes bid and ask quotes, trading volume, open interest and underlying asset price at 1-minute intervals. We define the price of an option as the bid and ask midpoint. We select all dates

between January 6 2012 and July 3 2023 for which 0DTE SPXW options are available. The first weeklies were introduced by CBOE on October 28 2005 with Friday expirations. Wednesday, Monday, Tuesday and Thursday expirations followed with introduction dates February 23 2016, August 15 2016, April 18 2022 and May 11 2022, respectively. This means that our sample contains one day per week up until February 23 2016, then two days per week until August 15 2016, three days per week until April 18 2022, four days per week until May 11 2022, and all days of the week afterwards. We have 1,417 dates in total, where roughly 20% of those dates is between May 11 2022 and July 3 2023.

Figure 1 depicts the striking evolution of the 0DTE option market over the last years in terms of its fraction of trading volume relative to the entire S&P 500 option market. While 0DTEs accounted for only around 2% of total trading volume in 2012, today they represent nearly 45% of the entire option market and are the most traded maturity. As noted by [Bandi et al. \(2023\)](#), this corresponds to a daily notional dollar volume of around 1 trillion dollars. This meteoric increase can partially be explained by the daily availability of 0DTE options since May 2022, allowing investors to hedge and make leveraged bets on specific intra-daily market movements for any day of the week.

To clean the raw option data, we aim to avoid as much as possible options with small trading volume and zero bid, while also selecting a comparable set of strikes over time. What mainly defines the range of strikes being traded on a given day is volatility: on days where volatility is high (low), large return realizations are more (less) likely to occur, such that investors trade options for a larger (smaller) range of strikes. For this reason, we classify options in terms of their standardized log-moneyness $k_{std} = \frac{\log(K/S_t)}{\sigma_{BS}(0)\sqrt{\tau}}$, which controls for the level of volatility as $\sigma_{BS}(0)$ is the ATM implied volatility for time of the day t and maturity τ .

By analyzing the 0DTE option data, we identify that the range of k_{std} between -6 and 6 strikes a good balance between trading volume and low proportion of zero bids.

This can be visualized in the upper subplots of Figure 2, which report, for different times of the day and bins of k_{std} , the average trading volume and proportion of contracts with zero bid over time. As can be seen, the bulk of trading volume is within the k_{std} range of -3 and 3 , whereas the volume outside the interval $[-6, 6]$ is negligible. At the same time, zero bids are essentially inexistent for $k_{std} \in [-3, 2]$ and then occur increasingly more for more extreme moneyness levels. Our interval of $[-6, 6]$ avoids the extremely large proportion of zero bids of deeper OTM put options.¹³

The bottom subplots of Figure 2 further display the average implied volatilities (IVs) and average number of strikes for each bin. For all times of the day, 0DTE IVs display a smile across moneyness, where OTM puts and calls are equally expensive in terms of IV. This is in contrast to the usual smirk observed for longer-maturity S&P 500 options (where IVs are much higher for OTM puts than for OTM calls). The average IVs outside $k_{std} \in [-6, 6]$ are extremely high, which highlights the importance of excluding these options that are not traded and would contaminate results.¹⁴ As for the average number of strikes, it is approximately constant across moneyness bins and decreases from around 4 at 10:00 to 2 strikes per bin at 14:00.

Having defined our option sample, we proceed to choose a set of times of the day for our analysis with the goal of being representative while feasible to report results. To guide our choice, Figure 3 reports, for our option sample, the 0DTE trading volume and relative bid-ask spread over the day. While trading volume is higher at market open and close, these times of the day are also the ones with highest bid-ask spread over the day. The bid-ask spread tends to be relatively stable at its minimum between 10:00 and 14:00.

¹³While we could have focused on the smaller interval $k_{std} \in [-3, 2]$, this would have resulted in selecting a very narrow range of strikes in days with small volatility. Note also that working with the larger interval $[-6, 6]$ has little effect in our results as most of our findings are driven by what happens around the ATM region, for which zero bids are not observed.

¹⁴IVs outside $k_{std} \in [-6, 6]$ are that high due to the deeper OTM options that have very small, but still positive prices, while the probability of returns occurring such that they finish in-the-money is virtually zero. Since these deeper OTM options are not traded, their extremely high IVs are artificial and do not reflect market expectations.

For this reason, we select this range of times of the day for our analysis, with intervals of 30 minutes to keep it feasible to report the results.

Finally, we obtain high-frequency data on the S&P 500 index, spanning January 1996 to July 2023, and E-mini S&P 500 futures data for the same period as our option data from Refinitiv Tick History. We perform standard cleaning procedures and sample our data at 1-minute intervals during regular trading hours (9:30 to 16:00 EST). The E-mini futures data are front-month contracts with roll-over on contract expiry. The risk-free rate, which we obtain for the same sample period, is the daily one-month Treasury bill rate from the FRED (Federal Reserve Bank of St. Louis) website. We assume that the risk-free rate remains constant throughout the day.

3.2 Option returns

For a given day in our sample and time of the day t , we compute hold-until-maturity returns of different portfolios of options expiring at the end of the day. We first calculate the return of the call (put) option with price $O_{t,T}^c$ ($O_{t,T}^p$) and strike K as:

$$R_c = \frac{\max(S_T - K, 0)}{O_{t,T}^c} - 1, \quad R_p = \frac{\max(K - S_T, 0)}{O_{t,T}^p} - 1, \quad (8)$$

where S_T is the market index at the maturity of the option. We will analyze how call and put returns vary with the strike price over our sample, which is informative about the shape of the intra-day pricing kernel. To have returns of options with exactly the same moneyness for each day of our sample, we use the interpolated option prices coming from the procedure described in the next subsection.

Then, we compute the returns of different option strategies that are insightful about the variance risk premium. First, we consider simple straddle returns obtained from the

simultaneous purchase of an ATM call and ATM put with strike K :

$$\text{Simple-Straddle} = \frac{\max(S_T - K, 0) + \max(K - S_T, 0)}{O_{t,T}^c + O_{t,T}^p} - 1. \quad (9)$$

We focus on the ATM straddle that is more exposed to volatility risk. In addition, we calculate the return of an exactly delta-neutral ATM straddle:

$$\text{Straddle} = wR_c + (1 - w)R_p, \quad w = -\frac{\Delta_p/O_{t,T}^p}{\Delta_c/O_{t,T}^c - \Delta_p/O_{t,T}^p}, \quad (10)$$

where Δ_c (Δ_p) is the call (put) Black-Scholes delta.

Finally, we follow [Bakshi and Kapadia \(2003\)](#) by calculating ATM delta-hedged call returns as:¹⁵

$$\Delta\text{-Hedged} = \frac{\max(S_T - K, 0) - O_{t,T}^c - \Delta_c(S_T - S_t) - r_t^f(O_{t,T}^c - \Delta_c S_t) \times \frac{\tau}{(24 \times 365)}}{S_t}, \quad (11)$$

where r_t^f is the annualized risk-free rate and τ is the time to maturity expressed in hours, e.g., $\tau = 6$. Since the straddle and delta-hedged strategies are essentially long positions in market volatility, we will analyze their average returns over our sample to extract information about the variance risk premium over the day. For these strategies, we use the observed price of the option closest to ATM.

3.3 Risk-neutral distribution

For a given day of our sample and time of the day, we estimate the risk-neutral distribution from the cross-section of ODTE option prices. [Breen and Litzenberger \(1978\)](#) show that, under no-arbitrage and in the presence of a continuum of options

¹⁵Results are very similar if, instead of the underlying price S_t , we set the denominator to be the initial investment absolute cost $|O_{t,T}^c - \Delta_c S_t|$.

across strikes, risk-neutral probabilities are equal to the risk-free rate times the second derivative of option prices with respect to the strike price:

$$\pi_{t,T}^{\mathbb{Q}}(R_{t,T}) = S_t \times R_f \times \left. \frac{\partial^2 O_{t,T}^c(K)}{\partial K^2} \right|_{K=S_T}, \quad (12)$$

where the strikes represent different states of the underlying asset price at maturity and multiplying by S_t performs a change of variables from $\pi_{t,T}^{\mathbb{Q}}(S_T)$ to $\pi_{t,T}^{\mathbb{Q}}(R_{t,T})$.

In practice, however, we only observe a discrete set of strikes that sometimes does not cover the whole range of moneyness. For this reason, it is necessary to interpolate and extrapolate observed option prices to compute the derivative and estimate $\pi_{t,T}^{\mathbb{Q}}(R_{t,T})$. To do so, we follow the standard practice in the literature of converting option prices to IVs using the [Black and Scholes \(1973\)](#) formula, fitting an interpolant to them, using the interpolant to generate IVs for a fine grid of strikes, translating IVs back to option prices, and computing (12) over the fine grid of strikes via finite differences.¹⁶ Only OTM options are used to fit the interpolant, as they are more liquid than in-the-money (ITM) options, which should contain redundant information by put-call parity.

We fit the IV curve across strikes using the parsimonious Stochastic Volatility Inspired (SVI) method of [Gatheral \(2004\)](#). This method has also been used by [Beason and Schreindorfer \(2022\)](#) to estimate the risk-neutral distribution and combines reliable interpolation of the IV curve with well-behaved extrapolation for extreme moneyness levels. More specifically, the SVI describes the square of IV with the function:

$$\sigma_{BS}^2(k) = a + b \left\{ \rho(k - m) + \sqrt{(k - m)^2 + \sigma^2} \right\}, \quad (13)$$

¹⁶It is important to emphasize that this approach does not assume that the [Black and Scholes \(1973\)](#) model is valid. Rather, the [Black and Scholes \(1973\)](#) formula is only used as a one-to-one mapping between option prices and IVs. This is done because fitting the IV curve is much easier than fitting option prices as IVs are comparable across strikes.

where $k = \log(K/S_t)$ is the log-moneyness and a, b, ρ, m and σ are parameters.¹⁷ We fit (13) to the cross-section of observed 0DTE IVs and estimate the parameters by minimizing the IV mean squared error with a constrained nonlinear programming solver.¹⁸

Figure 4 plots the average over our sample of the observed 0DTE IVs and the fitted IVs using the SVI method. The SVI provides an excellent fit to the smile displayed by 0DTEs, with average OLS R^2 's higher than 95%. This, in turn, results in well-behaved risk-neutral distributions, as can be seen in Figure OA.1 for a particular day of our sample. These distributions are obtained by interpolating and extrapolating the IVs in a grid using the SVI method, mapping the IVs back to option prices and then computing $\pi_{t,T}^{\mathbb{Q}}(R_{t,T})$ via the Breeden and Litzenberger (1978) formula. The plot shows how the broad range of standardized log-moneyness we consider translates to a narrow range of return states that investors believe the market can experience over this particular day. As the time gets closer to market close, the risk-neutral distribution gets narrower, reflecting that there is less room for large return realizations.

3.4 Physical distribution

To calculate the pricing kernel projection and the SSD option price bounds, we need to estimate the conditional physical distribution $\pi_{t,T}^{\mathbb{P}}(R_{t,T})$. Aït-Sahalia and Lo (2000) and Jackwerth (2000) rely directly on the historical market return distribution, where there is a trade-off between using a short sample, which makes the distribution conditional, and using a long sample, which improves the estimation precision, especially for the tails. A more recent approach has been to take the unconditional return distribution from a long

¹⁷More specifically, a controls the IV level, b the IV slope, ρ the asymmetry of the IV slope for negative and positive k , m the horizontal location of the IV curve, and σ the ATM curvature of the IV curve.

¹⁸The SVI is well defined for $a \in \mathbb{R}$, $b \geq 0$, $|\rho| \leq 1$, $m \in \mathbb{R}$, $\sigma > 0$ and $a + b\sigma\sqrt{1-\rho^2} \geq 0$. We impose these constraints in the optimization, with two small modifications: we replace $a + b\sigma\sqrt{1-\rho^2} \geq 0$ with the slightly stronger restriction $a \geq 0$, which yields better behaved extrapolations for the right tail, and we impose $\sigma \geq 0.01$, which helps discipline the IV ATM curvature.

sample and make it conditional by adjusting for the conditional volatility at time t using GARCH models (see, e.g., [Barone-Adesi, Engle, and Mancini, 2008](#)). This preserves the empirical patterns of skewness, kurtosis, and tail probabilities, but has the drawback that GARCH models can be misspecified. Even if one uses realized variance instead to make the distribution conditional, there is still a potential mismatch of conditioning sets in comparing backward-looking information from historical returns with forward-looking information from option prices (see, e.g., [Linn et al., 2018](#)).

We follow a similar approach to [Almeida and Freire \(2022\)](#) to overcome the issues above. For a given day in our sample and time of the day, we first estimate the historical return distribution as the histogram of past market returns from t to T over a long sample starting on January 1996.¹⁹ Then, we make the return distribution conditional by setting its volatility equal to the ATM IV at time t of the current day, which is forward-looking. That is, we use minimal option information to make the conditioning sets of $\pi_{t,T}^{\mathbb{P}}(R_{t,T})$ and $\pi_{t,T}^{\mathbb{Q}}(R_{t,T})$ comparable. While the ATM IV contains a risk premium, other papers such as [Dew-Becker, Giglio, and Kelly \(2021\)](#) also use it as a proxy for the expected physical volatility given its good performance in forecasting future realized volatility. In [Section 5.4](#), we show that our results are robust to alternative specifications of the conditional physical distribution of market returns.

The resulting conditional physical distribution is an unsmoothed histogram. For the purpose of estimating the pricing kernel projection, it is necessary to smooth it to obtain a well-behaved PDF. We follow [Jackwerth \(2000\)](#) in fitting a kernel density with a Gaussian kernel to smooth the histogram and obtain $\pi_{t,T}^{\mathbb{P}}(R_{t,T})$.²⁰ [Figure OA.1](#) displays, for a

¹⁹Following [Almeida and Freire \(2022\)](#), we also impose the sensible economic restriction of a 5% lower bound on the annualized equity premium over the risk-free rate. That is, if the annualized mean of the unconditional return distribution generates a premium less than 5% over the risk-free rate, we demean the returns and reintroduce a 5% equity premium. [Jackwerth \(2000\)](#) imposes a similar restriction.

²⁰For a given day, the obtained histogram is smoothed with a Gaussian kernel with bandwidth $\frac{x\sigma}{\sqrt[5]{n}}$, where σ is the volatility of the return distribution, n is the number of observations, and we set $x = 2.8$ and $m = 5$, which strikes a good balance between smoothness and fit.

representative date of our sample, the conditional physical distribution together with the risk-neutral distribution for different times of the day. The (discounted) ratio between the risk-neutral and physical PDFs gives the estimate of the pricing kernel projection $m_{t,T}(R_{t,T})$ for that day and horizon.

3.5 SSD bounds and Risk premia

Using the estimated physical distribution, we compute, for each call option with strike K , the following SSD upper and lower price bounds as in [Ritchken \(1985\)](#):²¹

$$C_{\max} = \mathbb{E}[\max(S_t R_{t,T} - K, 0)] / \mathbb{E}[R_{t,T}], \quad (14)$$

$$C_{\min} = \mathbb{E}[\max(S_t R_{t,T} - K, 0) | S_t R_{t,T} < s_j^*] \frac{1}{R_f}, \quad (15)$$

where s_j^* is chosen such that $\mathbb{E}[S_t R_{t,T} | S_t R_{t,T} < s_j^*] = S_t R_f$. All expectations are calculated under the estimated $\pi_{t,T}^{\mathbb{P}}(R_{t,T})$ over the grid of states $R_{t,T}$. Equation (14) says that the maximum price of the call should be the price such that the expected call return equals the expected return on the market. The interpretation for the lower bound is less straightforward. [Ritchken \(1985\)](#) uses linear programming techniques to show that these bounds contain all prices for the call option consistent with the set of pricing kernels that are monotonically decreasing in $R_{t,T}$. The bounds for the put option with strike K can be obtained via put-call parity. We will compare the bounds to the observed option prices to identify any potential mispricing of the 0DTEs.

We follow [Beason and Schreindorfer \(2022\)](#) to implement the decomposition of the unconditional equity premium as in (7). First, we estimate the unconditional risk-neutral distribution $\pi^{\mathbb{Q}}(\tilde{R})$ over our sample as the average of the conditional risk-neutral distributions, i.e., the average of $\pi_{t,T}^{\mathbb{Q}}(\tilde{R}_{t,T})$ state by state. With that, we can evaluate

²¹[Perrakis and Ryan \(1984\)](#) and [Levy \(1985\)](#) derive the same bounds following different approaches.

$\int_{-1}^{\infty} \tilde{R} \pi^{\mathbb{Q}}(\tilde{R}) d\tilde{R}$ numerically over the grid of \tilde{R} (which is equal to the grid of $R_{t,T} - 1$). Then, we compute $\int_{-1}^{\infty} \tilde{R} \pi^{\mathbb{P}}(\tilde{R}) d\tilde{R}$ from the unconditional empirical distribution of market returns (of the horizon of the option) over our sample as $(1/\mathcal{T}) \sum_{i=1}^{\mathcal{T}} \tilde{R}_{i,t,T} \mathbf{1}\{\tilde{R}_{i,t,T} \leq x\}$, where \mathcal{T} denotes the total number of days, $\tilde{R}_{i,t,T}$ is the realized market return of day i from time of the day t to market close T , and we consider x 's over the grid of \tilde{R} . This is equivalent to computing the integral under the unconditional physical distribution.

Finally, we compute a measure of the variance risk premium similarly to [Bollerslev et al. \(2009\)](#). They calculate it for the one-month horizon as the risk-neutral expected variance over the next month minus the physical expected variance proxied by the realized variance from the previous month to the current one. Analogously, our ex-ante variance risk premium $VRP_{t,T}$ from time t to T is the expected risk-neutral variance implied by the cross-section of 0DTEs at time t minus the realized variance from t to T of the previous day. The expected risk-neutral variance is computed as in [Bakshi et al. \(2003\)](#):

$$V_{t,T}^{\mathbb{Q}} = \int_{S_t}^{\infty} \frac{2[1 - \log(K/S_t)]}{K^2} O_{t,T}^c(K) dK + \int_0^{S_t} \frac{2[1 + \log(S_t/K)]}{K^2} O_{t,T}^p(K) dK, \quad (16)$$

where we compute the integrals using the interpolated and extrapolated option prices from the SVI method. The realized variance ([Andersen, Bollerslev, Diebold, and Labys, 2003](#)) $RV_{t,T}$ is the sum of 1-minute squared log-returns on the market index. To disentangle the compensation demanded by investors to bear variation risk in positive and negative market returns, we also compute the “good” and the “bad” variance risk premium in the spirit of [Kilic and Shaliastovich \(2019\)](#). The former (latter), $VRP_{t,T}^+$ ($VRP_{t,T}^-$), is defined as the first (second) integral in (16) minus the sum of 1-minute squared market returns times an indicator function for a positive (negative) return. Naturally, $VRP_{t,T} = VRP_{t,T}^+ + VRP_{t,T}^-$.

4 Empirical results

4.1 0DTE option returns

We start by analyzing the returns of different option strategies. Focusing first on how the returns of vanilla call and put options change across strike prices, Figures 5 and 6 plot their average returns over our sample, respectively, together with 90% i.i.d. bootstrap confidence bands.²² Observed patterns are very similar across different times of the day. Call options experience low returns overall, which are decreasing with the strike starting from the ATM region and eventually become negative with statistical significance. Following the rationale of Bakshi et al. (2010), this provides evidence that the intra-day pricing kernel is increasing in a region of positive return states. Taking at face value, these results indicate that writing naked OTM calls is usually profitable in the 0DTE option market. This would be aligned with investors requiring compensation for bearing positive return variation risk. As for put options, average returns are always negative, with statistical significance for all OTM puts. This is consistent with a steep, monotonically decreasing pricing kernel in the region of negative market returns, reflecting investors' aversion to downside risk. Again, writing naked OTM puts is usually a profitable strategy, compatible with compensation for negative return variation risk.

We further consider the returns of ATM delta-hedged call options, straddles and delta-neutral straddles. Figure 7 displays their average returns together with 90% i.i.d. bootstrap confidence bands for different times of the day. All strategies produce negative average returns, consistent with a positive intra-day variance risk premium. Since these strategies are essentially long positions in volatility, a negative average return indicates

²²Due to the small dollar prices of 0DTE options, for a few days returns can be quite extreme. To minimize the effect of these outliers and produce a smoother plot, for each moneyness we winsorize the right-tail of the time-series of the option returns at 0.5%. This has no qualitative effect on the conclusions we obtain from the average returns. In Figures OA.2 and OA.3 in the Online Appendix, we also report percentiles of the returns instead of averages, which convey the same message.

that investors are willing to pay a premium to be protected against variance risk over the day. On the other hand, an investor willing to be exposed to this risk would profit, on average, from shorting delta-hedged calls and straddles in the 0DTE market. The confidence bands indicate that the statistical significance of the negative average returns is stronger from 12:00 to 14:00.

4.2 Intra-day risk premia

We now investigate the implications of 0DTE options for intra-day market risk premia. Figure 8 plots the decomposition of the equity premium across return states for different times of the day. Most of the equity premium stems from compensation to market returns between -5% and 0%. Strikingly, these states account for 300% (3000%) of the total equity premium from 10:00 (14:00) to close, which would amount to an annualized premium of 40% (150%), as seen in the right axis of the plot. More important than the magnitude of these values, however, is the reason why they exceed 100%: positive market returns contribute negatively to the equity premium. This is consistent with a U-shaped pricing kernel, as discussed in Section 2.3. Since marginal utility is high for positive market returns, investors are willing to pay a premium to hold Arrow-Debreu securities paying in these states, such that these states have a negative contribution to the equity premium. In other words, the intra-day equity premium, which is around 10% annualized over our sample, would be much higher if the pricing kernel were monotonically decreasing.

For comparison, [Beason and Schreindorfer \(2022\)](#) show that at the one-month horizon, most of the equity premium stems from returns between -30% and -10%. Moreover, states up to a monthly market return of 5% account for around 120% of the equity premium, while higher returns contribute negatively to it. The more extreme values we find for the intra-day equity premium reveal how the nonmonotonicity of the pricing kernel is relatively much more pronounced at the ultra-short horizons we consider.

We next estimate, for each day and time of the day, the $VRP_{t,T}$ and its two components, $VRP_{t,T}^+$ and $VRP_{t,T}^-$. Table 1 reports summary statistics over our sample of each of these measures for different times of the day. Across all times, the average $VRP_{t,T}$ is high and significantly positive, confirming the evidence from option returns that investors require substantial compensation to bear variance risk over the day. In fact, the annualized $VRP_{t,T}$ varies from 2.87% to 8.00% depending on the time of the day, which is considerably larger than the usual 2% variance risk premium at the monthly horizon (Bollerslev et al., 2009). Interestingly, both $VRP_{t,T}$ components are significantly positive as well. In particular, our evidence that $VRP_{t,T}^+$ is positive is consistent with a U-shaped pricing kernel: investors require compensation to bear variation risk in the region of positive market returns. This is in contrast to the one-month VRP^+ , which Kilic and Shaliastovich (2019) estimate to be slightly negative on average.²³ Even more striking is that the compensation for upside risk is of similar magnitude, and in fact, from 11:30 onwards, larger than the compensation for downside risk as captured by $VRP_{t,T}^-$.

Figure OA.4 in the Online Appendix plots, for different times of the day, the time series of the (one-week moving average of the) $VRP_{t,T}$, $VRP_{t,T}^+$ and $VRP_{t,T}^-$. Up to 2022, the three measures are almost always positive, while afterwards negative values become more frequent. The largest spike in $VRP_{t,T}$ is associated with the COVID-19 crisis. During this period, the total variance risk premium reached extreme values such as 150%. Importantly, this spike is driven by both the upside risk premium ($VRP_{t,T}^+$) and the downside risk premium ($VRP_{t,T}^-$).

4.3 Intra-day return predictability

We further investigate whether ex-ante measures capturing information from 0DTE options are able to predict realized excess market returns from t to T with predictive

²³Note that a slightly negative average one-month VRP^+ does not necessarily mean that the one-month pricing kernel is monotonically decreasing (see footnote 11).

regressions of the following form:

$$R_{t,T} - R_f = a + bX_t + \epsilon_{t,T}, \quad (17)$$

where X_t collects different predictors available in real time at t . We first focus on the variance risk premium and its components, given that these summarize the compensation investors require to bear variation risk in different regions of the market return space. Then, we analyze as predictors the risk-neutral variance, skewness and kurtosis as in [Bakshi et al. \(2003\)](#), the realized variance, and the SVIX of [Martin \(2017\)](#), which represents a lower bound for the equity premium under a negative correlation assumption.

Table 2 reports the results for univariate predictive regressions based on $VRP_{t,T}$, $VRP_{t,T}^+$ and $VRP_{t,T}^-$, and a multivariate regression including both $VRP_{t,T}^+$ and $VRP_{t,T}^-$. The total variance risk premium negatively predicts the intra-day equity premium, with statistical significance at 11:00, 11:30, 12:00 and 13:00. This is at odds with the positive relation documented at monthly or longer horizons ([Bollerslev et al., 2009](#)). Using the two components of the $VRP_{t,T}$ in the predictive regressions sheds light on this finding. The $VRP_{t,T}^+$ is a strong predictor of market returns, with a negative coefficient that is statistically significant at all times of the day, except for 14:00. In contrast, there is a positive relation between $VRP_{t,T}^-$ and future returns, which is generally insignificant and of weaker magnitude. That is, the results for $VRP_{t,T}$ are driven by $VRP_{t,T}^+$. The R^2 of the regressions is also much higher than for the total $VRP_{t,T}$, which mainly comes from the predictive power of $VRP_{t,T}^+$, as can be seen from the univariate regressions.

The negative relation between $VRP_{t,T}^+$ and market returns from t to T is consistent with a U-shaped pricing kernel. As discussed in Section 2.3, the equity premium can be seen as the aggregate compensation required for holding assets paying R in each market return state, and zero otherwise. When marginal utility for positive market return states is high, such assets paying in positive return states are hedges (as they pay

a high payoff when the pricing kernel is high). This is such that investors are willing to give up compensation to hold them, resulting in a smaller equity premium. As $VRP_{t,T}^+$ summarizes the compensation for positive return variation risk, this explains the negative relation with future returns. On the other hand, since the Arrow-Debreu-like assets paying R for negative return states behave as speculative assets (as they pay a low payoff when the pricing kernel is high), investors require compensation to hold them and they contribute positively to the equity premium. This explains the positive relation between $VRP_{t,T}^-$ and future market returns. The fact that the effect of $VRP_{t,T}^+$ is dominant over that of $VRP_{t,T}^-$ reflects the exceptional role that the U-shape plays in the intra-daily horizons. This, again, is in contrast to the one-month horizon or longer where the positive coefficient of $VRP_{t,T}^-$ is predominant and drives the positive relation between the total $VRP_{t,T}$ and the equity premium.

Table OA.1 contains the results for the univariate predictive regressions based on RV , risk-neutral skewness and kurtosis, and $SVIX$. As can be seen, none of these measures is able to predict the intra-day equity premium. That is, we find no evidence of an intra-day risk-return trade-off. In particular, the fact that the lower bound of Martin (2017), $SVIX$, does not predict future returns, could either mean that the bound is not tight at the intra-daily horizons we consider, or that the negative correlation assumption under which the bound is derived is not valid. More specifically, this assumption states that the covariance between $R_{t,T}$ and $m_{t,T}(R_{t,t}) \times R_{t,T}$ must be negative. While this condition is valid under most macro-finance models, it would be violated under a pricing kernel with pronounced nonmonotonicities. Given the extensive evidence from our analysis in favor of such nonmonotonicity over intra-daily horizons, this seems like a plausible explanation.

Table OA.2 further considers predictive regressions based on the variance risk premium measures including the risk-neutral moments as controls.²⁴ We consider times of the day

²⁴Results including RV instead of $SVIX$ are very similar.

for which $VRP_{t,T}$ was significant in the univariate regressions. As can be observed, the inclusion of the controls only makes the negative relation between the total variance risk premium and future market returns stronger and more significant. Among the controls, $SVIX$ becomes marginally significant with a positive coefficient at 11:30 and 13:00, while the other variables have no predictive power. When we replace $VRP_{t,T}$ with its two components, we see that the R^2 increases substantially and only $VRP_{t,T}^+$ is a statistically significant predictor of returns, where a higher value of this variable leads to a lower equity premium. This provides additional robustness to the findings of this subsection.

4.4 Investors' risk preferences

The previous subsections provide indirect evidence that the pricing kernel as a function of market returns is nonmonotonic and high for positive market returns. In this subsection, we estimate the pricing kernel directly for each time of the day and analyze its average shape over our sample. Figure 9 displays the results for a range of (gross) market returns between 0.98 and 1.02. This is the largest range for which the pricing kernels over different days of our sample (including days with low volatility) are well-defined. As can be seen, regardless of the time of the day, there is a pronounced hump around the at-the-money region, i.e., around returns close to 0 (or, equivalently, gross returns close to 1). This is consistent with our previous results: in the region where returns are more likely to occur, marginal utility is higher for positive returns than for negative returns.

To be able to better compare across days with different volatility and incorporate information from the whole range of the return space of each day, Figure 10 reports average pricing kernels over standardized return states (i.e., in the standardized log-moneyness space). Similarly to before, the feature that is pervasive across all times of the day is a pronounced hump around the at-the-money region, where the pricing kernel increases from k_{std} equal to -1 up to 1 , and then decreases. However, nonmonotonicity

is also observed outside this region. For higher return states, depending on the time of the day, there are either additional humps or an overall increasing shape. We refer to all of those patterns generally as a U-shape, in the sense that there are increasing segments in marginal utility for positive returns, even though the exact pattern can sometimes be better described by an S- or W-shape. For negative returns, further nonmonotonicities can occur at extreme negative k_{std} states.

To assess the statistical significance of the observed nonmonotonicities, we test with bootstrap replications the null hypothesis that the difference between the average pricing kernel at k_{std} and $k_{std} - 1$ is negative (i.e., the pricing kernel is decreasing in this region), against the alternative hypothesis that the difference is positive (i.e., the pricing kernel is increasing in this region), for each k_{std} among $-5, -4, \dots$, up to 6. Table 3 reports the results. For nearly all times of the day, we reject that the pricing kernel is decreasing from k_{std} levels -1 to 0 and from 0 to 1 . That is, the hump around the at-the-money region is statistically significant. We also reject monotonicity from 3 to 4 and 5 to 6 standardized return states for most of the times of the day. Again, nonmonotonicity for negative returns is concentrated in the extremely negative standardized return states.

4.5 ODTE (mis)pricing according to risk-averse investors

So far, we have shown that ODTE option prices can only be jointly reconciled with the physical distribution of market returns under a pricing kernel displaying pronounced nonmonotonicities. In this subsection, we address the problem from a different angle. For each option, we compute SSD bounds from the physical distribution consistent with all pricing kernels that are monotonically decreasing in market returns. In other words, we entertain the possibility that the option market is segmented and test whether, for each option, a risk-averse investor that is marginal in the market index and the risk-free rate would also be marginal in the option. As discussed in Section 2.5, a nonmonotonic pricing

kernel projection does not necessarily imply that option prices violate SSD bounds.

Table 4 reports, for each time of the day, the percentage over our sample of options on a given category that: satisfy the SSD bounds; violate the upper bound; or violate the lower bound. Patterns are similar across the day. Strikingly, only around 40% (30%) of the call (put) prices are consistent with monotonically decreasing pricing kernels. This is mainly driven by the ATM category, where only around 7% of the prices satisfy the SSD bounds. In particular, we observe mainly upper bound violations, meaning that ATM option prices are usually too high, in the sense that any risk-averse investor would improve expected utility by selling ATM options. This is consistent with the pronounced hump in the ATM region observed in the pricing kernel projection. On the other hand, OTM calls and puts almost never violate the upper bound, while their prices are below the lower bound reasonably often, i.e., they are generally too cheap from the perspective of these investors. ITM options frequently violate both the upper and lower bound, which can be due to the fact that they are less liquid and may present unreliable prices. Overall, 0DTE options are mostly “mispriced”, in that they do not reflect the risks implied by the time series of intra-day market returns under risk-averse risk preferences.

Figures OA.5 and OA.6 help visualize these patterns by plotting the average over our sample of the SSD bounds and the call and put prices, respectively, in IV space across standardized log-moneyness. Focusing first on the calls, it can be seen that average ATM option prices are above the upper bound, while OTM option prices are close to the lower bound, reflecting the patterns in violations discussed above. ITM option prices are on average within the bounds, which reflects the fact that upper and lower bound violations for this category occur in similar proportions. As for the puts, it is evident how both ATM and deep ITM options are too expensive on average relative to the bounds, while OTM put prices are close to the lower bound. An important takeaway from the figures is that the reason 0DTEs violate SSD bounds is not due to the existence of a smile of

the IV curve, as the monotonically decreasing pricing kernels also generate a smile.²⁵

To assess the economic significance of the mispricing we identify, we build a trading strategy that exploits the SSD bounds violations. The strategy focuses on the ATM option and is defined as follows. If the option is overpriced (underpriced) with respect to the underlying asset according to risk-averse investors – i.e., if the option represents an upper (lower) bound violation – we write (purchase) the option and buy (sell) delta shares of the underlying. If, instead, the option price is inside the SSD bounds, we go long on the risk-free rate. As a benchmark, we consider the strategy that always sells the ATM delta-hedged option, which amounts to exploiting the variance risk premium. The SSD violation strategy would be equivalent to this benchmark if the ATM option always violated the upper bound.

Table 5 reports the Sharpe ratio of the SSD violation strategy and the benchmark for different times of the day, before and after transaction costs. We focus on the ATM call, as results for the ATM put are very similar. To incorporate transaction costs, whenever we buy (sell) the option, we consider the ask (bid) price instead of the bid-ask midpoint, that is, we consider the worst case scenario for the strategy. As can be seen, the SSD violation strategy is highly profitable, yielding Sharpe ratios in the range of 0.2 to 0.3 for very short-term horizons. These Sharpe ratios are an order of magnitude larger than those obtained from selling the ATM delta-hedged call. In particular, transaction costs have a small impact on the performance of the SSD strategy, while the benchmark produces mostly negative Sharpe ratios net of costs.

Figure 11 further plots the cumulative returns of the strategies at two different times of the day, before and after transaction costs. Returns are adjusted to have unit stan-

²⁵This also illustrates why there is no inconsistency between options being “cheap” according to risk-averse investors and “expensive” in terms of high IVs as each criteria benchmark prices relative to those implied by different distributions: the former with respect to the physical distribution adjusted by monotonically decreasing pricing kernels and the latter to a log-normal distribution. Options can also be “expensive” in terms of low returns on average but “cheap” according to risk-averse investors if the SSD lower bound price is already enough to generate such low returns.

standard deviation to reflect risk-adjusted performance. The SSD violation strategy has a remarkably stable performance over time, while the benchmark is more erratic and experiences long periods with decreasing cumulative returns. This highlights the economic significance of the lower bound violations, which signal when the strategy should go long in the ATM delta-hedged call instead of going short. The profitability associated with the SSD violation strategy reinforces our finding that 0DTE options are mispriced. In fact, if our results were driven by misspecification of the estimated physical distribution, it is very unlikely that exploiting the violations would lead to such economic gains.

To shed further light on the economic significance of the mispricing we document, we compute Sharpe ratios after transaction costs conditioned to different variables being above the median (high) or below the median (low) in our sample.²⁶ Table 6 reports the results. The Sharpe ratio of the SSD violation strategy is higher when the 0DTE volume, realized variance and attention to the 0DTE option market (measured with Google trends) are low. In other words, the mispricing is more pronounced, in economic terms, when liquidity is low, uncertainty is low and agents are not paying attention to 0DTEs. On the other hand, conditioning on the *VRP* has only a small effect on the Sharpe ratio of the violation strategy, reinforcing that the strategy is not simply exploiting the variance risk premium. In fact, its risk-adjusted profitability is slightly larger when the *VRP* is low.

5 Robustness checks

In this section, we provide a number of robustness checks of our main results. More specifically, we investigate how our findings are affected if we: consider the sample before and after the daily availability of 0DTE options; remove days with FOMC announcements; use E-mini S&P 500 futures data instead of the cash index; and work with alternative estimates for the physical distribution. Tables and figures supporting this analysis

²⁶Figure OA.7 depicts the time-series of these variables together with their median.

are collected in the Online Appendix.

5.1 Before and after 2022

As previously mentioned, since the introduction of Thursday expirations in May 11 2022, 0DTE options on the S&P 500 index are available on a daily basis. This comes hand in hand with the largest increase in 0DTE trading volume in recent years, as seen in Figure 1. In this subsection, we investigate how our results are affected by splitting the sample before and after May 11 2022.

Table OA.3 reports the average VRP , VRP^+ and VRP^- on the two subsamples. The average variance risk premium is significantly positive in both cases, albeit larger before May 11 2022. Upon inspection of Panels B and C, we can see that this is mainly because the compensation for downside risk is larger before May 11 2022. In fact, after this date, the 25th percentile of VRP^- is already negative and statistical significance of its average is reduced. In contrast, the compensation for upside risk is similar for both subsamples, and actually slightly larger after May 11 2022.

Table OA.4 contains the market return predictability exercise with variance risk premium variables. Focusing first on the VRP , it negatively predicts market returns for both subsamples, but with much stronger statistical significance after 0DTEs are available on a daily basis. When decomposing the VRP into its two components, we can see that the VRP^+ negatively predicts the equity premium both before and after May 11 2022, with statistical significance for most times of the day. In contrast, VRP^- is insignificant before May 11 2022, while it significantly predicts market returns with a negative relation afterwards. This explains the stronger significance of the total variance risk premium in the most recent subsample.

Table OA.5 further reports the option price bounds results before and after May 11 2022. The proportion of calls inside the bounds increases to around 45% with the daily

availability of 0DTEs. This is mainly because bound violations for ATM and ITM calls decrease. Interestingly, lower bound violations for ATM calls become rare, as upper bound violations account for around 90% of violations. This indicates that ATM calls are most of the time expensive in the most recent subsample according to risk-averse investors. Results for puts are similar, with the exception that violations for OTM puts increase, which is mainly due to lower bound violations. That is, OTM puts are in general too cheap according to the stochastic dominance criterion.

Finally, Figures OA.8 and OA.9 depict the average pricing kernel before and after May 11 2022, respectively. While there seems to be less variation in the pricing kernel after the introduction of Thursday expirations, the pronounced hump in the ATM region is present in both subsamples. In sum, our results are largely similar in both subsamples, in that in both cases they are consistent with pronounced nonmonotonicities in the pricing kernel and a special compensation required by investors to bear upside risk.

5.2 FOMC announcements

0DTEs allow investors to hedge and make leveraged bets on specific intra-daily market movements and resolution of uncertainty, which is especially relevant for days with events that can affect financial markets. Announcements from Federal Open Market Committee (FOMC) meetings are arguably the most important kind of events happening during regular trading hours (around 14:00) that are relevant for 0DTEs. During our sample, there are 56 FOMC announcement days coinciding with dates for which 0DTEs were traded. Rather than analyzing effects in such a relatively small sample,²⁷ we investigate how our results are affected by excluding FOMC announcement dates from our sample.

²⁷For instance, Figure OA.10 plots average returns of ATM delta-hedged calls and straddles on FOMC announcement days. Before the announcement, average returns are more negative than for the whole sample, suggesting a high premium to be protected against variance risk. After the announcement and its resolution of uncertainty, average returns become positive, such that investors would not be willing to pay for protection against variance risk anymore. However, average returns are not statistically significant due to the small sample, making it difficult to draw conclusions.

Table OA.6 reports the average VRP , VRP^+ and VRP^- over our sample after removing days with FOMC announcements. The only change with respect to Table 1 is that the variance risk premium and its components are on average smaller once FOMC days are excluded. However, they are still economically large and statistically significant, and the VRP^+ continues to dominate the VRP^- from 11:30 onwards. That is, our conclusions regarding the compensation investors require to bear variance risk over the day are not affected by FOMC announcements.

Table OA.7 contains the market return predictability exercise with variance risk premium variables when removing FOMC announcement days. Focusing first on the VRP , it continues to negatively predict market returns, but with smaller statistical significance. In fact, the variance risk premium is only significant at 13:00. The VRP^+ also has its significance somewhat reduced compared to the total sample, but it is still significant for most times of the day in the regression including VRP^- . This indicates that FOMC announcement days are relevant for the predictive relation between compensation for variation risk and market returns, but do not solely account for this relation.

Tables OA.8 and OA.9 report the option price bounds results and the Sharpe ratios for the SSD violation strategy for days without FOMC announcements. Both tables show that these results are essentially unaffected by removing FOMC days. In other words, the stochastic dominance violations and the profitability associated with these violations are not explained by FOMC announcements. Finally, Figure OA.11 depicts the average pricing kernel over our sample without FOMC days, which is again largely the same as for the total sample.

5.3 E-mini S&P 500 futures

While the S&P 500 cash index is the underlying asset of the 0DTE options we analyze, it might be relevant to consider the E-mini S&P 500 futures instead for at least two

reasons. First, they are tradable assets that are often used by option market makers to hedge their 0DTE exposures. Second, most of the price discovery occurs in the E-mini market, that is, price movements and changes in market sentiment are first reflected in the E-mini futures before being incorporated into the cash index (see, e.g., [Hasbrouck, 2003](#)). Therefore, in this subsection, we analyze the implications of using E-mini futures returns to: compute the realized variance used in the variance risk premium measures, define the market returns used in the predictive regressions and perform the delta-hedge in the SSD violation trading strategy.

Table [OA.10](#) reports the average VRP , VRP^+ and VRP^- over our sample when using the E-mini futures. The main difference with respect to Table [1](#) is that the variance risk premium and its components are approximately halved on average, meaning that the realized variance measures are usually twice as large when computed using the E-mini futures returns instead of the cash index returns. Even so, the average VRP still obtains values as large as 5.44% annualized. Moreover, the main conclusions from this analysis remain: the variance risk premium and its components are positive and statistically significant, and the VRP^+ becomes greater than the VRP^- from 12:00 onwards.

Table [OA.11](#) reports results for predicting the same excess S&P 500 cash index returns as before, but using the variance risk premium measures computed from E-mini futures returns. Relative to Table [2](#), we find even stronger predictability coming from VRP and VRP^+ , as seen by mostly higher (absolute) t -statistics and R^2 s, while keeping the same interpretation as before as the signs of the coefficients are the same. Table [OA.12](#) further contains results for predicting the E-mini futures returns, instead of cash index returns, using the variance risk premium also computed from the futures returns. Again, we find overall stronger predictability compared to the original results in Table [2](#), with the same conclusion that the VRP negatively predicts market returns, which is driven by the strong negative relation between VRP^+ and future returns.

Finally, Table [OA.13](#) reports the Sharpe ratios associated with the SSD violation strategy when we use the E-mini futures to perform the delta-hedge. While theoretically the S&P 500 cash index should be used, it is more realistic to use the futures as they are directly tradable. As can be seen, the strategy exploiting violations of stochastic dominance has its Sharpe ratio substantially reduced when using futures to delta-hedge. Even so, before and after transaction costs, the Sharpe ratio is still always positive and considerably high given the short horizon, continuing to greatly outperform the benchmark strategy that always sells the delta-hedged ATM call. That is, our strategy is robust to using the futures in its implementation.

5.4 Alternative physical distributions

The estimation of the pricing kernel depends on the method used to estimate the conditional physical distribution from past market returns. In our main analysis, we take a conservative approach: we consider past returns (of the horizon of the option) over a long expanding window and make this distribution conditional by setting its volatility equal to the ATM implied volatility, which is forward-looking. In this subsection, we consider three alternatives instead: the same distribution but setting the volatility to be equal to the realized volatility of the horizon of the option in the previous day; a different distribution, in the spirit of [Jackwerth \(2000\)](#), given by the histogram of past returns over a short window of 2 months, with volatility unadjusted;²⁸ the latter distribution with volatility adjusted to be equal to the ATM implied volatility. For all these distributions on a given day, we impose the lower bound of 5% for the (annualized) equity premium and smooth the distribution with a Gaussian kernel.

Figures [OA.12](#), [OA.13](#) and [OA.14](#) plot the resulting average pricing kernel over our

²⁸[Jackwerth \(2000\)](#), focusing on the one-month horizon, uses a window of 4 years, consisting of past 48 non-overlapping monthly returns. We use 2 months, consisting of past 44 non-overlapping intra-daily returns of the horizon of the option.

sample using each of the methods above, respectively. As can be seen, the hump around the at-the-money region is a robust feature that is present in the three cases. Moreover, Figure OA.12 is similar to the main Figure 9, while Figures OA.13 and OA.14 display more variation in the pricing kernel with an overall W-shape. This indicates that using a short rolling window where higher-order physical moments vary more over time relative to a long expanding window leads to such a W-shape pattern. Importantly, any of these patterns is consistent with our empirical results as they all imply pricing kernels with high values for positive market return states, which we generally refer to as a U-shape.

6 Conclusion

We explore the asset pricing implications of the new 0DTE option market, which today accounts for around 45% of total S&P 500 option volume. These options contain new, valuable information about investors' risk preferences and risk premia over the intra-daily horizons for which the options expire. We extract this information from different perspectives and document a number of new asset pricing stylized facts.

First, most of the equity premium arises as compensation for market returns between -5% and 0%, where positive returns contribute negatively to the equity premium. Second, the average returns of calls, puts and different option strategies that are long in volatility are highly negative, which is consistent with the high variance risk premium we document over the day. Surprisingly, the total intra-day variance risk premium is mainly attributed to the compensation for bearing positive return variation risk. Third, the variance risk premium negatively predicts market returns over the day, which is mainly driven by the negative relation between future returns and the compensation for upside risk. Fourth, 0DTE options present severe violations of stochastic dominance bounds that are highly profitable if exploited in a trading strategy. 0DTEs can thus be seen as “mispriced” in

the sense that they do not reflect the risks implied by the time series of intra-day market returns under risk-averse preferences.

Our empirical results are all consistent with a strong U-shape in the intra-day pricing kernel as a function of market returns. We confirm that with direct estimates of the pricing kernel, which reveal pronounced and statistically significant nonmonotonicities, especially around the at-the-money region. While there is a large literature documenting pricing kernel nonmonotonicity at the one-month or longer horizons, we show that the implications of nonmonotonicity at the intra-daily horizons are vastly different, with a predominant role of compensation for upside risk relative to that for downside risk.

References

- Adams, G., Fontaine, J.-S., and Ornathanalai, C. (2024). The market for 0-days-to-expiration: The role of liquidity providers in volatility attenuation. *SSRN Working Paper*.
- Aït-Sahalia, Y., and Lo, A. W. (2000). Nonparametric risk management and implied risk aversion. *Journal of Econometrics*, 94(1-2), 9–51.
- Aleti, S., and Bollerslev, T. (2024). News and asset pricing: A high-frequency anatomy of the sdf. *Review of Financial Studies*, forthcoming.
- Aleti, S., Bollerslev, T., and Siggaard, M. (2023). Intraday market return predictability culled from the factor zoo. *SSRN Working Paper*.
- Alexiou, L., Bevilacqua, M., and Hizmeri, R. (2023). Uncovering the asymmetric information content of high-frequency options. *SSRN Working Paper*.
- Almeida, C., Ardison, K., Freire, G., Garcia, R., and Orlowski, P. (2023). High-frequency tail risk premium and stock return predictability. *Journal of Financial and Quantitative Analysis*, forthcoming.
- Almeida, C., and Freire, G. (2022). Pricing of index options in incomplete markets. *Journal of Financial Economics*, 144(1), 174–205.
- Almeida, C., and Freire, G. (2023). Demand in the option market and the pricing kernel. *SSRN Working Paper*.
- Almeida, C., Freire, G., Garcia, R., and Hizmeri, R. (2023). Tail risk and asset prices in the short-term. *SSRN Working Paper*.
- Andersen, T. G., Bollerslev, T., Diebold, F. X., and Labys, P. (2003). Modeling and forecasting realized volatility. *Econometrica*, 71(2), 579–625.
- Andersen, T. G., Fusari, N., and Todorov, V. (2015). The risk premia embedded in index options. *Journal of Financial Economics*, 117(3), 558–584.
- Andersen, T. G., Fusari, N., and Todorov, V. (2017). Short-term market risks implied


- by weekly options. *The Journal of Finance*, 72(3), 1335–1386.
- Ait-Sahalia, Y., Fan, J., Xue, L., and Zhou, Y. (2022). How and when are high-frequency stock returns predictable? *NBER Working Paper*, 30366.
- Baele, L., Driessen, J., Ebert, S., Londono, J. M., and Spalt, O. G. (2019). Cumulative prospect theory, option returns, and the variance premium. *Review of Financial Studies*, 32(9), 3667-3723.
- Bakshi, G., and Kapadia, N. (2003). Delta-hedged gains and the negative market volatility risk premium. *Review of Financial Studies*, 16(2), 527-566.
- Bakshi, G., Kapadia, N., and Madan, D. (2003). Stock return characteristics, skew laws, and the differential pricing of individual equity options. *Review of Financial Studies*, 16(1), 101-143.
- Bakshi, G., Madan, D., and Panayotov, G. (2010). Returns of claims on the upside and the viability of U-shaped pricing kernels. *Journal of Financial Economics*, 97(1), 130-154.
- Baltussen, G., Da, Z., Lammers, S., and Martens, M. (2021). Hedging demand and market intraday momentum. *Journal of Financial Economics*, 142(1), 377-403.
- Bandi, F. M., Fusari, N., and Renò, R. (2023). Otdte option pricing. *SSRN Working Paper*.
- Barone-Adesi, G., Engle, R., and Mancini, L. (2008). A garch option pricing model with filtered historical simulation. *Review of Financial Studies*, 21(3), 1223-1258.
- Beason, T., and Schreindorfer, D. (2022). Dissecting the equity premium. *Journal of Political Economy*, 130(8), 2203-2222.
- Beckmeyer, H., Branger, N., and Gayda, L. (2023). Retail traders love Otdte options... but should they? *SSRN Working Paper*.
- Black, F., and Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3), 637–654. doi: <http://www.jstor.org/stable/1831029>

- Bollerslev, T., Tauchen, G., and Zhou, H. (2009). Expected stock returns and variance risk premia. *Review of Financial Studies*, 22(11), 4463–4492.
- Bollerslev, T., Todorov, V., and Xu, L. (2015). Tail risk premia and return predictability. *Journal of Financial Economics*, 118(1), 113–134.
- Breeden, D. T., and Litzenberger, R. H. (1978). Prices of state-contingent claims implicit in option prices. *The Journal of Business*, 51(4), 621-651.
- Brogaard, J., Han, J., and Won, P. Y. (2023). How does zero-day-to-expiry options trading affect the volatility of underlying assets? *SSRN Working Paper*.
- Chong, C. H., and Todorov, V. (2024). Do equity and options markets agree about volatility? *SSRN Working Paper*.
- Constantinides, G. M., Jackwerth, J. C., and Perrakis, S. (2009). Mispricing of S&P 500 Index Options. *Review of Financial Studies*, 22(3), 1247-1277.
- Coval, J. D., and Shumway, T. (2001). Expected option returns. *The Journal of Finance*, 56(3), 983-1009.
- Cuesdeanu, H., and Jackwerth, J. C. (2018). The pricing kernel puzzle: survey and outlook. *Annals of Finance*, 14(3), 289-329.
- Dew-Becker, I., Giglio, S., and Kelly, B. (2021). Hedging macroeconomic and financial uncertainty and volatility. *Journal of Financial Economics*, 142(1), 23-45.
- Dim, C., Eraker, B., and Vilkov, G. (2024). Odttes: Trading, gamma risk and volatility propagation. *SSRN Working Paper*.
- Gao, L., Han, Y., Zhengzi Li, S., and Zhou, G. (2018). Market intraday momentum. *Journal of Financial Economics*, 129(2), 394-414.
- Gatheral, J. (2004). A parsimonious arbitrage-free implied volatility parameterization with application to the valuation of volatility derivatives. *Presentation at Global Derivatives*.
- Hasbrouck, J. (2003). Intraday price formation in u.s. equity index markets. *The Journal*

- of Finance*, 58(6), 2375-2399.
- Jackwerth, J. C. (2000). Recovering risk aversion from option prices and realized returns. *Review of Financial Studies*, 13(2), 433-451.
- Johannes, M., Kaeck, A., Seeger, N., and Shah, N. (2024). Expected 1dte option returns. *SSRN Working Paper*.
- Kilic, M., and Shaliastovich, I. (2019). Good and bad variance premia and expected returns. *Management Science*, 67(6), 2522-2544.
- Levy, H. (1985). Upper and lower bounds of put and call option value: Stochastic dominance approach. *The Journal of Finance*, 40(4), 1197-1217.
- Linn, M., Shive, S., and Shumway, T. (2018). Pricing kernel monotonicity and conditional information. *Review of Financial Studies*, 31(2), 493-531.
- Martin, I. (2017). What is the expected return on the market? *The Quarterly Journal of Economics*, 132(1), 367-433.
- Perrakis, S., and Ryan, P. J. (1984). Option pricing bounds in discrete time. *The Journal of Finance*, 39(2), 519-25.
- Ritchken, P. H. (1985). On Option Pricing Bounds. *The Journal of Finance*, 40(4), 1219-1233.
- Rosenberg, J. V., and Engle, R. F. (2002). Empirical pricing kernels. *Journal of Financial Economics*, 64(3), 341-372.
- Vilkov, G. (2023). 0dte trading rules. *SSRN Working Paper*.

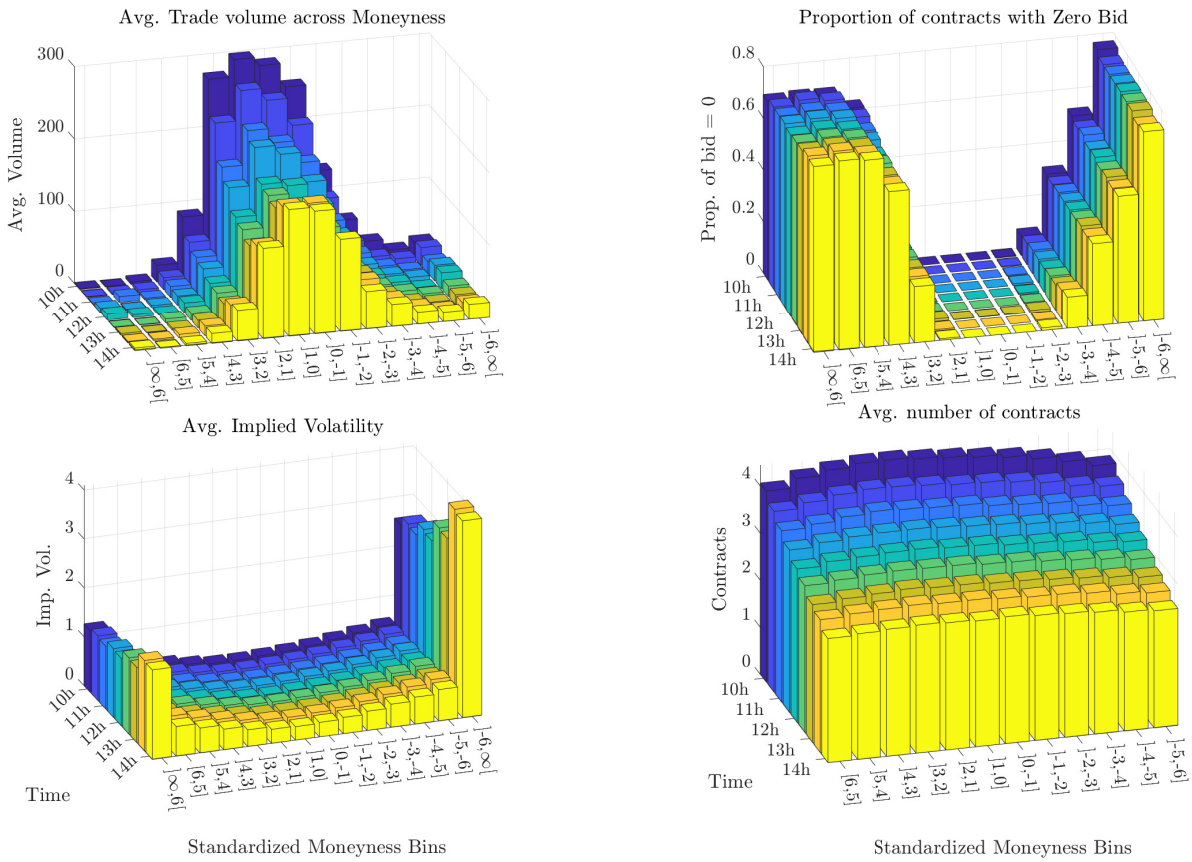
A Figures and tables

Figure 1: Yearly fraction of trading volume in 0DTE options




Note: The figure depicts the yearly fraction of trading volume in 0DTE S&P 500 options relative to the entire S&P 500 option market. The sample ranges from January 6 2012 to July 3 2023.

Figure 2: Descriptive statistics of raw 0DTE option data



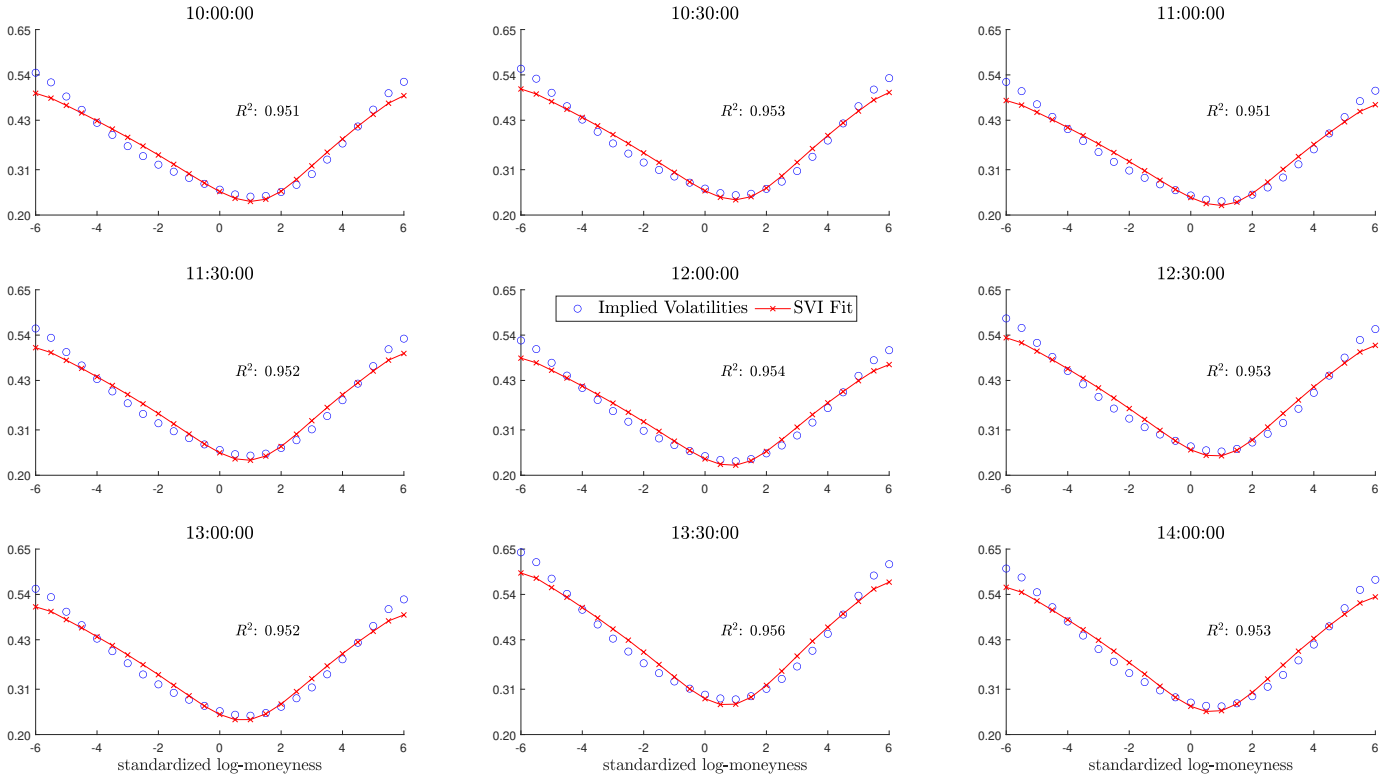
Note: The figure plots, in four subplots, for different times of the day and different standardized log-moneyness bins, the average trading volume, proportion of contracts with zero bid, average implied volatility and average number of strikes. The sample ranges from January 6 2012 to July 3 2023.

Figure 3: Trading volume and relative bid-ask spread of 0DTE options over the day



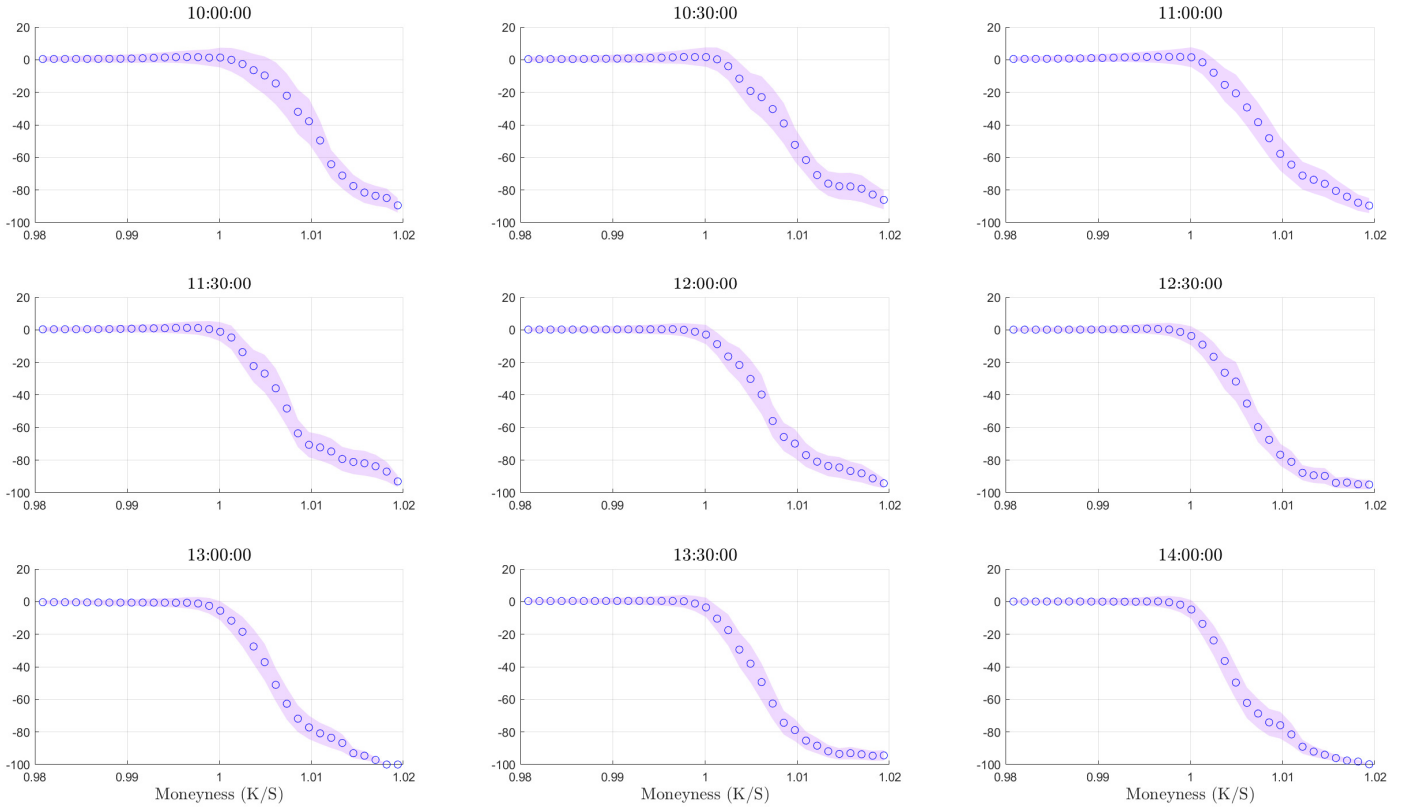
Note: The left plot of this figure depicts the time series average of the trading volume in terms of number of contracts of 0DTE S&P 500 options over the day. The right plot depicts the time series average of the average relative bid-ask spread across all 0DTE options over the day, where the relative spread is computed as $(Ask - Bid)/MidQuote$. The sample ranges from January 6 2012 to July 3 2023.

Figure 4: SVI fit for 0DTE options



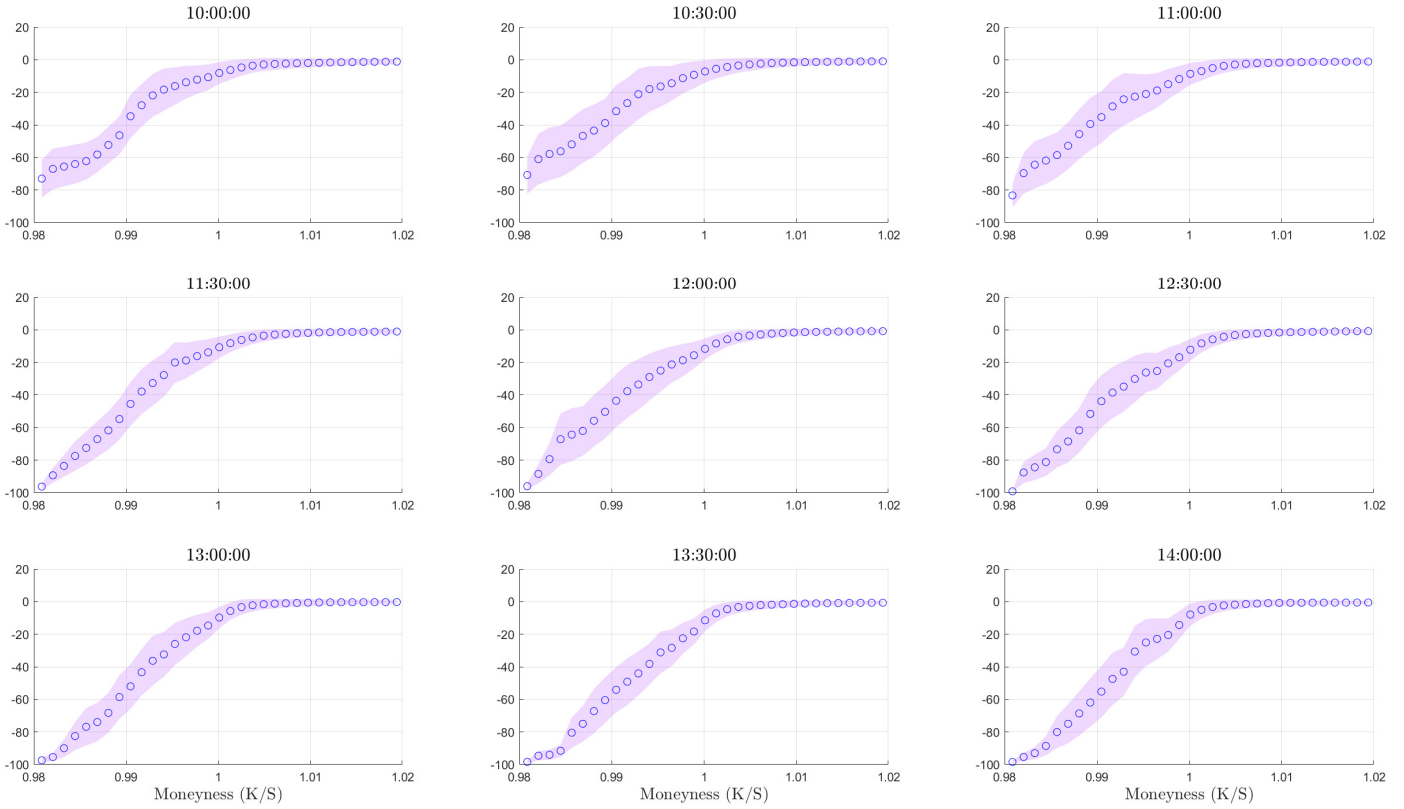
Note: The figure plots the average over our sample of the observed 0DTE IVs and the fitted IVs using the SVI method for different times of the day. The average OLS R^2 fit is also reported. Standardized log-moneyness is defined as $\frac{\log(K/S_t)}{\sigma_{BS}(0)\sqrt{\tau}}$.

Figure 5: 0DTE average call returns across strikes




Note: The figure plots the average returns (in %) of call options with different strikes together with 90% confidence bands. For a given moneyness, option returns over the time-series are winsorized in the right-tail at 0.5%, that is, returns above the 99.5% quantile are set equal to that quantile. Confidence bands are based on 2,500 bootstrap replications. The sample ranges from January 6 2012 to July 3 2023.

Figure 6: 0DTE average put returns across strikes



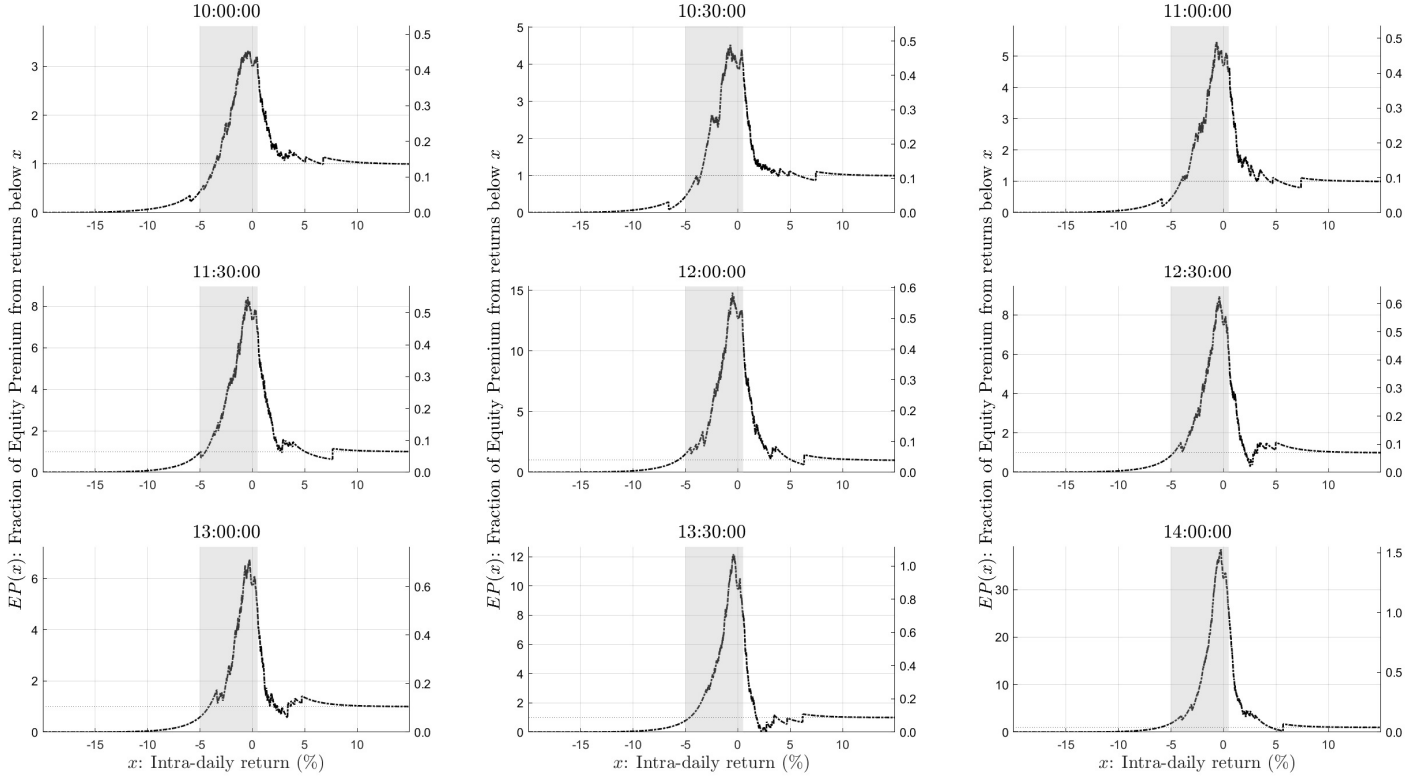
Note: The figure plots the average returns (in %) of put options with different strikes together with 90% confidence bands. For a given moneyness, option returns over the time-series are winsorized in the right-tail at 0.5%, that is, returns above the 99.5% quantile are set equal to that quantile. Confidence bands are based on 2,500 bootstrap replications. The sample ranges from January 6 2012 to July 3 2023.

Figure 7: Average returns of 0DTE option strategies



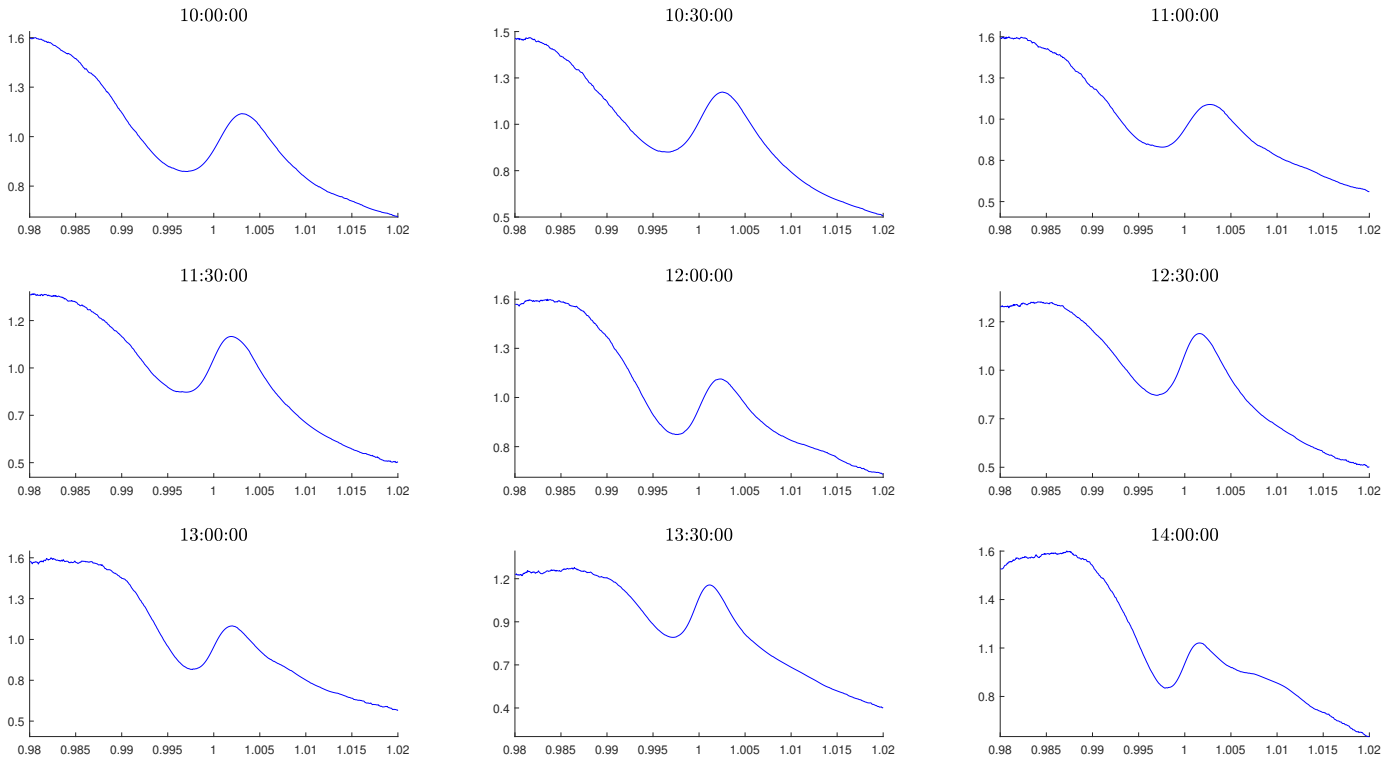
Note: The figure plots, for different times of the day, average returns together with 90% confidence bands for ATM delta-hedged calls, delta-neutral straddles and simple straddles. Confidence bands are based on 2,500 bootstrap replications. The sample ranges from January 6 2012 to July 3 2023.

Figure 8: Intra-day equity premium decomposition



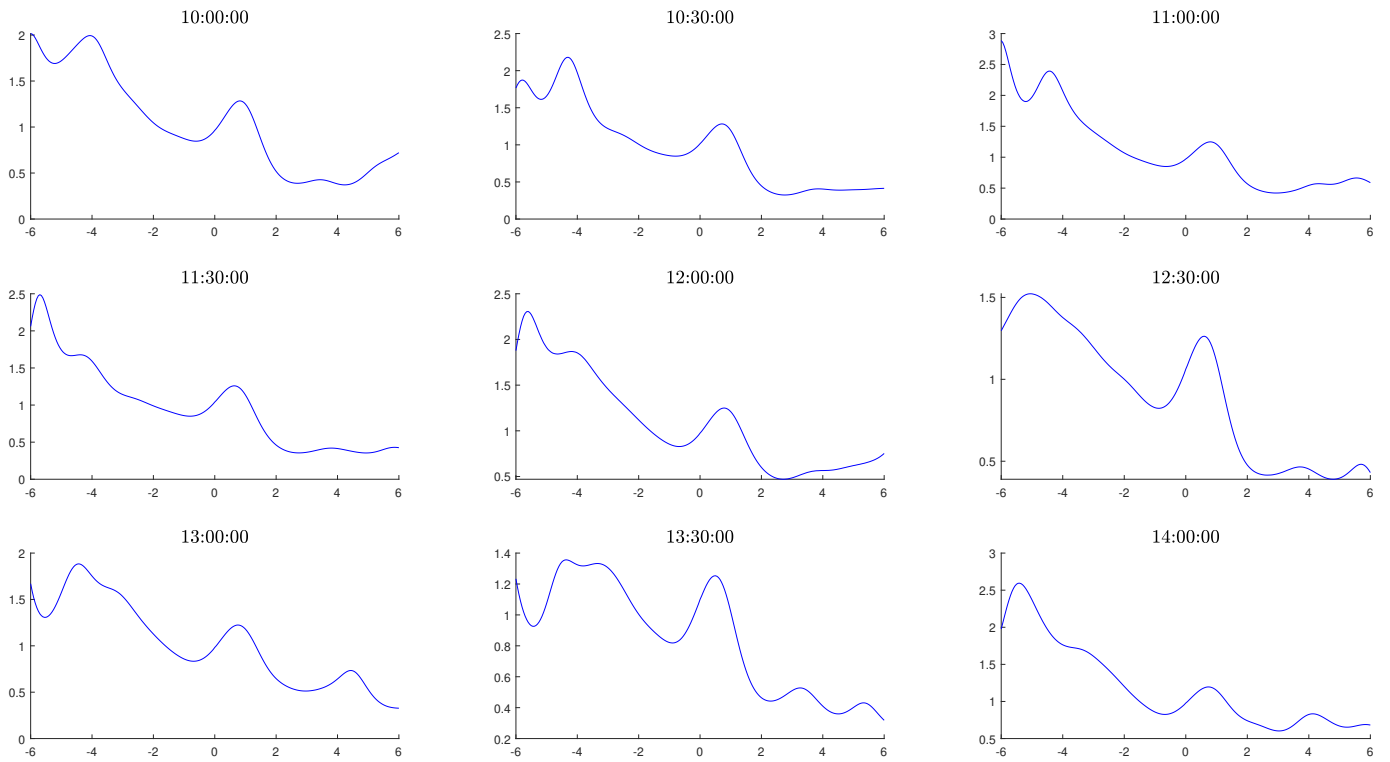
Note: The figure plots, for different times of the day, the intra-day equity premium decomposition implied by 0DTE options. The left vertical axis displays the fraction of the equity premium stemming from market returns below x , while the right vertical axis displays the total annualized equity premium that would be implied from market returns up to x . The sample ranges from January 6 2012 to July 3 2023.

Figure 9: Average intra-day pricing kernels over return states



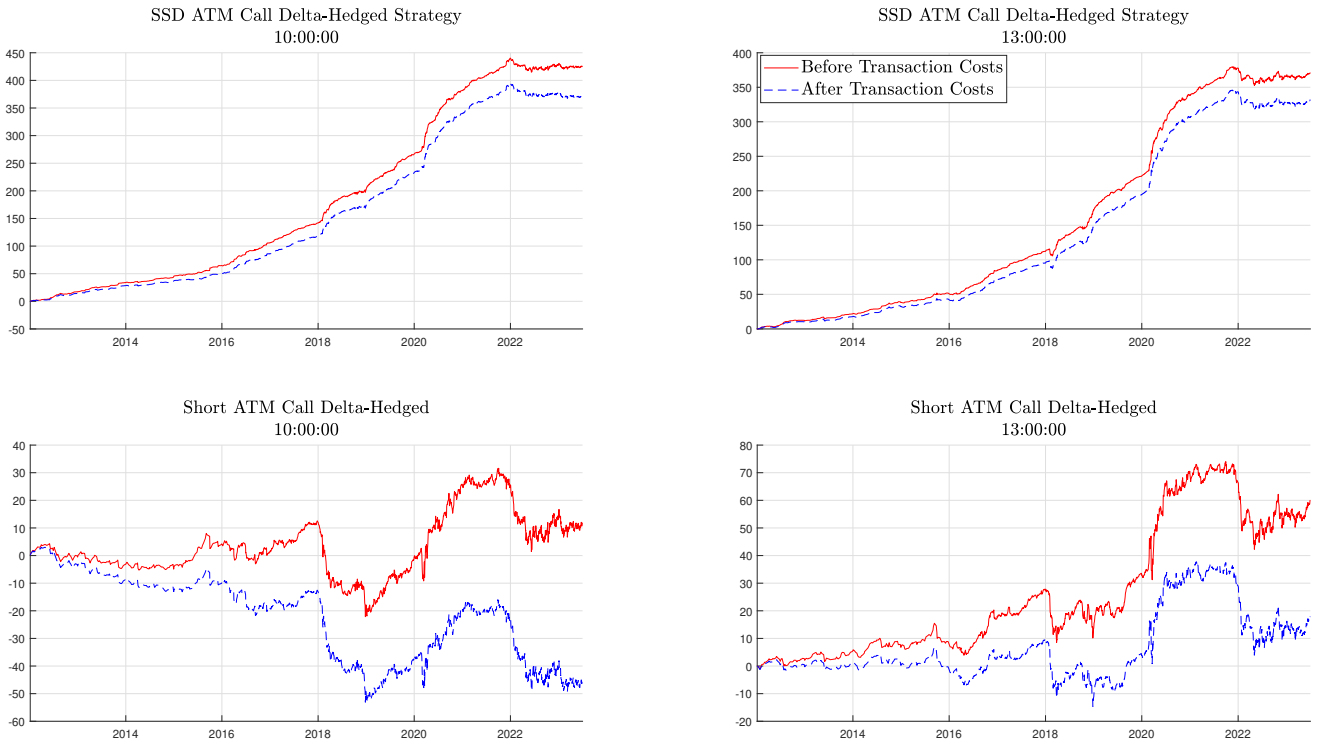
Note: The figure plots, for different times of the day, the average of the pricing kernel, as a function of market returns, over time. The sample ranges from January 6 2012 to July 3 2023.

Figure 10: Average intra-day pricing kernels over standardized log-moneyness



Note: The figure plots, for different times of the day, the average of the pricing kernel, as a function of standardized log-moneyness, over time. The sample ranges from January 6 2012 to July 3 2023.

Figure 11: Cumulative returns of SSD violation strategy



Note: The figure plots, for two different times of the day, the cumulative returns of the SSD violation strategy based on the ATM call (in the upper panels) and the benchmark strategy writing the ATM call delta-hedged (in the lower panels). To incorporate transaction costs, whenever we buy (sell) the option, we consider the ask (bid) price instead of the bid-ask midpoint, that is, we consider the worst case scenario for the strategy. Returns are adjusted to have standard deviation of one. The sample ranges from January 6 2012 to July 3 2023.

Table 1: Intra-day variance risk premium

	10:00:00	10:30:00	11:00:00	11:30:00	12:00:00	12:30:00	13:00:00	13:30:00	14:00:00
Panel A: VRP									
Mean	3.275	4.023	2.879	4.214	3.085	5.275	4.209	8.005	6.388
St. Dev.	10.141	11.716	9.381	11.332	10.595	15.658	17.667	24.041	23.191
25th Percentile	0.853	1.087	0.779	1.092	0.810	1.308	0.984	1.929	1.456
75th Percentile	4.078	4.705	3.870	5.124	4.129	5.818	4.736	7.988	6.475
p -value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Panel B: VRP^+									
Mean	1.459	1.889	1.394	2.110	1.614	2.789	2.259	4.278	3.483
St. Dev.	4.757	5.505	4.671	5.515	5.315	7.615	8.679	11.463	10.448
25th Percentile	0.570	0.711	0.603	0.808	0.673	0.999	0.811	1.507	1.273
75th Percentile	1.935	2.353	1.988	2.654	2.262	3.250	2.789	4.583	3.851
p -value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Panel C: VRP^-									
Mean	1.816	2.134	1.485	2.104	1.456	2.486	1.946	3.727	2.905
St. Dev.	5.928	6.720	5.247	6.288	5.732	8.492	9.493	13.096	13.738
25th Percentile	0.301	0.324	0.169	0.278	0.111	0.282	0.123	0.416	0.213
75th Percentile	2.191	2.424	1.893	2.402	1.875	2.651	2.042	3.506	2.734
p -value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Note: The table reports, for each time of the day, summary statistics of the $VRP_{t,T}$, $VRP_{t,T}^+$ and $VRP_{t,T}^-$ over our sample. The variance risk premium measures are annualized and expressed in percentage points. The p -values for the test with null hypothesis that the mean is smaller than or equal to zero, against the alternative that the mean is positive, are implemented using bootstrapped standard errors with 2,500 replications. The sample ranges from January 6 2012 to July 3 2023.

Table 2: Predicting excess market returns with variance risk premium

	10:00:00				10:30:00				11:00:00			
VRP	-0.031				-0.015				-0.047*			
t -stat	-0.643				-0.442				-1.776			
VRP^+	-0.050		-0.110*		-0.038		-0.139*		-0.068**		-0.132**	
t -stat	-1.258		-1.657		-1.244		-1.820		-2.400		-2.112	
VRP^-	-0.013		0.075		0.005		0.121		-0.023		0.081	
t -stat	-0.238		0.786		0.120		1.297		-0.807		1.212	
R^2 (%)	0.160	0.417	0.027	0.758	0.045	0.287	0.005	1.171	0.501	1.059	0.122	1.622
	11:30:00				12:00:00				12:30:00			
VRP	-0.046**				-0.054**				-0.026			
t -stat	-2.087				-2.029				-0.607			
VRP^+	-0.071***		-0.187***		-0.074***		-0.156**		-0.049		-0.218**	
t -stat	-3.892		-2.821		-2.850		-2.105		-1.367		-2.413	
VRP^-	-0.020		0.137*		-0.031		0.099		-0.004		0.190*	
t -stat	-0.649		1.759		-1.004		1.243		-0.075		1.720	
R^2 (%)	0.542	1.307	0.105	2.717	0.851	1.608	0.278	2.504	0.200	0.723	0.004	3.034
	13:00:00				13:30:00				14:00:00			
VRP	-0.081***				-0.026				-0.032			
t -stat	-2.657				-0.526				-0.793			
VRP^+	-0.097***		-0.203***		-0.045		-0.238**		-0.042		-0.075	
t -stat	-3.765		-2.716		-1.166		-2.166		-1.123		-1.471	
VRP^-	-0.062*		0.119		-0.008		0.211		-0.023		0.040	
t -stat	-1.727		1.389		-0.136		1.540		-0.550		0.639	
R^2 (%)	2.137	3.084	1.245	4.051	0.250	0.771	0.022	3.451	0.500	0.838	0.247	1.076

Note: The table reports, for each time of the day, the results from different predictive regressions over our sample using $VRP_{t,T}$, $VRP_{t,T}^+$ and $VRP_{t,T}^-$ to predict the excess market return from t to T . Regressors are standardized to have mean zero and unit variance. We compute the t -statistics using Newey-West robust standard errors with a lag length equal to 5. We denote with *, **, and *** significance at the 10%, 5% and 1% level, respectively. The sample ranges from January 6 2012 to July 3 2023.

Table 3: Intra-day pricing kernel nonmonotonicity

		-5	-4	-3	-2	-1	0	1	2	3	4	5	6
10:00:00	Mean	-0.296	0.270***	-0.576	-0.368	-0.172	0.089***	0.285***	-0.731	-0.116	-0.020	0.126***	0.213***
	St. Dev.	0.587	0.343	0.443	0.283	0.210	0.191	0.341	0.428	0.389	0.196	0.213	0.253
10:30:00	Mean	-0.108	0.314***	-0.753	-0.211	-0.154	0.162***	0.179***	-0.748	-0.114	0.071***	-0.010	0.020***
	St. Dev.	0.460	0.352	0.543	0.239	0.213	0.208	0.375	0.429	0.265	0.191	0.235	0.177
11:00:00	Mean	-0.914	0.098***	-0.657	-0.344	-0.193	0.095***	0.239***	-0.641	-0.148	0.123***	0.040***	0.004
	St. Dev.	1.128	0.452	0.589	0.296	0.248	0.215	0.382	0.470	0.347	0.324	0.352	0.261
11:30:00	Mean	-0.318	-0.151	-0.450	-0.153	-0.132	0.183***	0.095***	-0.674	-0.097	0.050***	-0.059	0.071***
	St. Dev.	0.504	0.429	0.356	0.197	0.223	0.234	0.407	0.442	0.217	0.134	0.209	0.122
12:00:00	Mean	0.033***	-0.058	-0.396	-0.338	-0.265	0.121	0.235***	-0.610	-0.115	0.080***	0.054***	0.133***
	St. Dev.	0.364	0.649	0.552	0.341	0.271	0.225	0.415	0.481	0.368	0.414	0.395	0.370
12:30:00	Mean	0.226***	-0.144	-0.182	-0.195	-0.174	0.235***	0.046***	-0.631	-0.052	0.028***	-0.056	0.034***
	St. Dev.	0.252	0.408	0.349	0.262	0.257	0.250	0.425	0.481	0.267	0.241	0.205	0.139
13:00:00	Mean	-0.097	0.180***	-0.221	-0.406	-0.270	0.124***	0.197***	-0.532	-0.132	0.131***	-0.115	-0.203
	St. Dev.	0.472	0.340	0.338	0.325	0.299	0.231	0.428	0.512	0.324	0.229	0.287	0.305
13:30:00	Mean	-0.160	0.250***	-0.020	-0.301	-0.184	0.279***	-0.081	-0.556	0.047***	-0.100	-0.008	-0.083
	St. Dev.	0.407	0.204	0.264	0.229	0.245	0.261	0.438	0.454	0.234	0.179	0.133	0.101
14:00:00	Mean	0.394***	-0.606	-0.154	-0.408	-0.346	0.121***	0.170***	-0.408	-0.136	0.223***	-0.157	0.013***
	St. Dev.	0.332	0.548	0.299	0.281	0.329	0.239	0.441	0.545	0.308	0.217	0.327	0.194

Note: The table reports the mean and standard deviation of the difference between the pricing kernel at k_{std} and $k_{std} - 1$, for k_{std} among $-5, -4, \dots, 6$. We denote with *, **, and *** significance at the 10%, 5% and 1% level, respectively, of the test of the null hypothesis that the difference in the average pricing kernel at k_{std} and $k_{std} - 1$ is negative (i.e., the pricing kernel is decreasing in this region), against the alternative hypothesis that the difference is positive (i.e., the pricing kernel is increasing in this region). The test is based on 2,500 bootstrap replications. The sample ranges from January 6 2012 to July 3 2023.

Table 4: Option price bounds for 0DTE options

		10:00:00	10:30:00	11:00:00	11:30:00	12:00:00	12:30:00	13:00:00	13:30:00	14:00:00
Panel A: Calls										
All Calls	In	37.406	40.729	40.814	40.158	35.509	34.420	37.863	43.902	34.852
	Upper	28.467	26.889	26.373	27.621	26.697	27.666	27.537	24.152	26.378
	Lower	34.127	32.382	32.813	32.221	37.795	37.914	34.600	31.947	38.771
OTM	In	60.819	64.409	65.970	61.882	57.971	54.621	58.203	67.859	54.723
	Upper	9.110	8.444	8.643	12.748	7.622	9.392	7.824	3.504	5.279
	Lower	30.071	27.147	25.387	25.370	34.407	35.987	33.973	28.637	39.998
ATM	In	7.715	8.378	8.557	9.925	7.126	7.709	8.862	9.200	5.204
	Upper	64.834	65.432	65.017	63.927	67.019	68.082	69.377	68.169	69.596
	Lower	27.451	26.190	26.425	26.148	25.855	24.209	21.761	22.632	25.199
ITM	In	18.239	23.645	21.911	25.306	18.374	20.247	24.632	28.929	22.608
	Upper	36.193	31.663	29.792	28.083	31.229	30.786	32.021	28.708	32.184
	Lower	45.568	44.692	48.298	46.611	50.397	48.967	43.347	42.363	45.208
Panel B: Puts										
All Puts	In	24.595	29.450	29.561	32.180	27.719	30.562	35.369	40.521	29.170
	Upper	35.515	33.962	34.106	33.876	36.096	34.050	31.874	30.102	33.778
	Lower	39.891	36.588	36.333	33.945	36.185	35.388	32.758	29.377	37.052
OTM	In	42.950	50.125	49.237	55.027	47.552	52.695	60.898	66.601	49.033
	Upper	5.460	5.165	4.750	4.496	6.784	5.930	5.017	3.930	6.951
	Lower	51.590	44.710	46.013	40.477	45.663	41.375	34.084	29.469	44.016
ATM	In	7.296	8.024	7.881	8.904	5.550	6.344	8.338	9.676	4.812
	Upper	69.558	70.597	72.181	72.136	77.160	75.485	72.324	71.270	75.515
	Lower	23.146	21.379	19.939	18.961	17.290	18.171	19.338	19.054	19.673
ITM	In	8.882	12.391	14.295	13.195	12.613	13.034	14.446	20.800	14.395
	Upper	57.746	53.529	53.382	53.201	53.450	49.794	46.536	43.525	48.123
	Lower	33.373	34.080	32.324	33.604	33.938	37.172	39.018	35.675	37.482

Note: The table reports, for each time of the day and for each class of options, the percentage of options over our sample for which prices fall within the SSD bounds (In), above the SSD upper bound (Upper) and below the SSD lower bound (Lower). The OTM put (ITM call), ATM and ITM put (OTM call) categories are defined as standardized log-moneyness below -1 , between -1 and 1 , and above 1 , respectively. The sample ranges from January 6 2012 to July 3 2023.

Table 5: Sharpe ratios for SSD violation strategy

		10:00:00	10:30:00	11:00:00	11:30:00	12:00:00	12:30:00	13:00:00	13:30:00	14:00:00
Panel A: Before Transaction Costs										
Short ATM Call Delta-Hedge		0.008	0.006	0.004	0.013	0.006	0.035	0.043	0.026	0.018
SSD ATM Call Delta-Hedge		0.301	0.293	0.255	0.271	0.264	0.270	0.264	0.318	0.334
Panel B: After Transaction Costs										
Short ATM Call Delta-Hedge		-0.032	-0.023	-0.023	-0.017	-0.022	0.006	0.013	-0.007	-0.030
SSD ATM Call Delta-Hedge		0.262	0.257	0.230	0.249	0.238	0.233	0.238	0.277	0.290

Note: The table reports, for each time of the day, before and after transaction costs, the Sharpe ratio associated with the SSD violation strategy based on the ATM call and the benchmark strategy writing the ATM call delta-hedged. To incorporate transaction costs, whenever we buy (sell) the option, we consider the ask (bid) price instead of the bid-ask midpoint, that is, we consider the worst case scenario for the strategy. The sample ranges from January 6 2012 to July 3 2023.

Table 6: Sharpe ratios for SSD violation strategy conditional on variables

	10:00:00	10:30:00	11:00:00	11:30:00	12:00:00	12:30:00	13:00:00	13:30:00	14:00:00
High 0DTE log-Volume	0.141	0.124	0.108	0.110	0.121	0.126	0.128	0.123	0.150
Low 0DTE log-Volume	0.261	0.278	0.255	0.280	0.259	0.264	0.259	0.304	0.295
High RV	0.133	0.141	0.123	0.137	0.142	0.139	0.133	0.171	0.171
Low RV	0.358	0.354	0.350	0.359	0.333	0.333	0.347	0.336	0.369
High VRP	0.179	0.173	0.141	0.144	0.145	0.153	0.168	0.201	0.198
Low VRP	0.201	0.201	0.213	0.243	0.226	0.208	0.187	0.201	0.236
High Google Trend Index	0.037	0.066	0.043	0.061	0.062	0.056	0.036	0.057	0.065
Low Google Trend Index	0.307	0.271	0.256	0.271	0.248	0.246	0.267	0.304	0.309

Note: The table reports, for each time of the day, after transaction costs, the Sharpe ratio associated with the SSD violation strategy based on the ATM call, computed on days with a high (above median) or low (below median) value of a conditioning variable. To incorporate transaction costs, whenever we buy (sell) the option, we consider the ask (bid) price instead of the bid-ask midpoint, that is, we consider the worst case scenario for the strategy. 0DTE log-volume is the daily S&P 500 0DTE option log-volume, RV is the average of the realized variances we compute over the different times of the day, VRP is the average of the VRPs we compute over the different times of the day, and Google Trend Index is a sentiment index created using Google trends with the keyword “0DTE SPX options”. The total sample ranges from January 6 2012 to July 3 2023.

Online Appendix to 0DTE Asset Pricing

Caio Almeida[†]

Gustavo Freire[‡]

Rodrigo Hizmeri[§]

October 18, 2024

Abstract

This online appendix collects additional empirical results supporting the main paper. More specifically, Section [OA.1](#) contains evidence complementing the main analysis, while Sections [OA.2](#), [OA.3](#), [OA.4](#) and [OA.5](#) report robustness results by splitting the sample before and after 2022, removing FOMC announcement days, working with E-mini S&P 500 futures and considering alternative physical distributions, respectively.

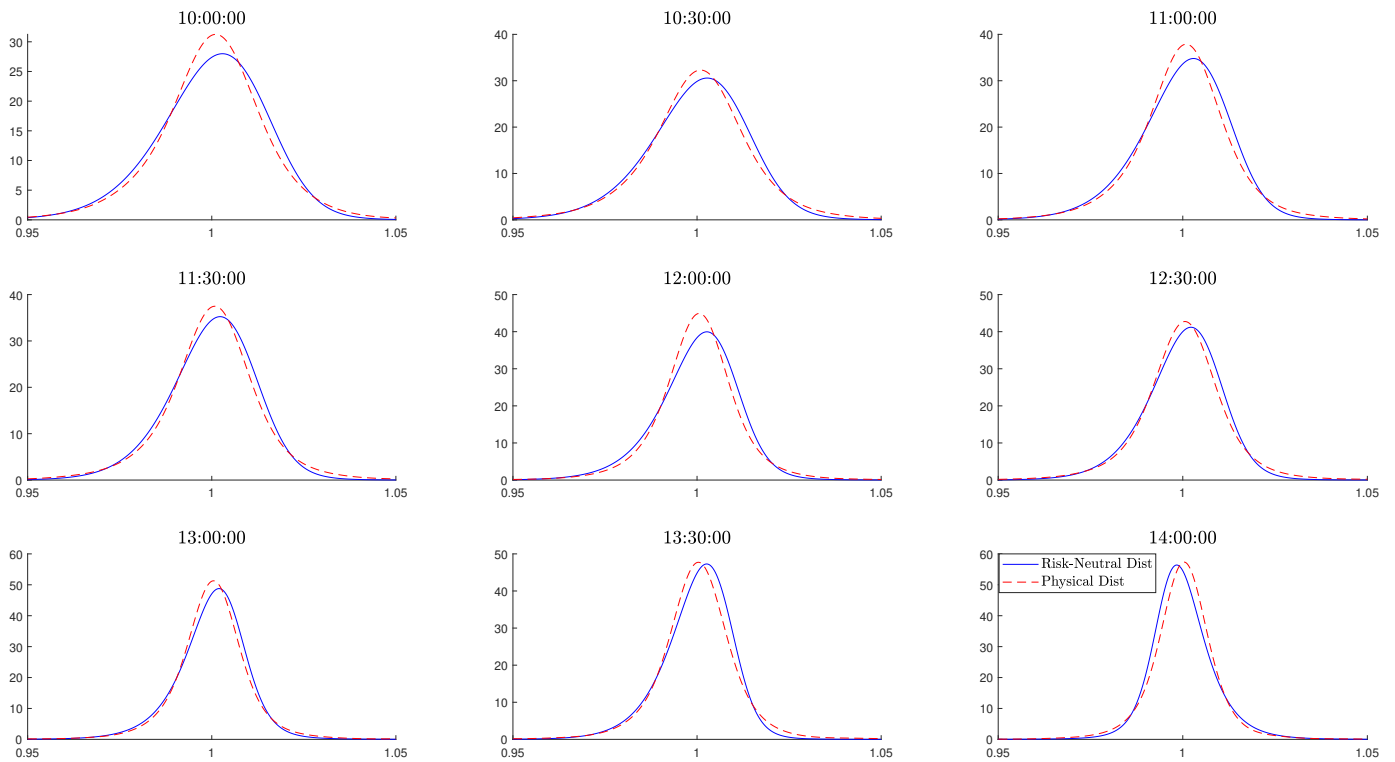
[†]E-mail: calmeida@princeton.edu, Department of Economics, Princeton University.

[‡]E-mail: freire@ese.eur.nl, Erasmus School of Economics - Erasmus University Rotterdam, Tinbergen Institute and Erasmus Research Institute of Management (ERIM).

[§]E-mail: r.hizmeri@liverpool.ac.uk, University of Liverpool Management School, University of Liverpool.

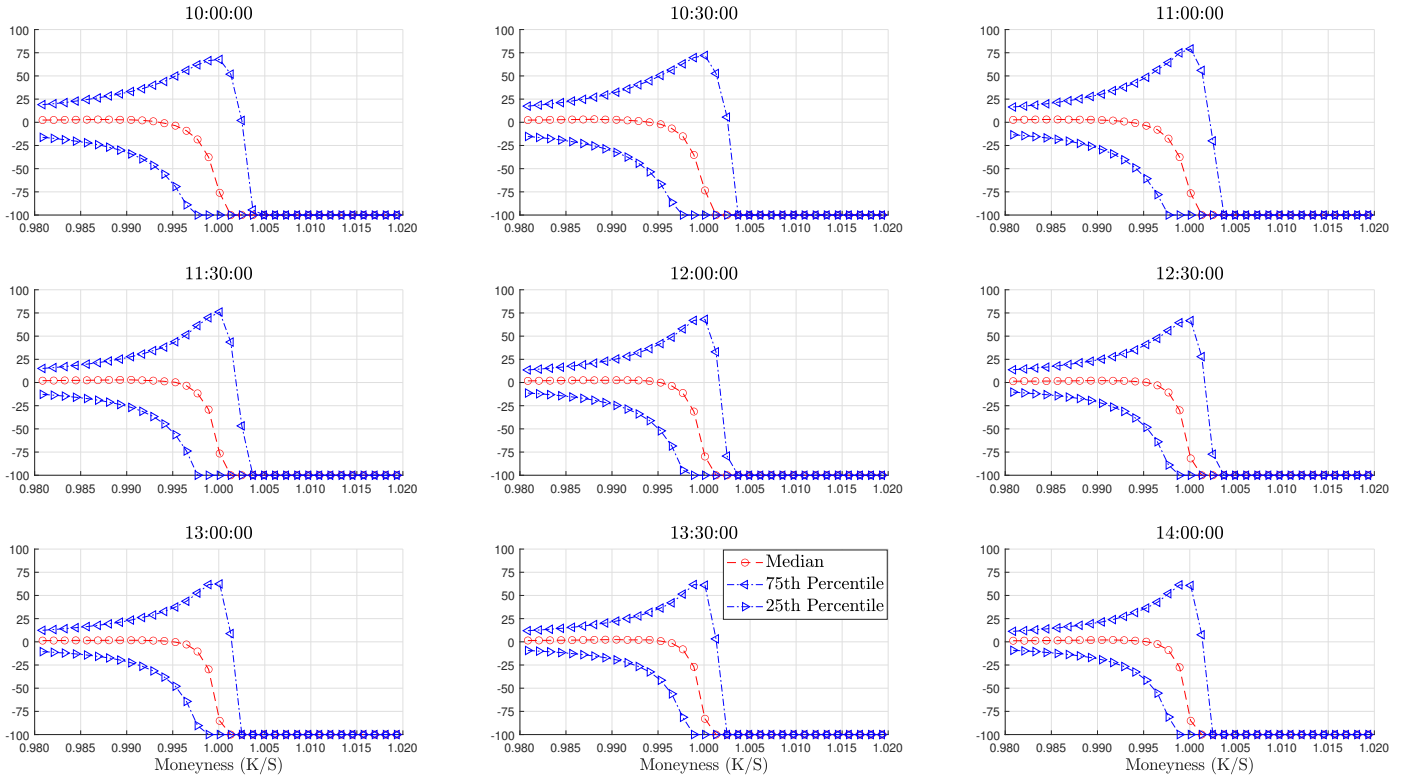
OA.1 Main analysis - additional results

Figure OA.1: Market return risk-neutral and physical distributions over intra-daily horizons



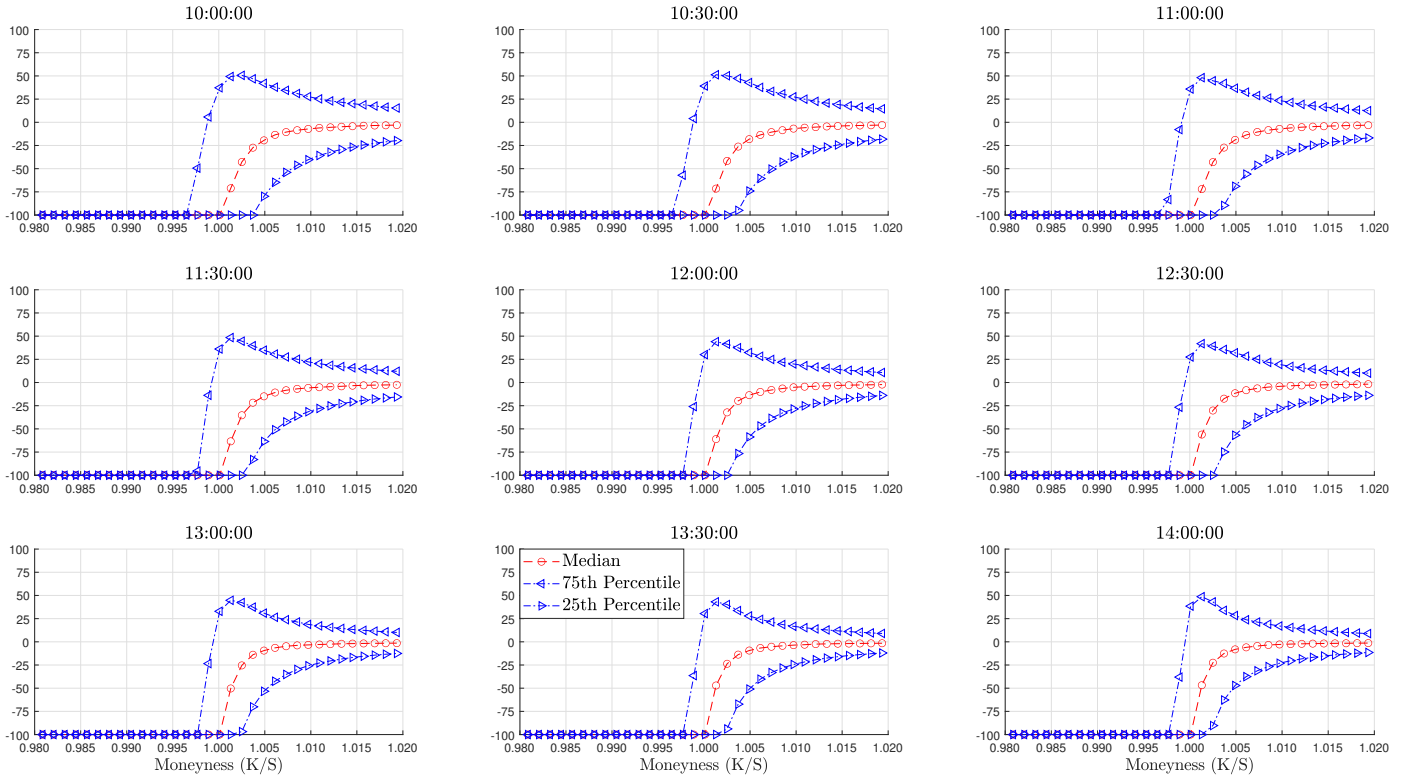
Note: The figure plots, for a representative date of our sample and for different times of the day, the market return risk-neutral distribution and physical distribution estimated as described in Section 3.3 and 3.4, respectively. The horizontal axis represents gross market return states.

Figure OA.2: 0DTE call returns across strikes



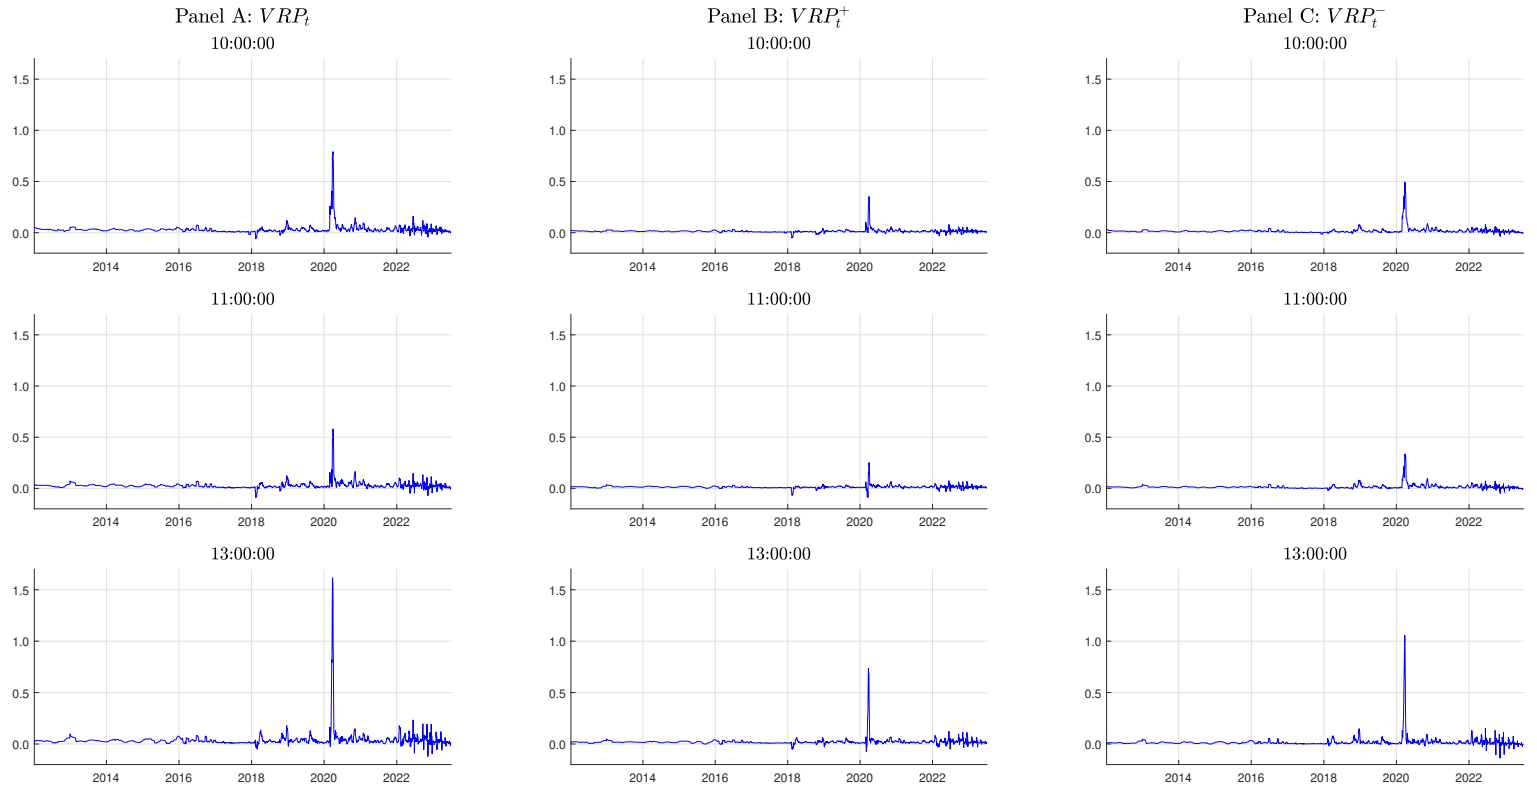
Note: The figure plots the 25th, median and 75th percentiles of returns (in %) of call options with different strikes over our sample. The sample ranges from January 6 2012 to July 3 2023.

Figure OA.3: 0DTE put returns across strikes



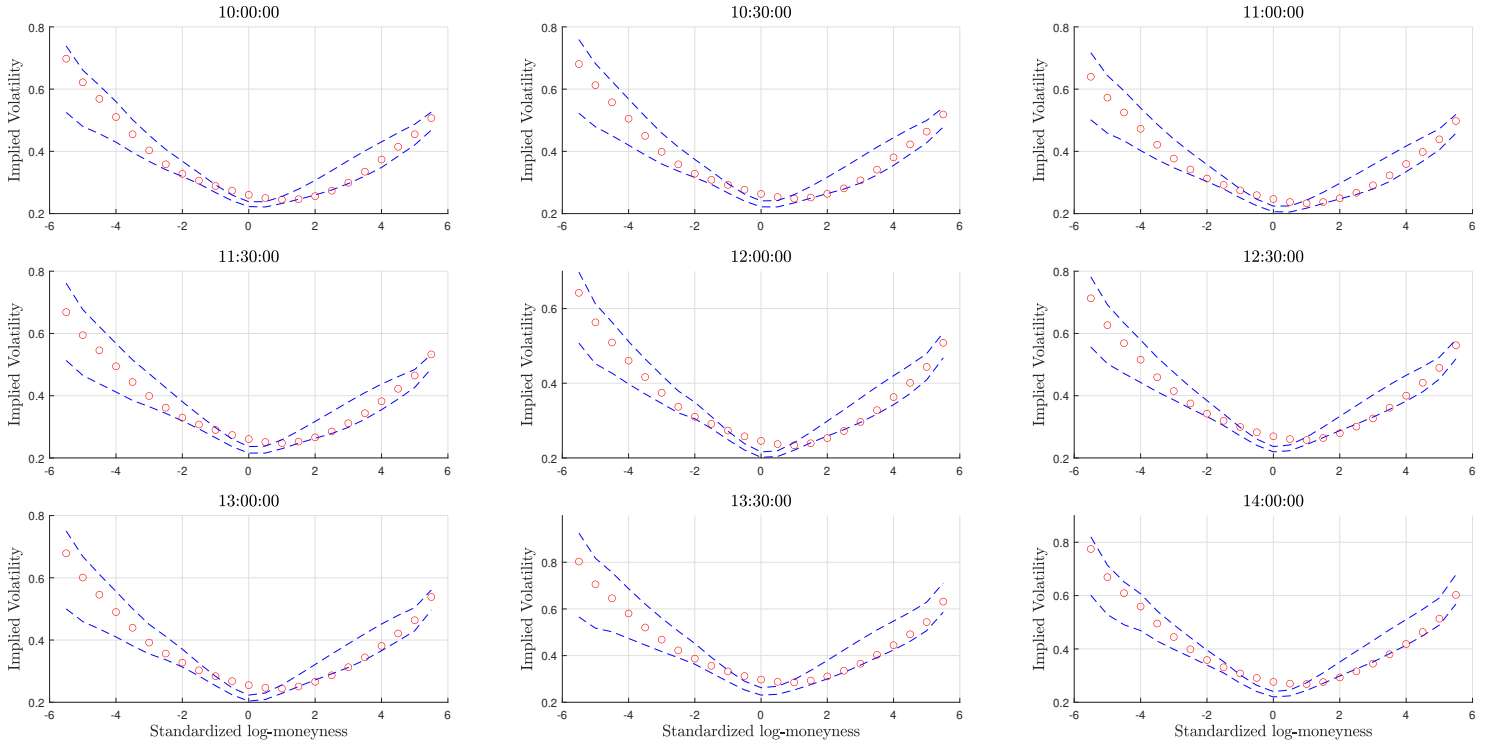
Note: The figure plots the 25th, median and 75th percentiles of returns (in %) of put options with different strikes over our sample. The sample ranges from January 6 2012 to July 3 2023.

Figure OA.4: Intra-day variance risk premium over time



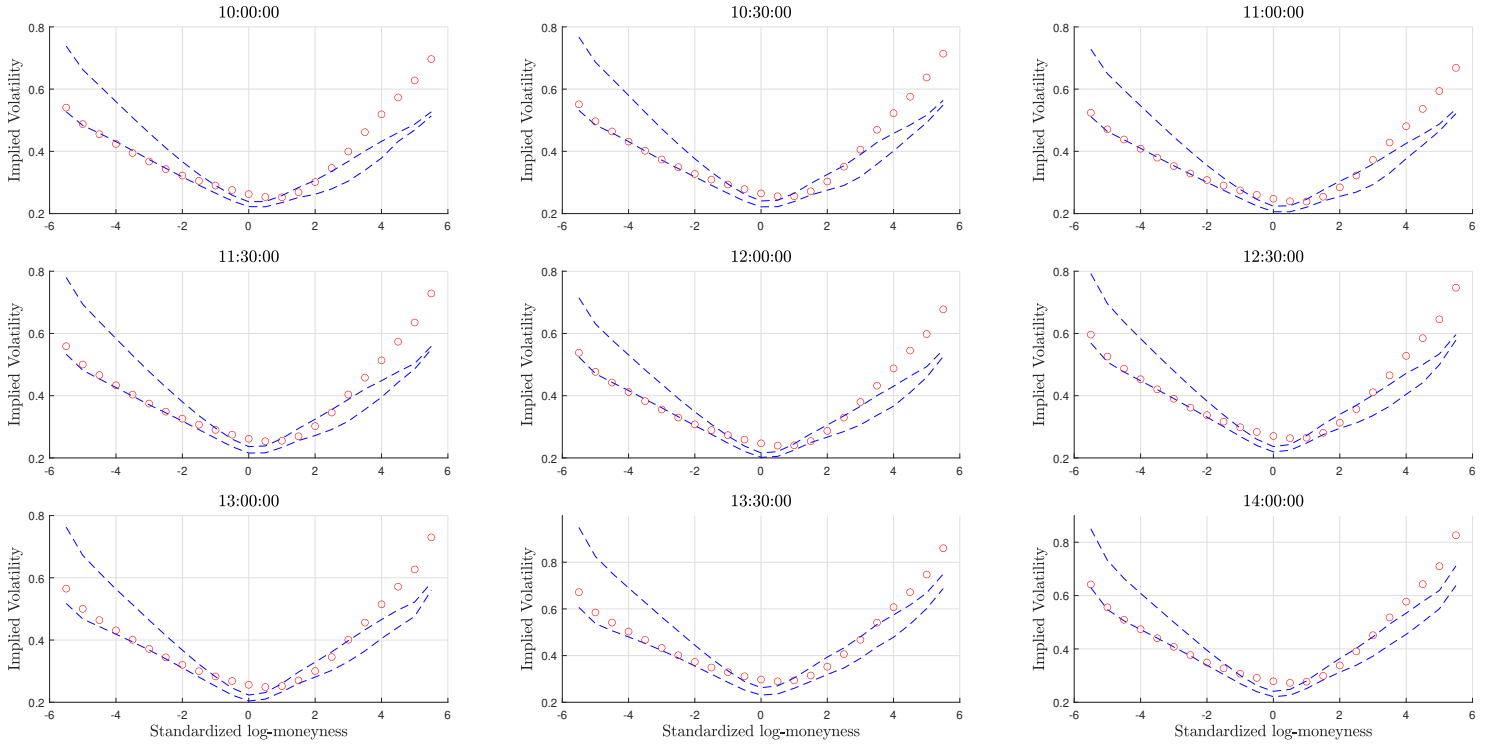
Note: The figure plots, for different times of the day, the one-week moving average (for ease of visualization) of the $VRP_{t,T}$, $VRP_{t,T}^+$ and $VRP_{t,T}^-$ over time. The measures are annualized. The sample ranges from January 6 2012 to July 3 2023.

Figure OA.5: Average SSD bounds and prices for 0DTE call options



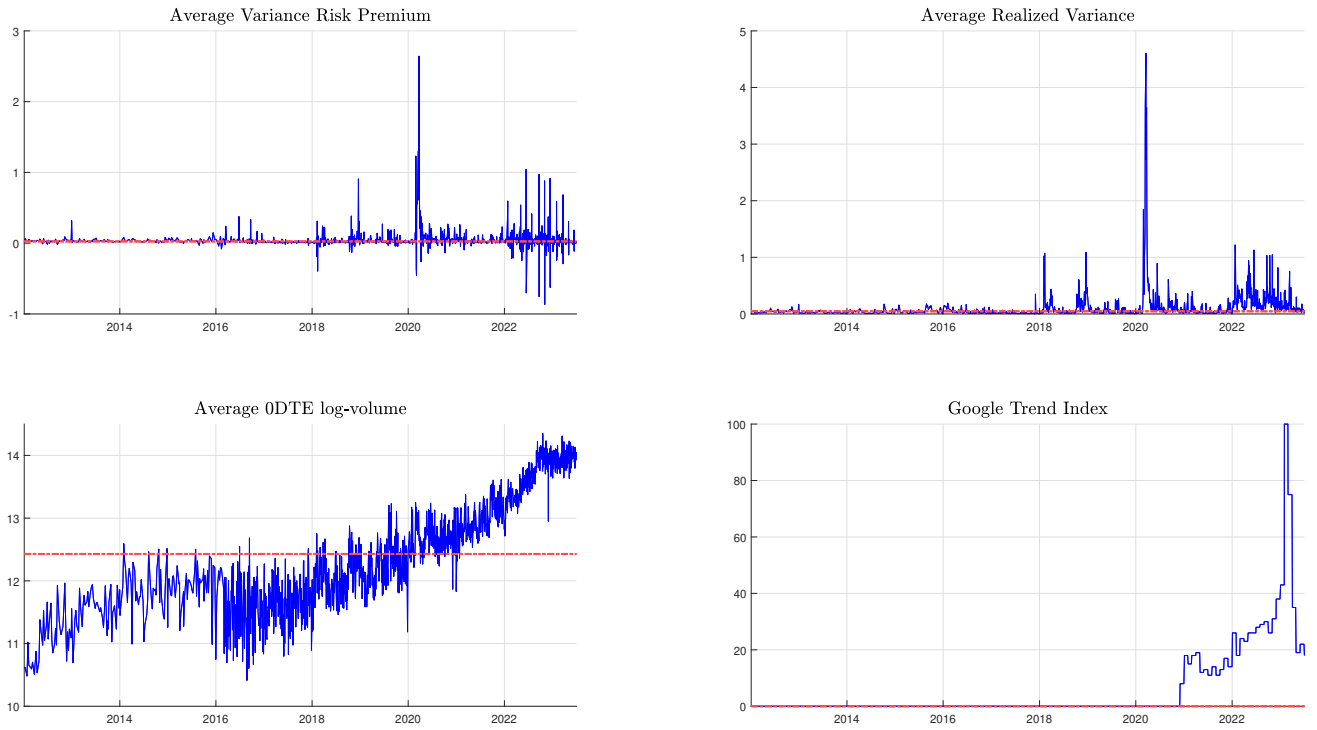
Note: The figure plots, for different times of the day, the average of the SSD option price bounds (in dashed blue) against the average option prices (in red circles) for calls, in implied volatility space, over standardized log-moneyness. The sample ranges from January 6 2012 to July 3 2023.

Figure OA.6: Average SSD bounds and prices for 0DTE put options



Note: The figure plots, for different times of the day, the average of the SSD option price bounds (in dashed blue) against the average option prices (in red circles) for puts, in implied volatility space, over standardized log-moneyness. The sample ranges from January 6 2012 to July 3 2023.

Figure OA.7: Conditioning variables



Note: The figure plots, for each of the conditioning variables in Table 6, their time-series together with a red dashed line indicating their median. The sample ranges from January 6 2012 to July 3 2023.

Table OA.1: Predicting excess market returns with risk measures

	10:00:00	10:30:00	11:00:00	11:30:00	12:00:00	12:30:00	13:00:00	13:30:00	14:00:00
<i>RV</i>	0.022	0.050	0.036	0.046	0.025	0.028	0.036	0.060	0.052
<i>t</i> -stat	0.383	1.180	0.856	1.129	0.535	0.572	0.740	1.272	1.334
R^2 (%)	0.082	0.505	0.294	0.546	0.188	0.237	0.426	1.366	1.312
<i>MFIS</i>	-0.005	0.000	-0.008	-0.013	-0.011	-0.004	-0.008	0.000	0.004
<i>t</i> -stat	-0.251	-0.007	-0.469	-0.896	-0.710	-0.287	-0.551	-0.038	0.327
R^2 (%)	0.003	0.000	0.014	0.045	0.033	0.006	0.020	0.000	0.008
<i>MFIK</i>	0.003	0.009	0.009	0.007	0.003	-0.015	-0.011	-0.018	-0.013
<i>t</i> -stat	0.164	0.456	0.553	0.431	0.187	-0.953	-0.737	-1.068	-0.949
R^2 (%)	0.002	0.017	0.020	0.013	0.002	0.069	0.042	0.118	0.084
<i>SVIX</i>	0.001	0.028	0.006	0.009	-0.010	-0.001	-0.034	0.010	0.005
<i>t</i> -stat	0.017	0.606	0.169	0.229	-0.206	-0.019	-0.725	0.157	0.088
R^2 (%)	0.000	0.154	0.009	0.021	0.029	0.000	0.373	0.035	0.012

Note: The table reports, for each time of the day, the results from univariate predictive regressions over our sample using *RV*, *MFIS* (risk-neutral skewness), *MFIK* (risk-neutral kurtosis) and *SVIX* to predict the excess market return from t to T . Regressors are standardized to have mean zero and unit variance. We compute the t -statistics using Newey-West robust standard errors with a lag length equal to 5. We denote with *, **, and *** significance at the 10%, 5% and 1% level, respectively. The sample ranges from January 6 2012 to July 3 2023.

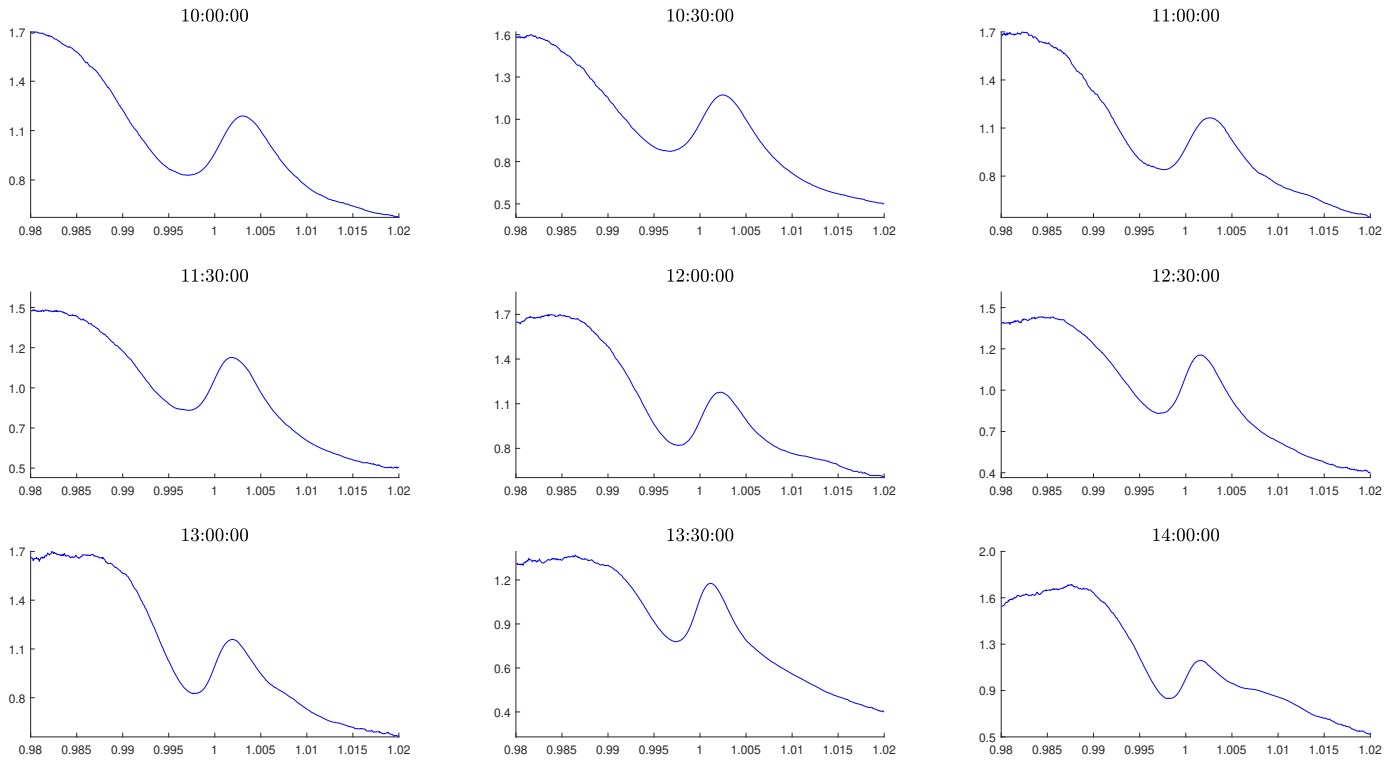
Table OA.2: Predicting excess market returns with variance risk premium and controls

	11:00:00		11:30:00		12:00:00		13:00:00	
VRP	−0.071***		−0.098***		−0.070***		−0.167***	
t -stat	−2.588		−3.840		−3.001		−3.459	
VRP^+		−0.136**		−0.186***		−0.161**		−0.214***
t -stat		−2.050		−2.644		−2.108		−2.717
VRP^-		0.084		0.121		0.098		0.066
t -stat		1.120		1.261		1.287		0.971
$MFIS$	−0.002	0.018	−0.009	0.016	−0.010	0.013	−0.011	0.007
t -stat	−0.117	0.832	−0.505	0.820	−0.544	0.689	−0.615	0.394
$MFIK$	0.016	0.021	0.011	0.016	−0.003	0.008	−0.013	−0.006
t -stat	0.706	0.940	0.515	0.837	−0.129	0.379	−0.656	−0.306
$SVIX$	0.050	0.006	0.081*	0.025	0.030	0.012	0.102*	0.074
t -stat	0.995	0.119	1.739	0.514	0.564	0.270	1.774	1.572
R_{adj}^2 (%)	0.585	1.343	1.123	2.481	0.768	2.217	3.090	4.329

Note: The table reports, for each time of the day, the results from multivariate predictive regressions over our sample based on $VRP_{t,T}$, $VRP_{t,T}^+$, $VRP_{t,T}^-$, RV , $MFIS$ (risk-neutral skewness), $MFIK$ (risk-neutral kurtosis) and $SVIX$ to predict the excess market return from t to T . Regressors are standardized to have mean zero and unit variance. We compute the t -statistics using Newey-West robust standard errors with a lag length equal to 5. We denote with *, **, and *** significance at the 10%, 5% and 1% level, respectively. The sample ranges from January 6 2012 to July 3 2023.

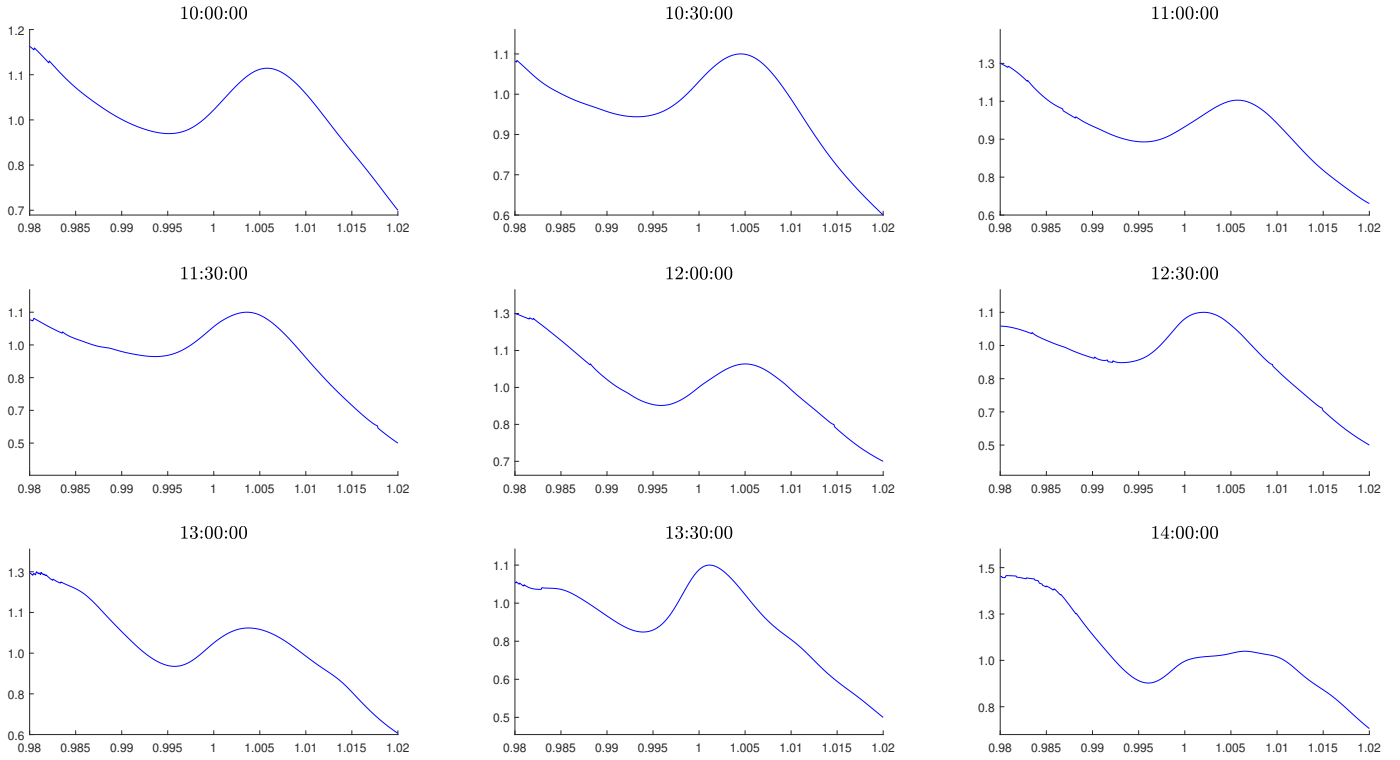
OA.2 Robustness - Before and after 2022

Figure OA.8: Average intra-day pricing kernels - before May 11 2022



Note: The figure plots, for different times of the day, the average of the pricing kernel, as a function of market returns, over time before May 11 2022. The sample begins on January 6 2012.

Figure OA.9: Average intra-day pricing kernels - after May 11 2022



Note: The figure plots, for different times of the day, the average of the pricing kernel, as a function of market returns, over time after May 11 2022. The sample ends on July 3 2023.

Table OA.3: Intra-day variance risk premium - before and after 2022

	10:00:00	11:00:00	12:00:00	13:00:00	14:00:00	10:00:00	11:00:00	12:00:00	13:00:00	14:00:00
	Before May 11, 2022					After May 11, 2022				
Panel A: VRP										
Mean	3.419	2.987	3.179	4.405	6.724	2.705	2.449	2.710	3.430	5.051
St. Dev.	10.282	8.854	9.529	17.107	20.701	9.557	11.241	14.067	19.748	31.186
25th Percentile	0.927	0.828	0.854	1.032	1.534	-0.230	0.271	0.256	0.322	0.462
75th Percentile	3.810	3.545	3.852	4.443	6.169	5.556	5.277	4.892	5.459	8.183
p -value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Panel B: VRP^+										
Mean	1.447	1.356	1.555	2.222	3.468	1.506	1.545	1.849	2.403	3.541
St. Dev.	4.736	4.399	4.865	8.439	9.005	4.845	5.631	6.820	9.585	14.877
25th Percentile	0.607	0.606	0.677	0.823	1.283	0.120	0.555	0.610	0.719	1.164
75th Percentile	1.825	1.813	2.063	2.533	3.659	2.981	2.818	3.011	3.608	4.983
p -value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Panel C: VRP^-										
Mean	1.972	1.632	1.606	2.183	3.256	1.199	0.904	0.861	1.004	1.510
St. Dev.	6.096	5.032	5.146	9.182	12.805	5.171	6.000	7.608	10.604	16.901
25th Percentile	0.352	0.224	0.161	0.186	0.274	-0.383	-0.369	-0.511	-0.576	-0.640
75th Percentile	2.060	1.773	1.846	2.049	2.641	2.723	2.344	1.947	2.032	2.995
p -value	0.000	0.000	0.000	0.000	0.000	0.000	0.005	0.029	0.058	0.061

Note: The table reports, for each time of the day, summary statistics of the $VRP_{t,T}$, $VRP_{t,T}^+$, and $VRP_{t,T}^-$ over our sample before May 11 2022 and after. The variance risk premium measures are annualized and expressed in percentage points. The p -values for the test with null hypothesis that the mean is smaller than or equal to zero, against the alternative that the mean is positive, are implemented using bootstrapped standard errors with 2,500 replications. The total sample ranges from January 6 2012 to July 3 2023.


Table OA.5: Option price bounds for 0DTE options - before and after 2022

		10:00:00	11:00:00	12:00:00	13:00:00	14:00:00	10:00:00	11:00:00	12:00:00	13:00:00	14:00:00
		Before May 11, 2022					After May 11, 2022				
Panel A: Calls											
All Calls	In	34.417	35.708	32.302	33.300	30.882	43.845	51.889	42.581	47.703	43.261
	Upper	30.090	28.806	28.127	30.671	26.969	24.970	21.097	23.542	20.779	25.125
	Lower	35.493	35.486	39.571	36.029	42.149	31.185	27.015	33.877	31.518	31.614
OTM	In	60.815	63.938	57.832	57.743	55.418	60.828	70.450	58.280	59.206	53.185
	Upper	12.323	11.668	10.349	10.838	6.553	1.953	1.976	1.610	1.258	2.459
	Lower	26.862	24.394	31.820	31.419	38.029	37.219	27.574	40.110	39.536	44.355
ATM	In	6.408	7.335	6.950	8.486	5.332	10.318	11.033	7.488	9.626	4.943
	Upper	54.869	54.668	55.597	60.278	57.919	84.678	85.984	90.528	87.843	93.558
	Lower	38.723	37.997	37.453	31.237	36.750	5.004	2.983	1.984	2.531	1.499
ITM	In	8.080	9.169	8.553	11.051	7.964	40.090	50.056	41.017	54.581	52.225
	Upper	43.846	39.687	38.816	43.124	40.073	19.733	7.934	13.736	7.536	16.228
	Lower	48.075	51.144	52.632	45.825	51.964	40.177	42.010	45.247	37.884	31.547
Panel B: Puts											
All Puts	In	27.250	31.253	29.969	35.800	29.428	19.160	26.126	23.194	34.492	28.626
	Upper	33.764	33.244	35.812	31.444	34.061	39.099	35.857	36.667	32.747	33.181
	Lower	38.987	35.503	34.219	32.756	36.511	41.741	38.018	40.138	32.761	38.193
OTM	In	53.224	59.097	57.300	67.835	55.742	22.083	28.896	27.271	46.469	34.836
	Upper	7.122	6.005	8.417	6.266	8.897	2.083	2.163	3.389	2.419	2.832
	Lower	39.653	34.898	34.283	25.900	35.360	75.834	68.941	69.340	51.111	62.332
ATM	In	6.270	7.283	5.464	8.213	4.571	9.337	9.091	5.728	8.592	5.306
	Upper	60.910	64.003	69.536	63.999	66.759	86.761	88.742	92.832	89.198	93.439
	Lower	32.820	28.714	25.000	27.788	28.670	3.902	2.167	1.440	2.210	1.256
ITM	In	3.076	5.136	4.679	4.965	5.075	21.110	32.485	27.616	33.091	34.358
	Upper	55.181	54.298	55.518	48.963	51.764	63.148	51.561	49.538	41.765	40.323
	Lower	41.743	40.566	39.802	46.073	43.161	15.741	15.954	22.847	25.144	25.318

Note: The table reports, for each time of the day and for each class of options, the percentage of options over our sample before May 11 2022 and after for which prices fall within the SSD bounds (In), above the SSD upper bound (Upper) and below the SSD lower bound (Lower). The OTM put (ITM call), ATM and ITM put (OTM call) categories are defined as standardized log-moneyness below -1 , between -1 and 1 , and above 1 , respectively. The total sample ranges from January 6 2012 to July 3 2023.

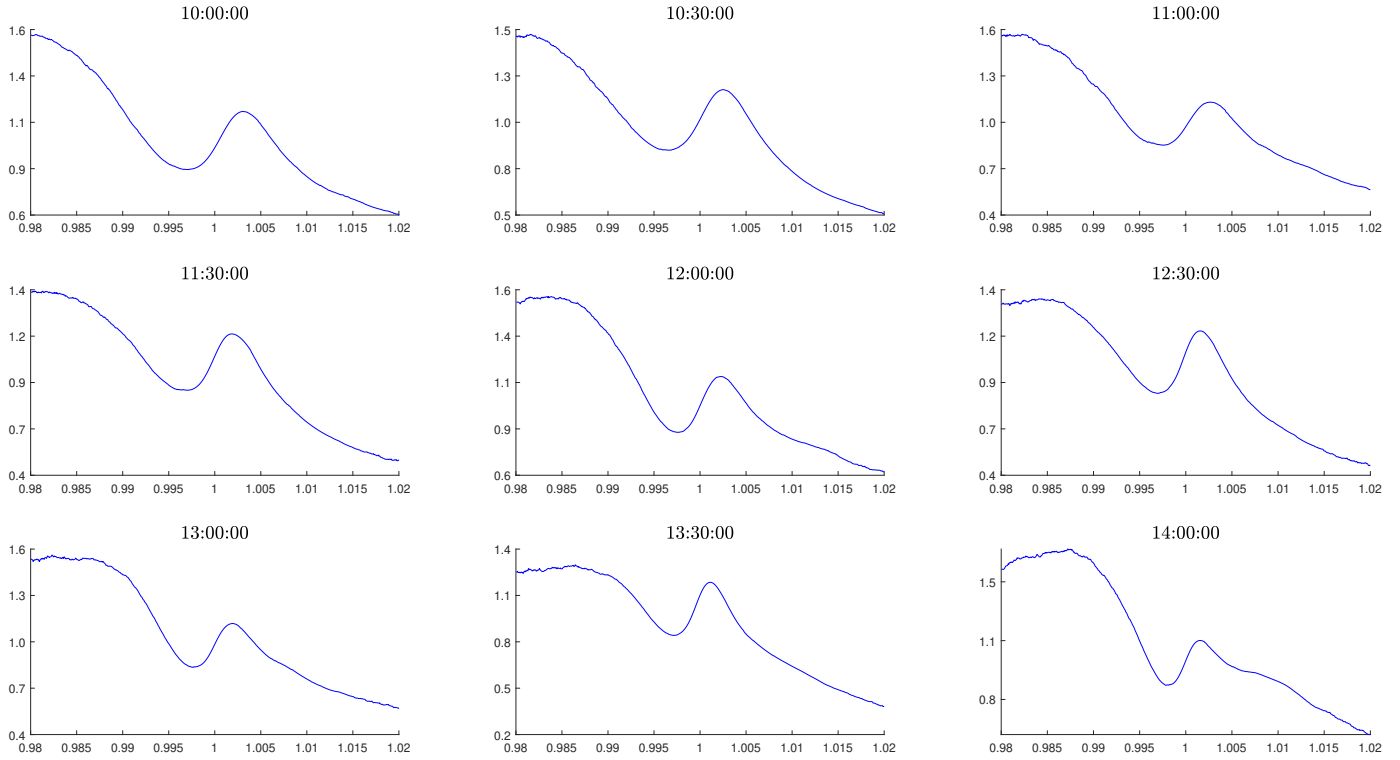
OA.3 Robustness - FOMC announcement days

Figure OA.10: Average returns of 0DTE option strategies - FOMC announcement days



Note: The figure plots, for different times of the day, the average returns together with 90% confidence bands for ATM delta-hedged calls, delta-neutral straddles and simple straddles, for FOMC announcement days. Confidence bands are based on 2,500 bootstrap replications. The total sample ranges from January 6 2012 to July 3 2023.

Figure OA.11: Average intra-day pricing kernels - removing FOMC days



Note: The figure plots, for different times of the day, the average of the pricing kernel, as a function of market returns, over time removing FOMC announcement days. The total sample ranges from January 6 2012 to July 3 2023.

Table OA.6: Intra-day variance risk premium - removing FOMC days

	10:00:00	10:30:00	11:00:00	11:30:00	12:00:00	12:30:00	13:00:00	13:30:00	14:00:00
Panel A: VRP									
Mean	2.920	3.588	2.419	3.642	2.478	4.462	3.366	6.699	4.764
St. Dev.	9.888	11.340	8.773	10.466	9.600	14.347	16.524	21.677	17.909
25th Percentile	0.816	1.043	0.722	1.051	0.742	1.252	0.932	1.906	1.391
75th Percentile	3.855	4.458	3.626	4.836	3.740	5.487	4.381	7.358	5.800
p -value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Panel B: VRP^+									
Mean	1.278	1.669	1.160	1.823	1.316	2.387	1.840	3.633	2.760
St. Dev.	4.598	5.282	4.348	5.055	4.821	6.908	8.081	10.177	8.586
25th Percentile	0.548	0.691	0.590	0.792	0.660	0.980	0.794	1.475	1.235
75th Percentile	1.841	2.239	1.901	2.528	2.129	3.103	2.582	4.304	3.598
p -value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Panel C: VRP^-									
Mean	1.642	1.919	1.259	1.819	1.148	2.075	1.520	3.066	2.004
St. Dev.	5.845	6.581	5.001	5.921	5.277	7.916	8.980	12.044	10.081
25th Percentile	0.278	0.303	0.146	0.258	0.084	0.261	0.091	0.390	0.186
75th Percentile	2.042	2.199	1.771	2.254	1.706	2.453	1.790	3.249	2.414
p -value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Note: The table reports, for each time of the day, summary statistics of the $VRP_{t,T}$, $VRP_{t,T}^+$ and $VRP_{t,T}^-$ over our sample removing FOMC days. The variance risk premium measures are annualized and expressed in percentage points. The p -values for the test with null hypothesis that the mean is smaller than or equal to zero, against the alternative that the mean is positive, are implemented using bootstrapped standard errors with 2,500 replications. The total sample ranges from January 6 2012 to July 3 2023.

Table OA.7: Predicting excess market returns with variance risk premium - removing FOMC days

	10:00:00				10:30:00				11:00:00			
<i>VRP</i>	-0.017				0.000				-0.032			
<i>t</i> -stat	-0.338				0.012				-1.216			
<i>VRP</i> ⁺		-0.037		-0.101		-0.023		-0.123		-0.054*		-0.110*
<i>t</i> -stat		-0.899		-1.550		-0.727		-1.626		-1.798		-1.783
<i>VRP</i> ⁻			0.001	0.080			0.019	0.120			-0.009	0.075
<i>t</i> -stat			0.009	0.833			0.451	1.296			-0.322	1.150
<i>R</i> ² (%)	0.050	0.239	0.000	0.650	0.000	0.111	0.077	1.072	0.240	0.681	0.020	1.240
	11:30:00				12:00:00				12:30:00			
<i>VRP</i>	-0.029				-0.048				-0.008			
<i>t</i> -stat	-1.265				-1.614				-0.171			
<i>VRP</i> ⁺		-0.055***		-0.151**		-0.069**		-0.135*		-0.034		-0.193**
<i>t</i> -stat		-2.865		-2.274		-2.435		-1.917		-0.797		-2.165
<i>VRP</i> ⁻			-0.005	0.118			-0.024	0.083			0.014	0.183*
<i>t</i> -stat			-0.160	1.514			-0.707	1.088			0.276	1.721
<i>R</i> ² (%)	0.238	0.821	0.008	2.101	0.728	1.500	0.186	2.307	0.022	0.368	0.067	2.972
	13:00:00				13:30:00				14:00:00			
<i>VRP</i>	-0.074**				-0.017				-0.026			
<i>t</i> -stat	-2.193				-0.289				-0.487			
<i>VRP</i> ⁺		-0.091***		-0.192***		-0.038		-0.214**		-0.040		-0.099**
<i>t</i> -stat		-3.331		-2.690		-0.822		-2.107		-0.892		-2.080
<i>VRP</i> ⁻			-0.053	0.115			0.002	0.195			-0.013	0.071
<i>t</i> -stat			-1.336	1.371			0.028	1.490			-0.217	0.922
<i>R</i> ² (%)	1.928	2.969	1.011	4.059	0.116	0.591	0.001	3.473	0.366	0.845	0.085	1.628

Note: The table reports, for each time of the day, the results from different predictive regressions over our sample removing FOMC announcement days using $VRP_{t,T}$, $VRP_{t,T}^+$ and $VRP_{t,T}^-$ to predict the excess market return from t to T . Regressors are standardized to have mean zero and unit variance. We compute the t -statistics using Newey-West robust standard errors with a lag length equal to 5. We denote with *, **, and *** significance at the 10%, 5% and 1% level, respectively. The total sample ranges from January 6 2012 to July 3 2023.

Table OA.8: Option price bounds for 0DTE options - removing FOMC days

		10:00:00	10:30:00	11:00:00	11:30:00	12:00:00	12:30:00	13:00:00	13:30:00	14:00:00
Panel A: Calls										
All Calls	In	38.023	41.388	41.589	40.965	36.410	35.439	39.100	45.717	36.444
	Upper	28.560	26.989	26.478	27.610	26.532	27.352	27.348	24.495	26.243
	Lower	33.417	31.623	31.934	31.425	37.058	37.209	33.552	29.789	37.313
OTM	In	61.559	65.315	66.994	62.884	59.204	55.871	59.784	70.268	57.131
	Upper	9.320	8.696	8.884	13.099	7.934	9.782	8.213	3.681	5.380
	Lower	29.122	25.989	24.121	24.017	32.862	34.347	32.002	26.051	37.489
ATM	In	7.664	8.413	8.571	9.850	7.102	7.836	8.974	9.485	5.458
	Upper	64.902	65.278	65.041	63.668	66.513	67.063	68.693	68.462	69.691
	Lower	27.434	26.308	26.388	26.482	26.385	25.101	22.333	22.054	24.851
ITM	In	18.592	23.932	22.436	26.085	18.997	21.077	25.768	30.272	23.385
	Upper	36.472	31.977	30.006	27.910	30.847	30.175	31.601	29.726	31.977
	Lower	44.936	44.091	47.558	46.005	50.156	48.748	42.631	40.002	44.638
Panel B: Puts										
All Calls	In	24.869	29.876	30.154	33.010	28.486	31.501	36.424	41.911	30.332
	Upper	35.881	34.301	34.378	34.110	36.371	34.564	32.354	29.967	33.680
	Lower	39.250	35.823	35.469	32.880	35.143	33.935	31.222	28.122	35.989
OTM	In	43.765	51.065	50.462	56.559	48.960	54.464	62.852	68.684	50.983
	Upper	5.500	5.188	4.828	4.574	6.888	6.033	5.036	3.750	6.580
	Lower	50.735	43.747	44.710	38.866	44.152	39.503	32.112	27.566	42.436
ATM	In	7.244	7.925	7.898	8.926	5.567	6.388	8.343	10.206	4.999
	Upper	69.484	70.356	72.025	71.893	77.022	75.782	72.343	70.384	75.178
	Lower	23.272	21.719	20.077	19.181	17.411	17.830	19.314	19.410	19.823
ITM	In	8.869	12.505	14.447	13.475	12.931	13.429	14.797	21.687	15.043
	Upper	58.466	54.584	54.074	54.021	54.240	51.001	48.084	43.999	48.602
	Lower	32.664	32.911	31.480	32.504	32.829	35.570	37.119	34.314	36.355

Note: The table reports, for each time of the day and for each class of options, the percentage of options over our sample, removing FOMC announcement days, for which prices fall within the SSD bounds (In), above the SSD upper bound (Upper) and below the SSD lower bound (Lower). The OTM put (ITM call), ATM and ITM put (OTM call) categories are defined as standardized log-moneyness below -1 , between -1 and 1 , and above 1 , respectively. The total sample ranges from January 6 2012 to July 3 2023.

Table OA.9: Sharpe ratios for SSD violation strategy - removing FOMC days

	10:00:00	10:30:00	11:00:00	11:30:00	12:00:00	12:30:00	13:00:00	13:30:00	14:00:00
Panel A: Before Transaction Costs									
Short ATM Call Delta-Hedge	0.003	0.003	0.003	0.013	0.005	0.031	0.043	0.035	0.027
SSD ATM Call Delta-Hedge	0.305	0.292	0.260	0.283	0.271	0.264	0.265	0.305	0.332
Panel B: After Transaction Costs									
Short ATM Call Delta-Hedge	-0.037	-0.026	-0.024	-0.017	-0.024	0.003	0.013	0.002	-0.014
SSD ATM Call Delta-Hedge	0.266	0.265	0.235	0.255	0.245	0.238	0.237	0.274	0.295

Note: The table reports, for each time of the day, before and after transaction costs, the Sharpe ratio associated with the SSD violation strategy based on the ATM call and the benchmark strategy writing the ATM call delta-hedged. To incorporate transaction costs, whenever we buy (sell) the option, we consider the ask (bid) price instead of the bid-ask midpoint, that is, we consider the worst case scenario for the strategy. The sample ranges from January 6 2012 to July 3 2023 excluding FOMC announcement days.

OA.4 Robustness - E-mini S&P 500 futures

Table OA.10: Intra-day variance risk premium computed with E-mini S&P 500 futures

	10:00:00	10:30:00	11:00:00	11:30:00	12:00:00	12:30:00	13:00:00	13:30:00	14:00:00
Panel A: VRP									
Mean	0.906	1.822	0.774	2.209	1.063	3.101	1.897	5.448	3.148
St. Dev.	9.969	10.638	9.880	10.293	10.901	13.230	15.581	21.187	22.700
25th Percentile	-0.559	-0.102	-0.416	0.084	-0.361	0.328	-0.174	0.828	0.004
75th Percentile	2.377	3.061	2.457	3.624	2.852	4.340	3.308	6.153	4.483
p -value	0.001	0.000	0.003	0.000	0.000	0.000	0.000	0.000	0.000
Panel B: VRP^+									
Mean	0.227	0.780	0.317	1.084	0.582	1.682	1.130	2.988	1.804
St. Dev.	5.703	5.396	5.349	5.445	5.924	6.863	7.797	10.273	10.888
25th Percentile	-0.218	0.127	-0.021	0.288	0.073	0.513	0.289	0.929	0.520
75th Percentile	1.224	1.634	1.351	2.018	1.675	2.596	2.103	3.823	2.926
p -value	0.133	0.000	0.027	0.000	0.000	0.000	0.000	0.000	0.000
Panel C: VRP^-									
Mean	0.679	1.042	0.457	1.125	0.481	1.419	0.767	2.459	1.344
St. Dev.	5.176	5.843	5.056	5.480	5.556	7.008	8.325	11.550	13.096
25th Percentile	-0.438	-0.232	-0.419	-0.215	-0.462	-0.221	-0.517	-0.145	-0.599
75th Percentile	1.238	1.474	1.082	1.652	1.203	1.948	1.383	2.656	1.743
p -value	0.000	0.000	0.001	0.000	0.001	0.000	0.000	0.000	0.000

Note: The table reports, for each time of the day, summary statistics of the $VRP_{t,T}$, $VRP_{t,T}^+$ and $VRP_{t,T}^-$ over our sample, where the realized variance measures are computed using E-mini S&P 500 futures returns, instead of returns on the S&P 500 index. The variance risk premium measures are annualized and expressed in percentage points. The p -values for the test with null hypothesis that the mean is smaller than or equal to zero, against the alternative that the mean is positive, are implemented using bootstrapped standard errors with 2,500 replications. The sample ranges from January 6 2012 to July 3 2023.

Table OA.11: Predicting excess market returns with variance risk premium based on E-mini S&P 500 futures

	10:00:00				10:30:00				11:00:00			
VRP	-0.036				0.024				-0.052*			
t -stat	-0.895				0.594				-1.697			
VRP^+		-0.068**		-0.134***		-0.003		-0.105*		-0.079***		-0.180***
t -stat		-2.387		-3.712		-0.084		-1.721		-2.624		-3.549
VRP^-			0.006	0.097*			0.046	0.129*			-0.019	0.126**
t -stat			0.142	1.697			1.112	1.880			-0.637	2.266
$R^2(\%)$	0.206	0.751	0.006	1.579	0.103	0.002	0.386	1.146	0.552	1.256	0.071	2.389
	11:30:00				12:00:00				12:30:00			
VRP	-0.050**				-0.052**				-0.053***			
t -stat	-2.288				-2.394				-2.764			
VRP^+		-0.082***		-0.183***		-0.077***		-0.171***		-0.076***		-0.170**
t -stat		-2.847		-3.022		-2.935		-2.660		-3.439		-2.572
VRP^-			-0.012	0.129**			-0.020	0.118*			-0.025	0.115
t -stat			-0.506	2.006			-0.912	1.721			-0.922	1.531
$R^2(\%)$	0.584	1.580	0.035	3.135	0.711	1.562	0.104	2.871	0.776	1.628	0.171	2.844
	13:00:00				13:30:00				14:00:00			
VRP	-0.077***				-0.045				-0.041*			
t -stat	-4.492				-1.585				-1.739			
VRP^+		-0.090***		-0.154**		-0.060**		-0.154***		-0.059***		-0.112***
t -stat		-4.148		-2.235		-2.307		-2.599		-2.631		-3.222
VRP^-			-0.060***	0.074			-0.030	0.106*			-0.022	0.066*
t -stat			-3.793	1.196			-0.884	1.688			-0.857	1.697
$R^2(\%)$	1.915	2.613	1.158	3.046	0.705	1.234	0.305	2.086	0.635	1.328	0.179	1.959

Note: The table reports, for each time of the day, the results from different predictive regressions over our sample using $VRP_{t,T}$, $VRP_{t,T}^+$ and $VRP_{t,T}^-$ to predict the excess market return from t to T , where the variance risk premium measures are computed using E-mini S&P 500 futures. Regressors are standardized to have mean zero and unit variance. We compute the t -statistics using Newey-West robust standard errors with a lag length equal to 5. We denote with *, **, and *** significance at the 10%, 5% and 1% level, respectively. The sample ranges from January 6 2012 to July 3 2023.

Table OA.12: Predicting E-mini S&P 500 futures returns with variance risk premium based on E-mini S&P 500 futures

	10:00:00				10:30:00				11:00:00			
VRP	-0.029				-0.034				-0.072**			
t -stat	-0.862				-1.150				-2.480			
VRP^+	-0.037		-0.049		-0.054**		-0.121*		-0.084***		-0.118*	
t -stat	-1.115		-0.698		-1.986		-1.782		-2.661		-1.782	
VRP^-	-0.015		0.018		-0.012		0.084		-0.052**		0.043	
t -stat	-0.360		0.221		-0.349		1.065		-2.018		0.734	
$R^2(\%)$	0.140	0.224	0.039	0.252	0.233	0.594	0.028	1.117	1.174	1.587	0.615	1.734
	11:30:00				12:00:00				12:30:00			
VRP	-0.070***				-0.049**				-0.064**			
t -stat	-3.212				-2.495				-2.485			
VRP^+	-0.099***		-0.183***		-0.078***		-0.187**		-0.091***		-0.198**	
t -stat	-3.528		-3.048		-3.101		-2.549		-2.650		-2.433	
VRP^-	-0.033		0.109*		-0.014		0.136*		-0.031		0.131*	
t -stat	-1.445		1.775		-0.583		1.658		-1.113		1.723	
$R^2(\%)$	1.257	2.519	0.280	3.750	0.718	1.777	0.058	3.714	1.242	2.518	0.302	4.234
	13:00:00				13:30:00				14:00:00			
VRP	-0.130***				-0.054*				-0.055**			
t -stat	-3.339				-1.760				-2.342			
VRP^+	-0.135***		-0.135**		-0.066**		-0.143**		-0.065***		-0.087**	
t -stat	-3.166		-2.483		-2.565		-2.239		-3.007		-2.184	
VRP^-	-0.117***		0.000		-0.040		0.087		-0.041		0.027	
t -stat	-3.208		0.003		-1.061		1.096		-1.533		0.534	
$R^2(\%)$	5.537	5.958	4.487	5.958	1.097	1.663	0.600	2.280	1.448	2.035	0.810	2.172

Note: The table reports, for each time of the day, the results from different predictive regressions over our sample using $VRP_{t,T}$, $VRP_{t,T}^+$ and $VRP_{t,T}^-$ to predict the E-mini S&P 500 futures return from t to T , where the variance risk premium measures are computed using E-mini S&P 500 futures. Regressors are standardized to have mean zero and unit variance. We compute the t -statistics using Newey-West robust standard errors with a lag length equal to 5. We denote with *, **, and *** significance at the 10%, 5% and 1% level, respectively. The sample ranges from January 6 2012 to July 3 2023.

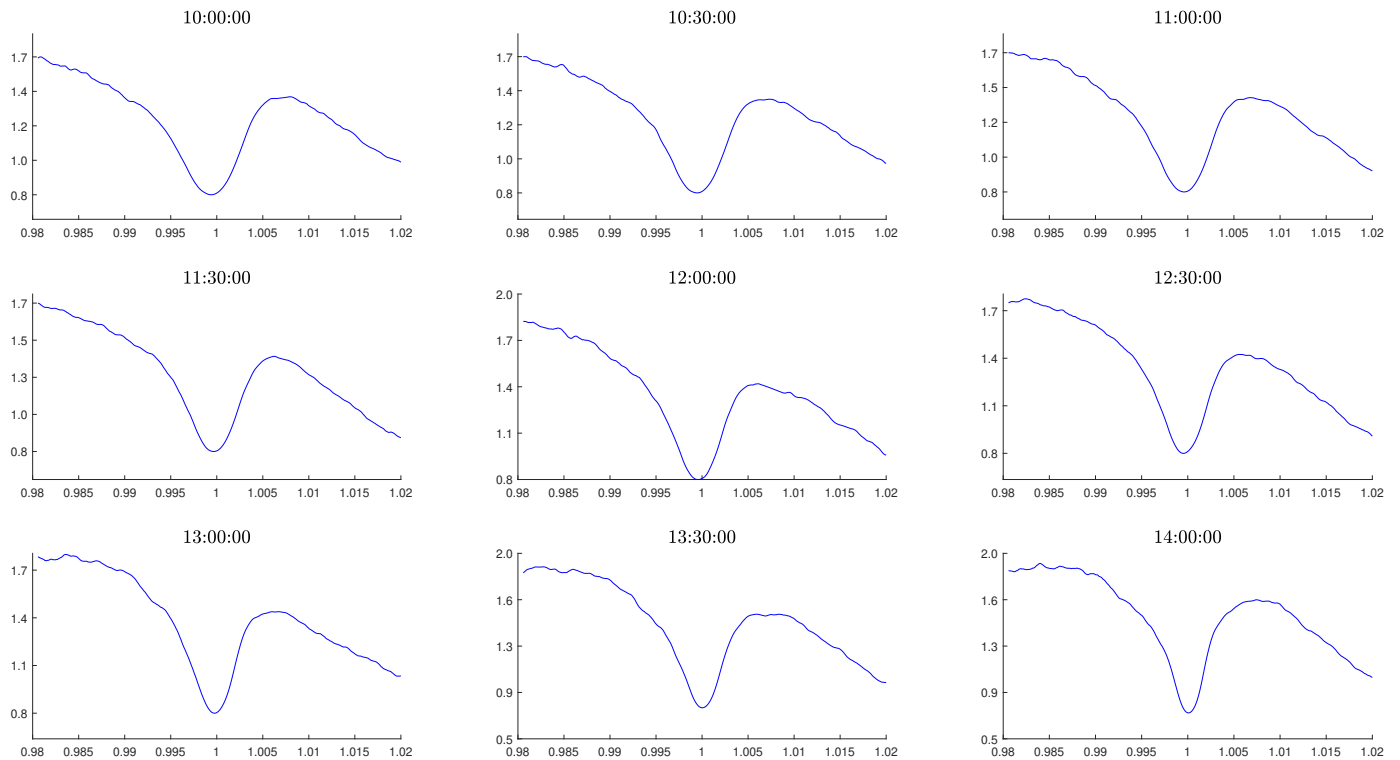
Table OA.13: Sharpe ratios for SSD violation strategy using E-mini S&P 500 futures to delta-hedge

	10:00:00	10:30:00	11:00:00	11:30:00	12:00:00	12:30:00	13:00:00	13:30:00	14:00:00
Panel A: Before Transaction Costs									
Short ATM Call Delta-Hedge	0.006	-0.002	0.007	0.020	0.037	0.043	0.030	0.015	-0.009
SSD ATM Call Delta-Hedge	0.081	0.066	0.076	0.068	0.090	0.105	0.090	0.056	0.038
Panel B: After Transaction Costs									
Short ATM Call Delta-Hedge	-0.035	-0.032	-0.020	-0.008	0.009	0.015	-0.001	-0.016	-0.047
SSD ATM Call Delta-Hedge	0.042	0.038	0.050	0.041	0.063	0.077	0.061	0.028	0.001

Note: The table reports, for each time of the day, before and after transaction costs, the Sharpe ratio associated with the SSD violation strategy based on the ATM call and the benchmark strategy writing the ATM call delta-hedged, when the delta-hedge is implemented using E-mini S&P 500 futures. To incorporate transaction costs, whenever we buy (sell) the option, we consider the ask (bid) price instead of the bid-ask midpoint, that is, we consider the worst case scenario for the strategy. The sample ranges from January 6 2012 to July 3 2023.

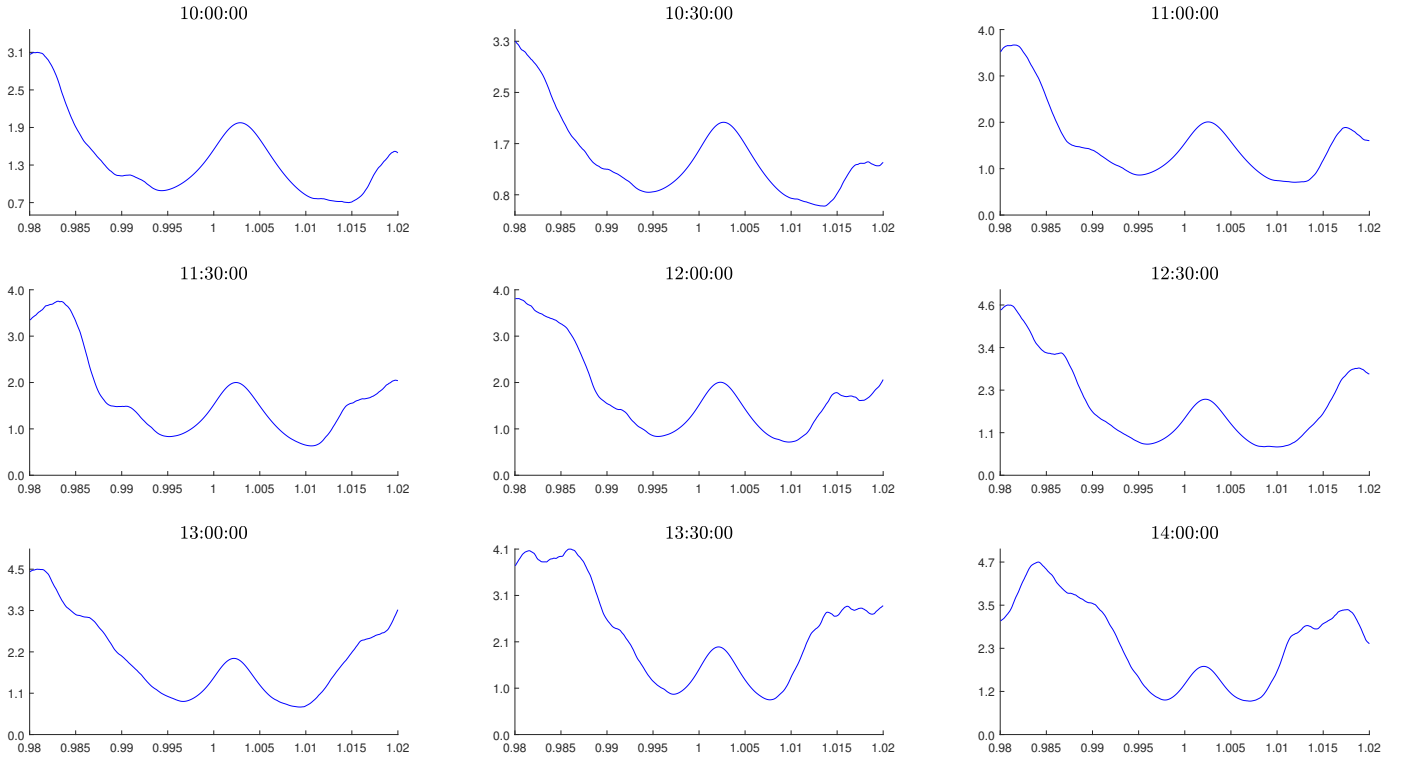
OA.5 Robustness - Alternative physical distributions

Figure OA.12: Average intra-day pricing kernels - conditional volatility with RV



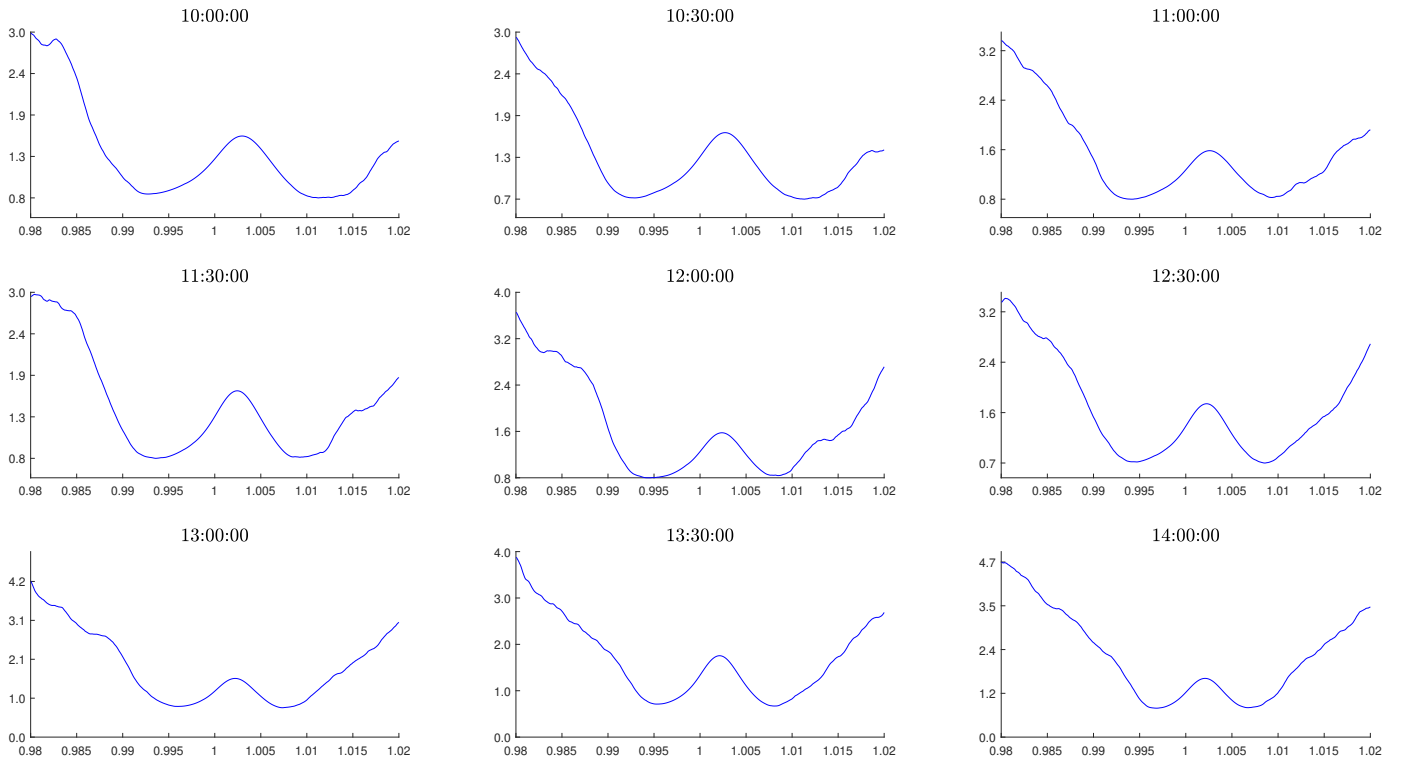
Note: The figure plots, for different times of the day, the average of the pricing kernel, as a function of market returns, over time. The conditional physical distribution is computed as detailed in Section 3.4 but with conditional volatility given by the realized volatility from the previous day for the horizon of the option, instead of the ATM implied volatility. The sample ranges from January 6 2012 to July 3 2023.

Figure OA.13: Average intra-day pricing kernels - Jackwerth's (2000) method



Note: The figure plots, for different times of the day, the average of the pricing kernel, as a function of market returns, over time. The conditional physical distribution is computed, similarly to [Jackwerth \(2000\)](#) but adapted to our horizons, as the histogram of past market returns over the horizon of the option over a 2-month window. For a given day, the obtained histogram is smoothed with a Gaussian kernel with bandwidth $\frac{x\sigma}{\sqrt{n}}$, where σ is the volatility of the return distribution, n is the number of observations, $x = 2.8$ and $m = 5$. As the pricing kernel is noisy in the tails, the plot is further smoothed by computing a moving average with 50 observations over the moneyness dimension. The sample ranges from January 6 2012 to July 3 2023.

Figure OA.14: Average intra-day pricing kernels - Jackwerth's (2000) method with volatility adjustment



Note: The figure plots, for different times of the day, the average of the pricing kernel, as a function of market returns, over time. The conditional physical distribution is computed, similarly to [Jackwerth \(2000\)](#) but adapted to our horizons, as the histogram of past market returns over the horizon of the option over a 2-month window, further adjusted to have volatility equal to the ATM implied volatility on day t . For a given day, the obtained histogram is smoothed with a Gaussian kernel with bandwidth $\frac{x\sigma}{\sqrt[n]{n}}$, where σ is the volatility of the return distribution, n is the number of observations, $x = 1.8$ and $m = 4$. As the pricing kernel is noisy in the tails, the plot is further smoothed by computing a moving average with 50 observations over the moneyness dimension. The sample ranges from January 6 2012 to July 3 2023.