0DTE Asset Pricing

Caio Almeida[*a], **Gustavo Freire**[†b], and **Rodrigo Hizmeri**[‡c]

[a]Princeton University
[b]Erasmus University Rotterdam, Tinbergen Institute & ERIM
[c]University of Liverpool

This draft: January 26, 2024
First draft: January 20, 2024

Abstract

We explore the asset pricing implications of zero days-to-expiration (0DTE) options. This new market has experienced remarkable growth in recent years and today accounts for 45% of total S&P 500 option volume. Extracting information about intra-day risk premia from 0DTEs, we document 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, consistent with extremely negative average returns of calls, puts and straddles; (iii) the intra-day pricing kernel displays pronounced nonmonotonicities, especially around the at-the-money region; (iv) 0DTE options are severely mispriced, in the sense that they do not reflect risks implied by the time series of intra-day underlying returns; and (v) the variance risk premium negatively predicts intra-day market returns, where this negative relation is driven by the premium to bear positive return variation risk.

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

JEL Classification: G11, G12, G13.

[*]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.

Electronic copy available at: https://ssrn.com/abstract=4701401

1 Introduction

Due to the contingent nature of options, their prices across strikes allow to recover investors' expectations (incorporating compensation for risk) of the probability distribution of the underlying return over the options maturity (Breeden and Litzenberger, 1978; Ross, 1976). Building on that, Jackwerth and Rubinstein (1996) and Aït-Sahalia and Lo (1998) show how the risk-neutral distribution implied by S&P 500 options became more left-skewed and leptokurtic after the 1987 crash, reflecting aversion to downside risk. Relatedly, Bakshi, Kapadia, and Madan (2003) provide a comprehensive analysis of the determinants of risk-neutral skewness based on similar spanning properties of options.

When confronted with the physical probabilities implied by the time series of market returns, the risk-neutral distribution is informative about investors' risk preferences. Jackwerth (2000), Aït-Sahalia and Lo (2000) and Rosenberg and Engle (2002) exploit this insight to estimate the projection of the economy-wide pricing kernel onto market return states. The shape of this projection is often nonmonotonic, which is puzzling under a representative investor framework.[1] From a different perspective, Almeida and Freire (2022) show that S&P 500 options can be reconciled with the time series of market returns in the sense that their prices satisfy bounds consistent with risk-averse investors, where the preferences of the marginal agent vary across the options.

The relation between expectations under risk-neutral and physical probabilities also sheds light on compensation for specific types of risk in the dynamics of the underlying asset returns. Beason and Schreindorfer (2022) decompose the one-month equity premium into different parts of the return state space, revealing it mainly arises as compensation for returns between -30% and -10%. Also leveraging the information content of options relative to realized returns, Bollerslev, Tauchen, and Zhou (2009) show that the premium for bearing variance risk is generally positive and helps predict future market returns. Bollerslev, Todorov, and Xu (2015), Andersen, Fusari, and Todorov (2015) and Andersen, Fusari, and Todorov (2017) document, instead, the special role of compensation for jump risk in determining the equity and variance risk premia.

[1] This is known as the pricing kernel puzzle. See Cuesdeanu and Jackwerth (2018) for a recent survey.

Electronic copy available at: https://ssrn.com/abstract=4701401

In this paper, we investigate the asset pricing implications of a new, 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 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.[2] 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 45% of total S&P 500 option volume, being thus by far the most traded maturity.[3] 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.[4] To extract this information, we explore how 0DTE option prices and implied risk-neutral probabilities relate to the time series of high-frequency market returns, leveraging different methodologies available from the literature. 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 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

[2]The first S&P 500 weekly options were introduced by CBOE on October 28 2005 with Friday expirations. Wednesday, Monday, Tuesday and Thursday expirations followed with introduction dates of February 23 2016, August 15 2016, April 18 2022 and May 11 2022, respectively.

[3]This increase in the relative importance of short-dated options does not mean that the traded volume in longer-dated options has decreased. On the contrary, their volume also increased, meaning that the size of the option market has been increasing over time, and relatively more for short-dated options.

[4]While 0DTE options are extremely popular among retail investors, accounting for more than 75% of their trading in S&P 500 options (Beckmeyer, Branger, and Gayda, 2023), the vast majority of 0DTE S&P 500 trading (around 94%) is still attributable to institutional investors.

Electronic copy available at: https://ssrn.com/abstract=4701401

(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.[5] 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, 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 further disentangle the compensation demanded by investors to bear variation risk in positive and negative market returns, we respectively compute the "good" and "bad" components of the total variance risk premium, in the spirit of 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 premium for positive return variation risk is positive and often larger than the premium for exposure to negative return variation. This, again, would be consistent with

[5] As Bollerslev et al. (2009), we define the variance risk premium as the difference between the risk-neutral and the physical expected variance of the market return.

Electronic copy available at: https://ssrn.com/abstract=4701401

a pricing kernel that is exceptionally high for positive market returns, such that investors demand compensation for these states of high marginal utility.

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. We also estimate directly 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 one. 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 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 option prices satisfy these bounds, which is again in stark contrast with evidence from longer horizons (Almeida and Freire, 2022). Most of the violations are concentrated in ATM options, where their prices are often too high. This is consistent with the pronounced hump observed in the pricing kernel projection. 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 reasonable risk preferences.

Finally, we investigate whether 0DTE options contain predictive information for excess market returns over the day. We consider predictive regressions using different regressors, such as the variance risk premium, risk-neutral moments (Bakshi et al., 2003) and the equity premium lower bound of Martin (2017). From these variables, only the variance

Electronic copy available at: https://ssrn.com/abstract=4701401

risk premium significantly predicts returns, but with a negative coefficient, which is at odds with the existing evidence for longer horizons. 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 positive return variation 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 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, while Section 4 contains our empirical analysis. Section 5 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 a few papers (so far) studying the new 0DTE option market from different lenses. Brogaard, Han, and Won (2023) show that more 0DTE option trading increases the volatility of the underlying asset, while Dim, Eraker, and Vilkov (2024) provide contrary evidence based on 0DTE aggregate gamma. Beckmeyer et al. (2023) document that 0DTE options are popular among retail traders, even though these investors mainly experience losses in this market. Bandi, Fusari, and Renò (2023) present a novel option pricing formula designed for 0DTEs and analyze the role of leverage and volatility-of-volatility in determining instantaneous market risk premia. Vilkov (2023) investigates the performance of different 0DTE option trading strategies. We contribute by exploring a broad set of asset pricing implications of 0DTE options, including investors' risk preferences implied by their prices and their implications for intra-day risk premia.

The second strand, already discussed in the introduction, recovers information about

Electronic copy available at: https://ssrn.com/abstract=4701401

investors' expectations and risk preferences from options. We provide novel asset pricing stylized facts for intra-daily horizons using information from the new 0DTE market, which is the most relevant option market today. In particular, we analyze the intra-day pricing kernel implied by 0DTEs and show how it is a nonmonotonic function of market returns. Relatedly, Aleti and Bollerslev (2023) 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 forward-looking and completely conditional on the time of the day that the options are observed.

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

Electronic copy available at: https://ssrn.com/abstract=4701401

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 $m_{t,T}$:

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

where s represents the state of the economy, $\mathcal{F}_t$ is the information available to investors at time t and $\pi^{\mathbb{P}}_{t,T}(s)$ 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 $\mathbb{E}_t[m_{t,T}]$:

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

where $\pi^{\mathbb{Q}}_{t,T}(s)$ 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, $\pi^{\mathbb{P}}_{t,T}(s)$ 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:

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

The economy-wide pricing kernel above depends on the realization of the state s of the

Electronic copy available at: https://ssrn.com/abstract=4701401

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), Aït-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:

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

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, $\pi^{\mathbb{Q}}_{t,T}(R_{t,T})$. On the other hand, historical market returns are informative about the physical distribution $\pi^{\mathbb{P}}_{t,T}(R_{t,T})$.[6] 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 a nonmonotonic (generally U-shaped) function of returns instead, characterizing the pricing kernel puzzle (Cuesdeanu and Jackwerth, 2018).[7] In the next subsections, we discuss how various objects of interest in our analysis relate to the pricing kernel and its shape.

[6] There is a potential mismatch of conditioning information sets when estimating $\pi^{\mathbb{Q}}_{t,T}(R_{t,T})$ with option prices, that are forward-looking, and $\pi^{\mathbb{P}}_{t,T}(R_{t,T})$ 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.

[7] More 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.

Electronic copy available at: https://ssrn.com/abstract=4701401

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:[8]

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

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.[9]

2.3 Market risk premia

The shape of $m_{t,T}(R_{t,T})$, or, equivalently, how $\pi^{\mathbb{Q}}_{t,T}(R_{t,T})$ relates to $\pi^{\mathbb{P}}_{t,T}(R_{t,T})$, 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:

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

[9] Since 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.

Electronic copy available at: https://ssrn.com/abstract=4701401

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

To decompose the unconditional equity premium, Beason and Schreindorfer (2022) take the unconditional expectation of (6) and consider net market returns to define:

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

where $\pi^{\mathbb{P}}(R) = \mathbb{E}[\pi^{\mathbb{P}}_{t,T}(R_{t,T})]$ and $\pi^{\mathbb{Q}}(R) = \mathbb{E}[\pi^{\mathbb{Q}}_{t,T}(R_{t,T})]$. $EP(x)$ 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 $R \approx 0$), 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 R contributes positively to the equity premium, i.e., $EP(x)$ is always increasing. To see that, note that in the negative return region, $R < 0$ and the pricing kernel is above 1, which means that $\pi^{\mathbb{P}}(R) - \pi^{\mathbb{Q}}(R) < 0$, such that $R\,[\pi^{\mathbb{P}}(R) - \pi^{\mathbb{Q}}(R)] > 0$ and $EP(x)$ is increasing. Analogously, in the positive return region, $R > 0$ and the pricing kernel is below 1, which means that $\pi^{\mathbb{P}}(R) - \pi^{\mathbb{Q}}(R) > 0$, such that $R\,[\pi^{\mathbb{P}}(R) - \pi^{\mathbb{Q}}(R)] > 0$ and $EP(x)$ is again increasing. Now, if we consider instead a U-shaped pricing kernel where $\pi^{\mathbb{Q}}(R)$ is above $\pi^{\mathbb{P}}(R)$ (i.e., pricing kernel is above 1) for some positive return $R > 0$, then $R\,[\pi^{\mathbb{P}}(R) - \pi^{\mathbb{Q}}(R)] < 0$ and $EP(x)$ is decreasing, i.e., such positive returns contribute negatively to the equity premium.

Our economic interpretation for these relations is as follows. For each state R, consider the asset that pays R in this state and zero otherwise, i.e., the asset buying R units of the Arrow-Debreu security of state R. Then, $R\,[\pi^{\mathbb{P}}(R) - \pi^{\mathbb{Q}}(R)]$ 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., $R\,[\pi^{\mathbb{P}}(R) - \pi^{\mathbb{Q}}(R)] > 0$. 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

Electronic copy available at: https://ssrn.com/abstract=4701401

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., $R\left[\pi^{\mathbb{P}}(R) - \pi^{\mathbb{Q}}(R)\right] < 0$. That is, investors would be willing to give up compensation to hedge against these states, such that they contribute negatively to the equity premium.

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 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 (negative) and the positive return variation premium is negative (positive). 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 re-

Electronic copy available at: https://ssrn.com/abstract=4701401

turns 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.

2.4 Option price bounds

The shape of $m_{t,T}(R_{t,T})$ says something about how the risk-neutral distribution implied by option prices compares to the physical distribution implied by market returns. A related, but alternative way of investigating this relation is by comparing observed option prices with option price bounds consistent with the physical distribution $\pi^{\mathbb{P}}_{t,T}(R_{t,T})$ 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 underlying asset and the risk-free rate. In other words, these bounds give 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 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

Electronic copy available at: https://ssrn.com/abstract=4701401

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.[10]

2.5 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.

[10]This 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.

Electronic copy available at: https://ssrn.com/abstract=4701401

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

Electronic copy available at: https://ssrn.com/abstract=4701401

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.

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.[11] 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. 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.

[11] IVs 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.

Electronic copy available at: https://ssrn.com/abstract=4701401

Finally, our high-frequency data for 1-minute market returns comes from Refinitiv Tick History, from January 1996 to July 2023. The risk-free rate, which we obtain for the same sample, is the daily 1-month Treasury bill from the FRED (Federal Reserve Bank of St. Louis) website. We assume that the risk-free rate is constant over 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 $O^c_{t,T}$ ($O^p_{t,T}$) and strike K as:

$$R_c = \frac{\max(S_T - K, 0)}{O^c_{t,T}} - 1, \quad R_p = \frac{\max(K - S_T, 0)}{O^p_{t,T}} - 1, \tag{8}$$

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:

$$\text{Simple-Straddle} = \frac{\max(S_T - K, 0) + \max(K - S_T, 0)}{O^c_{t,T} + O^p_{t,T}} - 1. \tag{9}$$

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:

$$\text{Straddle} = wR_c + (1 - w)R_p, \quad w = -\frac{\Delta_p / O^p_{t,T}}{\Delta_c / O^c_{t,T} - \Delta_p / O^p_{t,T}}, \tag{10}$$

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

Electronic copy available at: https://ssrn.com/abstract=4701401

returns as:

$$\Delta\text{-Hedged} = \frac{\max(S_T - K, 0) - O^c_{t,T} - \Delta_c(S_T - S_t) - r^f_t(O^c_{t,T} - \Delta_c S_t) \times \frac{\tau}{(24\times 365)}}{S_t}, \quad (11)$$

where r^f_t 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 0DTE option prices. Breeden and Litzenberger (1978) show that, under no-arbitrage and in the presence of a continuum of options 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:

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

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

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 $\pi^{\mathbb{Q}}_{t,T}(R_{t,T})$. 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.[12] Only OTM

[12]It 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

Electronic copy available at: https://ssrn.com/abstract=4701401

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:

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

where $k = \log(K/S_t)$ is the log-moneyness and a, b, ρ, m and σ are parameters.[13] 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.[14]

Figure 4 plots, for a representative date of our sample, 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 OLS R^2's higher than 95%, while filtering out noise in the observed IV curve.[15] This, in turn, results in well-behaved risk-neutral distributions, as can be seen in Figure 5. 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.

option prices as IVs are comparable across strikes.

[13] More 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.

[14] The 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.

[15] The average R^2 fit of the SVI over our sample is 95.13%, 95.25%, 95.11%, 95.23%, 95.39%, 95.13%, 94.07%, 94.23%, and 93.94%, at 10:00, 10:30, 11:00, 11:30, 12:00, 12:30, 13:00, 13:30, and 14:00 hours, respectively.

Electronic copy available at: https://ssrn.com/abstract=4701401

3.4 Physical distribution

To calculate the pricing kernel projection and the SSD option price bounds, we need to estimate the conditional physical distribution $\pi^{\mathbb{P}}_{t,T}(R_{t,T})$. 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 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.[16] 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 $\pi^{\mathbb{P}}_{t,T}(R_{t,T})$ and $\pi^{\mathbb{Q}}_{t,T}(R_{t,T})$ comparable. While other papers also use the ATM IV as a proxy for expected physical volatility (e.g., Dew-Becker, Giglio, and Kelly, 2021), it contains a risk premium and is, therefore, a biased proxy. We do not view this as an issue for our analysis, however, as we are being conservative in giving $\pi^{\mathbb{P}}_{t,T}(R_{t,T})$ the "best shot" at being close to $\pi^{\mathbb{Q}}_{t,T}(R_{t,T})$. In this sense, if we find nonmonotonicities in the pricing kernel and violations of the SSD bounds, we know that these patterns would be even more extreme under the return distribution with the "true" conditional expectation of physical variance.

[16] Following 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.

Electronic copy available at: https://ssrn.com/abstract=4701401

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 $\pi^{\mathbb{P}}_{t,T}(R_{t,T})$. Figure 5 displays, for a 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 Pricing kernel, risk premia and SSD bounds

Given the estimated risk-neutral distribution $\pi^{\mathbb{Q}}_{t,T}(R_{t,T})$ and physical distribution $\pi^{\mathbb{P}}_{t,T}(R_{t,T})$, we compute the pricing kernel on a given day as the ratio of the two probability distributions multiplied by the inverse of the appropriate risk-free rate, as in (4). We will investigate the shape of the pricing kernel over our sample.

Using the same estimated physical distribution, we compute, for each call option with strike K, the following SSD upper and lower price bounds as in Ritchken (1985):[17]

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

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

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 $\pi^{\mathbb{P}}_{t,T}(R_{t,T})$ 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

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

Electronic copy available at: https://ssrn.com/abstract=4701401

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 $\pi^{\mathbb{Q}}(R)$ over our sample as the average of the conditional risk-neutral distributions, i.e., the average of $\pi^{\mathbb{Q}}_{t,T}(R_{t,T})$ state by state. With that, we can evaluate $\int_{-1}^{\infty} R\pi^{\mathbb{Q}}(R)dR$ numerically over the grid of R (which is equal to the grid of $R_{t,T}-1$). Then, we compute $\int_{-1}^{\infty} R\pi^{\mathbb{P}}(R)dR$ from the unconditional empirical distribution of market returns (of the horizon of the option) over our sample as $(1/\mathcal{T})\sum_{i=1}^{\mathcal{T}} R_{i,t,T}\mathbf{1}\{R_{i,t,T}\leq x\}$, where $\mathcal{T}$ denotes the total number of days, $R_{i,t,T}$ 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 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 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):

$$V^{\mathbb{Q}}_{t,T} = \int_{S_t}^{\infty} \frac{2[1-\log(K/S_t)]}{K^2} O^c_{t,T}(K)dK + \int_0^{S_t} \frac{2[1+\log(S_t/K)]}{K^2} O^p_{t,T}(K)dK, \tag{16}$$

where we compute the integrals using the interpolated and extrapolated option prices from the SVI method. The realized variance $RV_{t,T}$ 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, $VRP_{t,T} = VRP^+_{t,T} + VRP^-_{t,T}$.

Electronic copy available at: https://ssrn.com/abstract=4701401

4 Empirical results

4.1 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 6 and 7 plot the median call and put returns over our sample, respectively, together with the corresponding 25th and 75th percentiles. We report percentiles of the returns instead of averages as these measures are more robust to outliers.[18] Observed patterns are very similar across different times of the day. Call options experience low returns overall, which are decreasing with the strike and eventually negative. Following the rationale of Bakshi et al. (2010), this provides evidence that the intra-day pricing kernel is a U-shaped function of market returns, i.e., that it is increasing in a region of positive return states. The percentiles reveal that there is some variation in call returns over time. The 25th percentile is consistent with a more pronounced U-shape. For the 75th percentile, returns are increasing up to the ATM region and then decrease, suggesting a much weaker, albeit still present, pricing kernel nonmonotonicity. 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.

Figure 7 reports the results for puts. Median returns are always negative, and in particular are equal to -100% 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. However, there is some variation in returns over time. Put returns can be even more negative, as seen with the 25th percentile, while in some occasions they can be positive and decreasing with the strike, as seen with the 75th percentile. The latter case suggests there may also be nonmonotonicities in negative market return states over part of the sample. Again, writing naked OTM puts is usually a profitable strategy, compatible with compensation for negative return variation risk.

[18] Due to the small prices of 0DTE options, for some days returns can be quite extreme.

Electronic copy available at: https://ssrn.com/abstract=4701401

We further consider the returns of ATM delta-hedged call options, straddles and delta-neutral straddles. As the return distribution of delta-neutral strategies is better behaved, we analyze average returns over our sample. Figure 8 displays the average returns together with 95% 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 that investors are willing to pay a premium to be protected against variance risk over the day. On the other hand, an investor disposed to be exposed to this risk would profit 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 9 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 equity premium from 10:00 (14:00) to close, which would amount to an annualized equity premium of 40% (150%), as can be seen in the right axis of the plot. These extreme values are due to the fact that positive market returns contribute negatively to the equity premium. This is consistent with a strongly 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,

Electronic copy available at: https://ssrn.com/abstract=4701401

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 negative on average. Even more striking is that this compensation is of similar magnitude, and in fact larger from 11:30 onwards, than the compensation for bearing negative return variation risk as captured by $VRP^{-}_{t,T}$.

Figure 10 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 that stands out in the plot 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 positive return variation risk premium ($VRP^{+}_{t,T}$) and the negative return variation risk premium ($VRP^{-}_{t,T}$).

4.3 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

Electronic copy available at: https://ssrn.com/abstract=4701401

its average shape over our sample. Figure 11 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 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 12 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 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, ..., up to 6. Table 2 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.

Electronic copy available at: https://ssrn.com/abstract=4701401

4.4 0DTE (mis)pricing

So far, we have shown that 0DTE 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 2.4, a nonmonotonic pricing kernel projection does not necessarily imply that option prices violate SSD bounds.

Table 3 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 reasonable risk preferences.

Figures 13 and 14 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 very close

Electronic copy available at: https://ssrn.com/abstract=4701401

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. The main takeaway from the figures is that 0DTEs do not violate SSD bounds due to the smile of the IV curve, as the monotonically decreasing pricing kernels also generate a smile. The mispricing comes from specific categories of options being either too expensive or cheap.

4.5 Return predictability

Finally, we investigate whether ex-ante measures based on 0DTE options are able to predict realized excess market returns. We first focus on the variance risk premium and its components, given that these measures 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. We consider predictive regressions over our sample for different times of the day, where all variables are available in real time at t and are used to predict the market return from t to T.

Table 4 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

Electronic copy available at: https://ssrn.com/abstract=4701401

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 total 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, these assets are hedges, 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 are 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 5 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 6 further considers predictive regressions for the variance risk premium mea-

Electronic copy available at: https://ssrn.com/abstract=4701401

sures including the risk-neutral moments as controls.[19] We consider times of the day 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.

5 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 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 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, direct estimates of the pricing kernel projection onto market returns reveal pronounced nonmonotonicities, especially around the at-the-money region.

Fourth, 0DTE options present severe violations of stochastic dominance bounds, and thus can 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. Fifth, 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 variation risk in

[19] Results including RV instead of $SVIX$ are very similar.

Electronic copy available at: https://ssrn.com/abstract=4701401

positive market returns. Our empirical results are all consistent with a strong U-shape in the pricing kernel as a function of market returns. While there is a large literature documenting the pricing kernel puzzle and attempting to explain nonmonotonic patterns at the one-month horizon, we find that the shape of the pricing kernel is even more puzzling at the intra-daily horizons associated with 0DTE options.

Electronic copy available at: https://ssrn.com/abstract=4701401

References

Aït-Sahalia, Y., and Lo, A. W. (1998). Nonparametric estimation of state-price densities implicit in financial asset prices. *The Journal of Finance*, *53*(2), 499–547.

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. (2023). News and asset pricing: A high-frequency anatomy of the sdf. *SSRN Working Paper*.

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., 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.

Aït-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.

Electronic copy available at: https://ssrn.com/abstract=4701401

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). 0dte 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 0dte 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*.

Constantinides, G. M., Jackwerth, J. C., and Perrakis, S. (2009). Mispricing of S&P 500

Electronic copy available at: https://ssrn.com/abstract=4701401

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). 0dtes: 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*.

Jackwerth, J. C. (2000). Recovering risk aversion from option prices and realized returns. *Review of Financial Studies*, *13*(2), 433-451.

Jackwerth, J. C., and Rubinstein, M. (1996). Recovering probability distributions from option prices. *The Journal of Finance*, *51*(5), 1611-1632.

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),

Electronic copy available at: https://ssrn.com/abstract=4701401

1219-1233.

Rosenberg, J. V., and Engle, R. F. (2002). Empirical pricing kernels. *Journal of Financial Economics*, *64*(3), 341-372.

Ross, S. A. (1976). Options and Efficiency. *The Quarterly Journal of Economics*, *90*(1), 75-89.

Vilkov, G. (2023). 0dte trading rules. *SSRN Working Paper*.

Electronic copy available at: https://ssrn.com/abstract=4701401

A Figures and tables

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

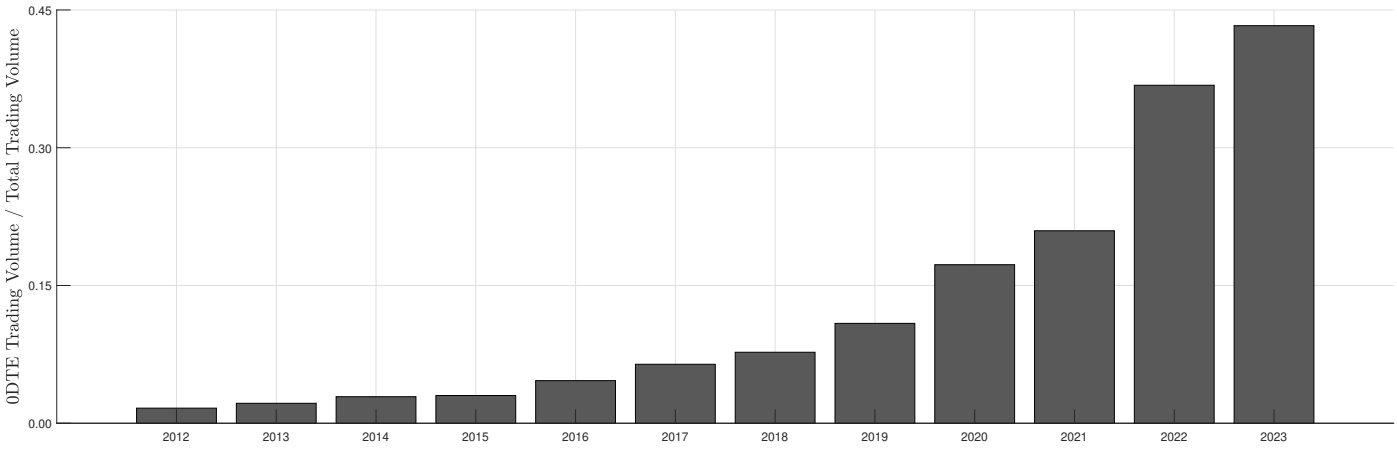

Note: The figure depicts the yearly fraction of trading volume in 0DTE S&P 500 options relative to the entire S&P 500 option market. The sample ranges from January 6 2012 to July 3 2023.

Electronic copy available at: https://ssrn.com/abstract=4701401

Figure 2: Descriptive statistics of raw 0DTE option data

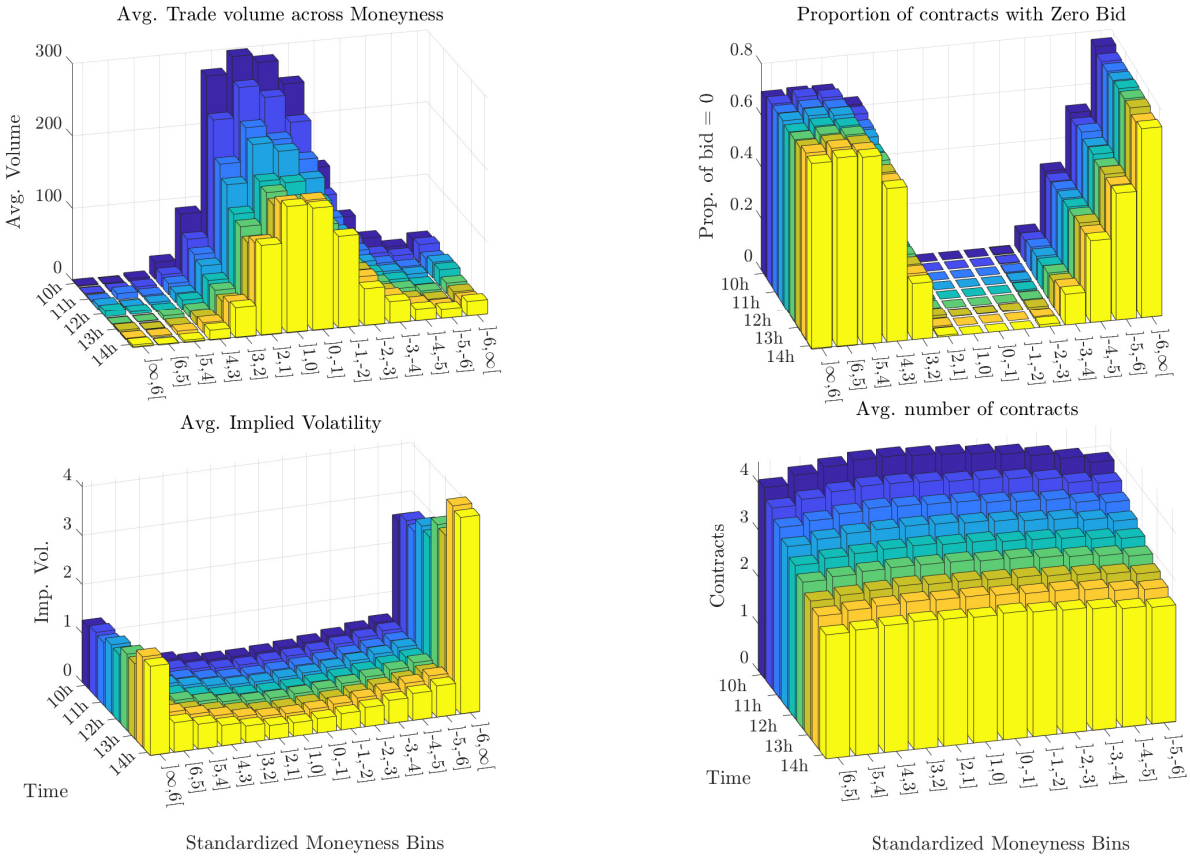

Note: The figure plots, in four subplots, for different times of the day and different standardized log-moneyness bins, the average trading volume, proportion of contracts with zero bid, average implied volatility and average number of strikes. The sample ranges from January 6 2012 to July 3 2023.

Electronic copy available at: https://ssrn.com/abstract=4701401

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

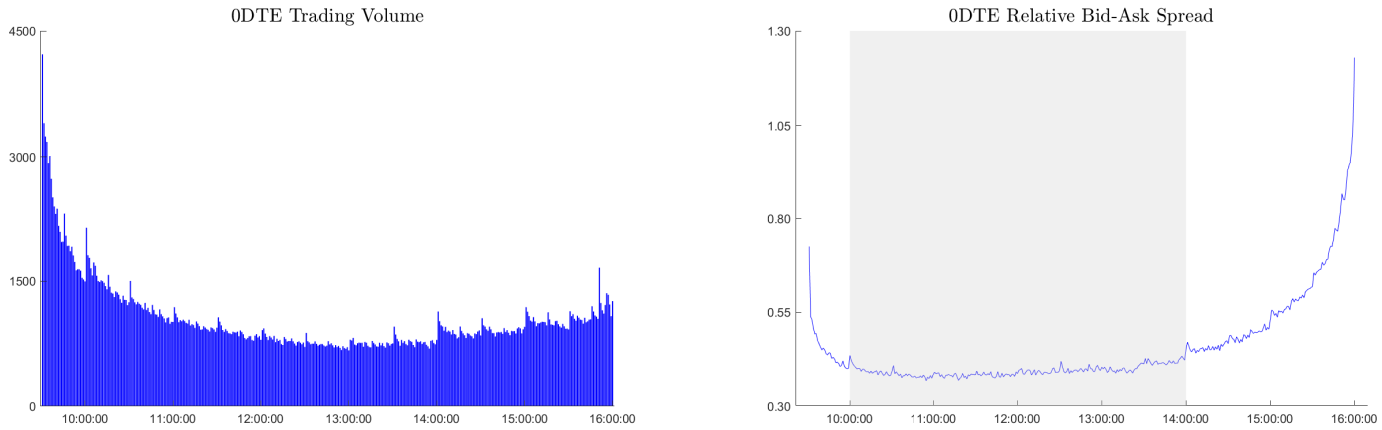

Note: The left plot of this figure depicts the time series average of the trading volume in terms of number of contracts of 0DTE S&P 500 options over the day. The right plot depicts the time series average of the average relative bid-ask spread across all 0DTE options over the day, where the relative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

Electronic copy available at: https://ssrn.com/abstract=4701401

Figure 4: SVI fit for 0DTE options

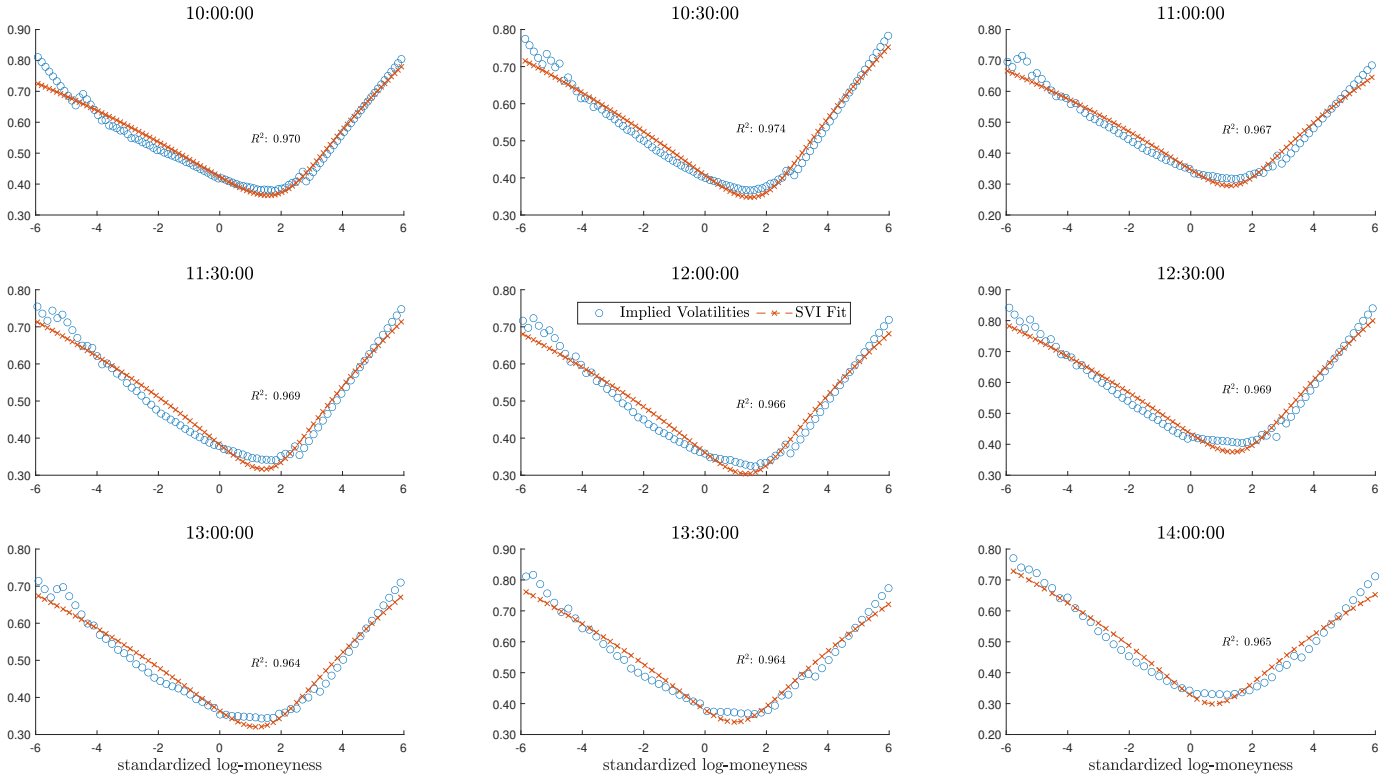

Note: The figure plots, for a representative date of our sample, the observed 0DTE IVs and the fitted IVs using the SVI method for different times of the day. The OLS R^2 fit is also reported. Standardized log-moneyness is defined as $\frac{\log(K/S_t)}{\sigma_{BS}(0)\sqrt{\tau}}$.

Electronic copy available at: https://ssrn.com/abstract=4701401

Figure 5: Market return risk-neutral and physical distributions over intra-daily horizons

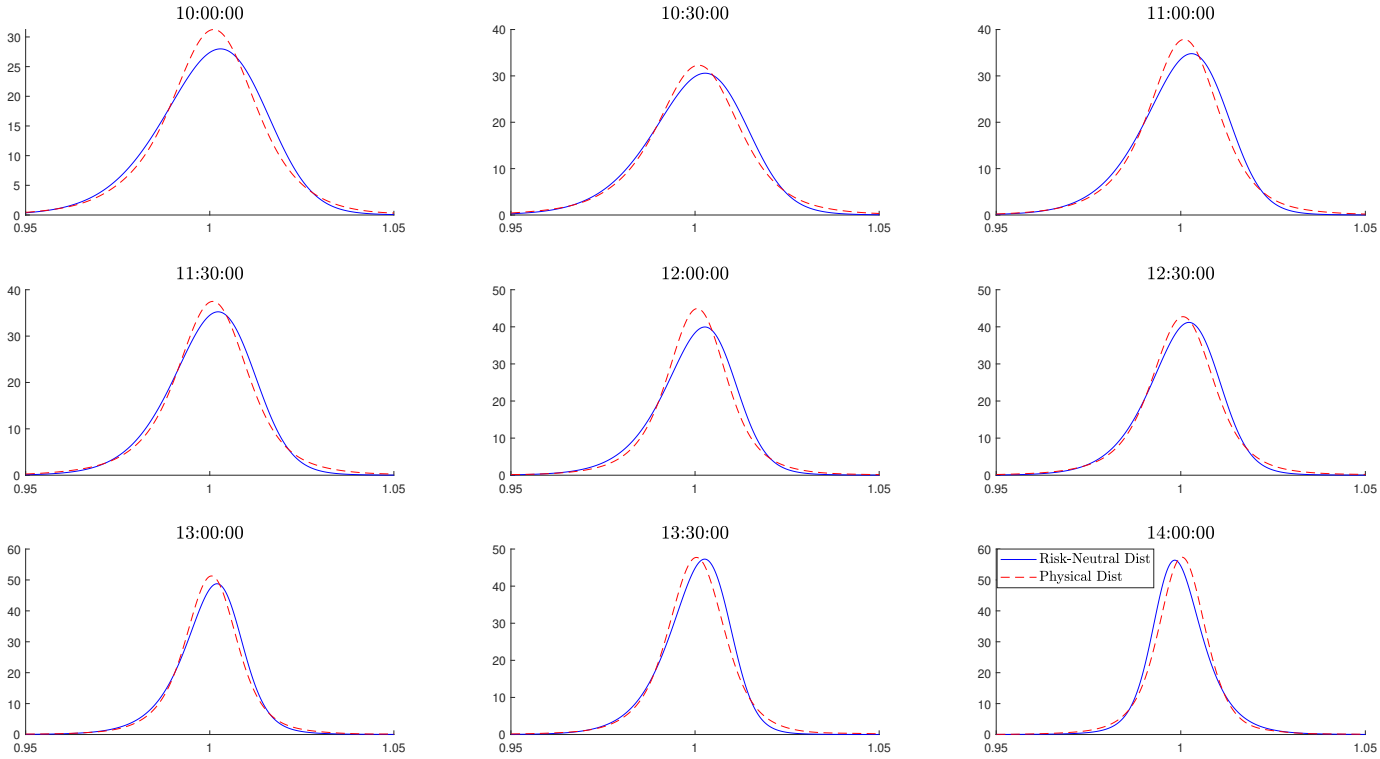

Note: The figure plots, for a representative date of our sample and for different times of the day, the market return risk-neutral distribution and physical distribution estimated as described in Section 3.3 and 3.4, respectively. The horizontal axis represents gross market return states.

Electronic copy available at: https://ssrn.com/abstract=4701401

Figure 6: 0DTE call returns across strikes

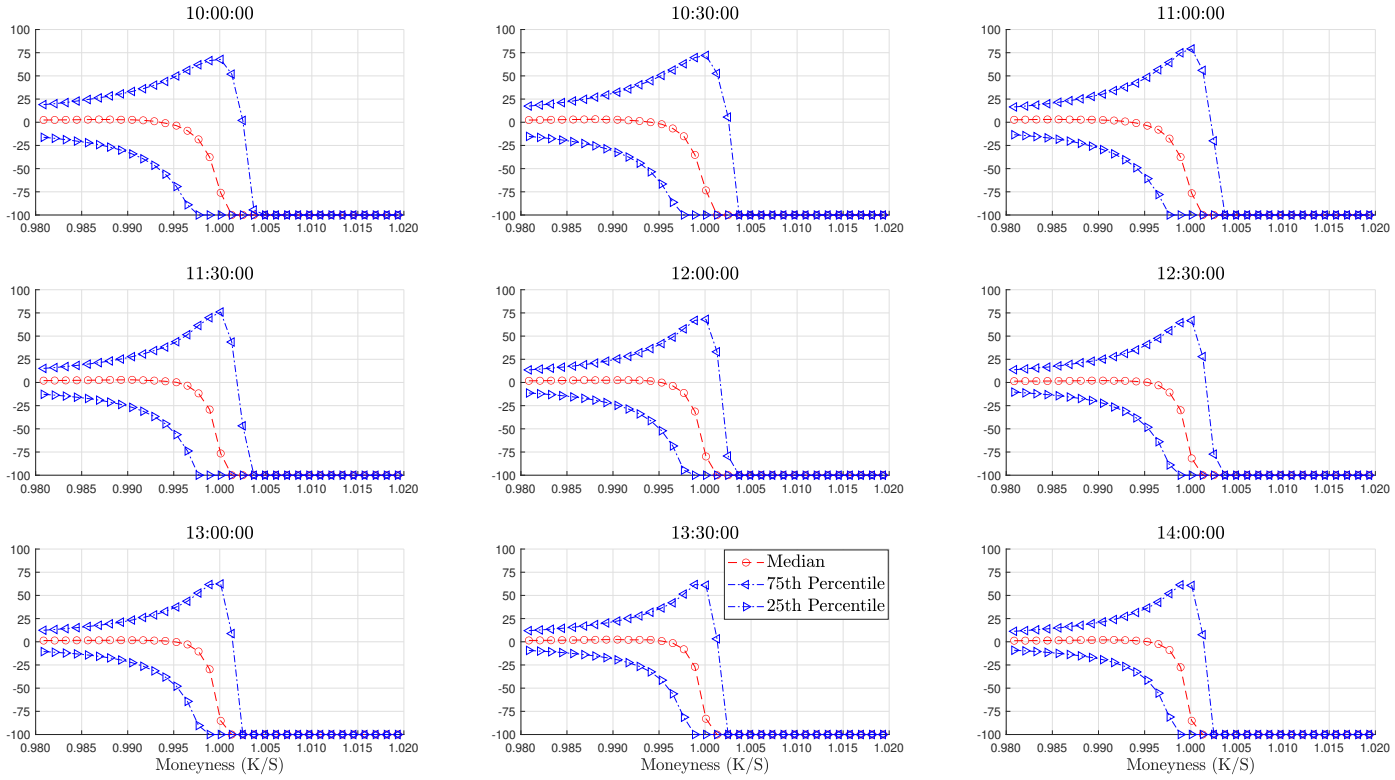

Note: The figure plots the 25th, median and 75th percentiles of returns (in %) of call options with different strikes over our sample. The sample ranges from January 6 2012 to July 3 2023.

Electronic copy available at: https://ssrn.com/abstract=4701401

Figure 7: 0DTE put returns across strikes

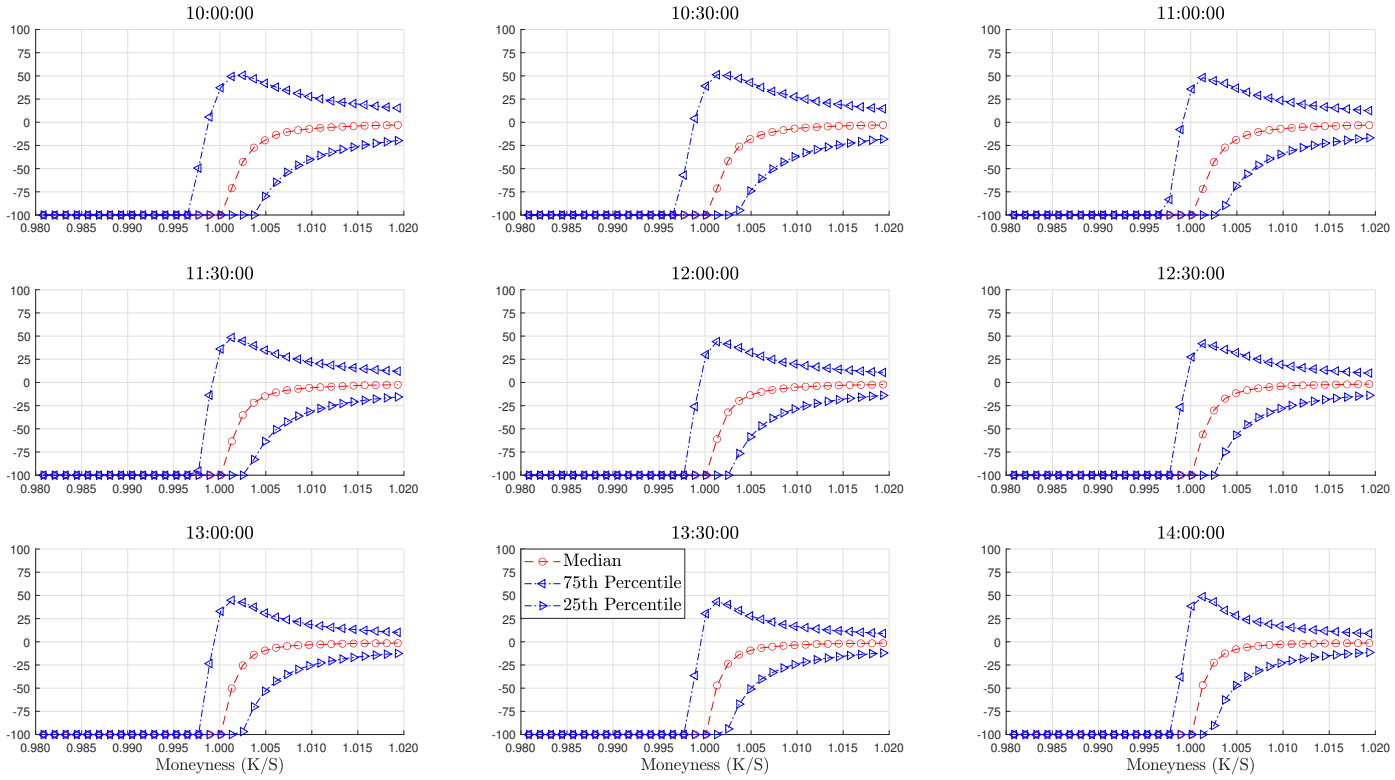

Note: The figure plots the 25th, median and 75th percentiles of returns (in %) of put options with different strikes over our sample. The sample ranges from January 6 2012 to July 3 2023.

Figure 8: Average returns of 0DTE option strategies

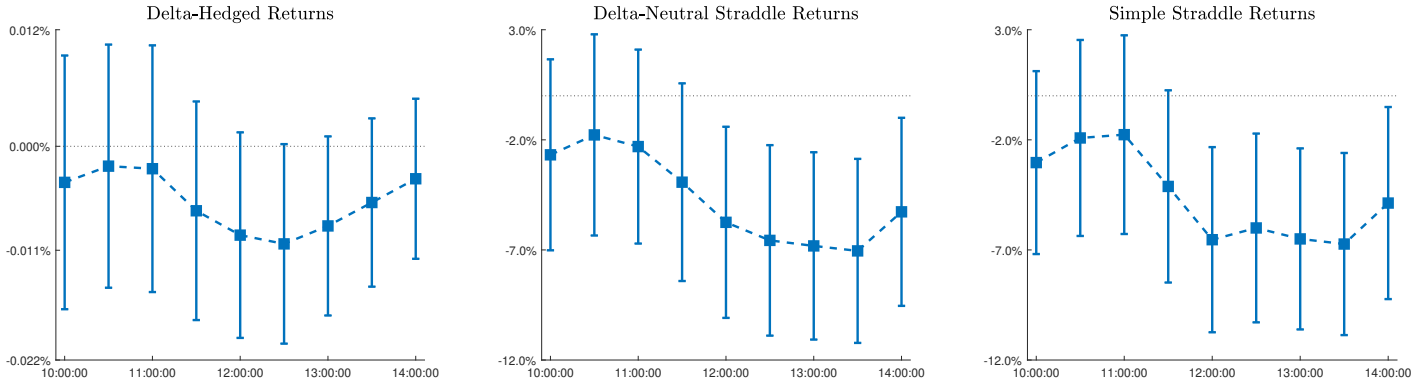

Note: The figure plots, for different times of the day, the average returns together with 95% confidence bands for ATM delta-hedged calls, delta-neutral straddles and simple straddles. Confidence bands are based on 2,500 bootstrap replications. The sample ranges from January 6 2012 to July 3 2023.

Electronic copy available at: https://ssrn.com/abstract=4701401

Figure 9: Intra-day equity premium decomposition

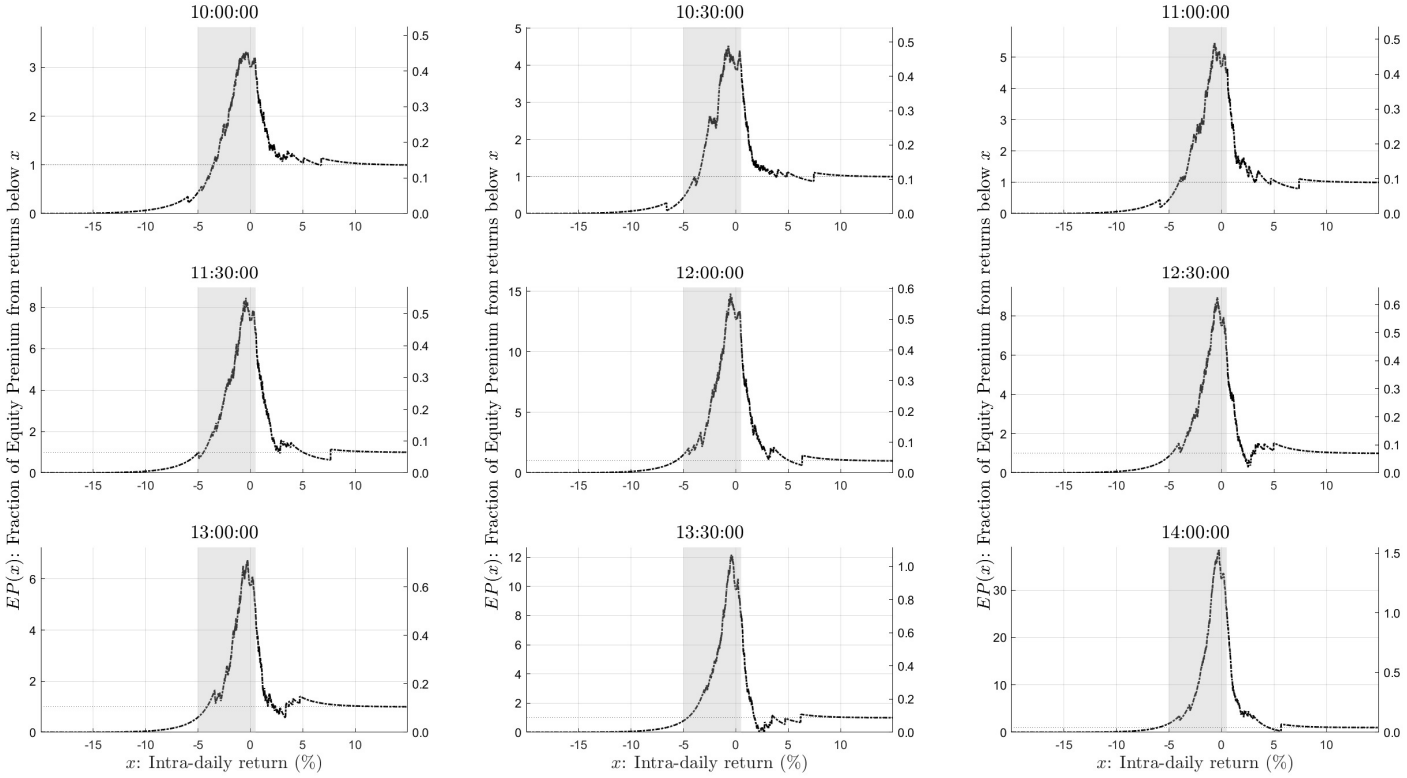

Note: The figure plots, for different times of the day, the intra-day equity premium decomposition implied by 0DTE options. The left vertical axis displays the fraction of the equity premium stemming from market returns below x, while the right vertical axis displays the total annualized equity premium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

Electronic copy available at: https://ssrn.com/abstract=4701401

Figure 10: Intra-day variance risk premium over time

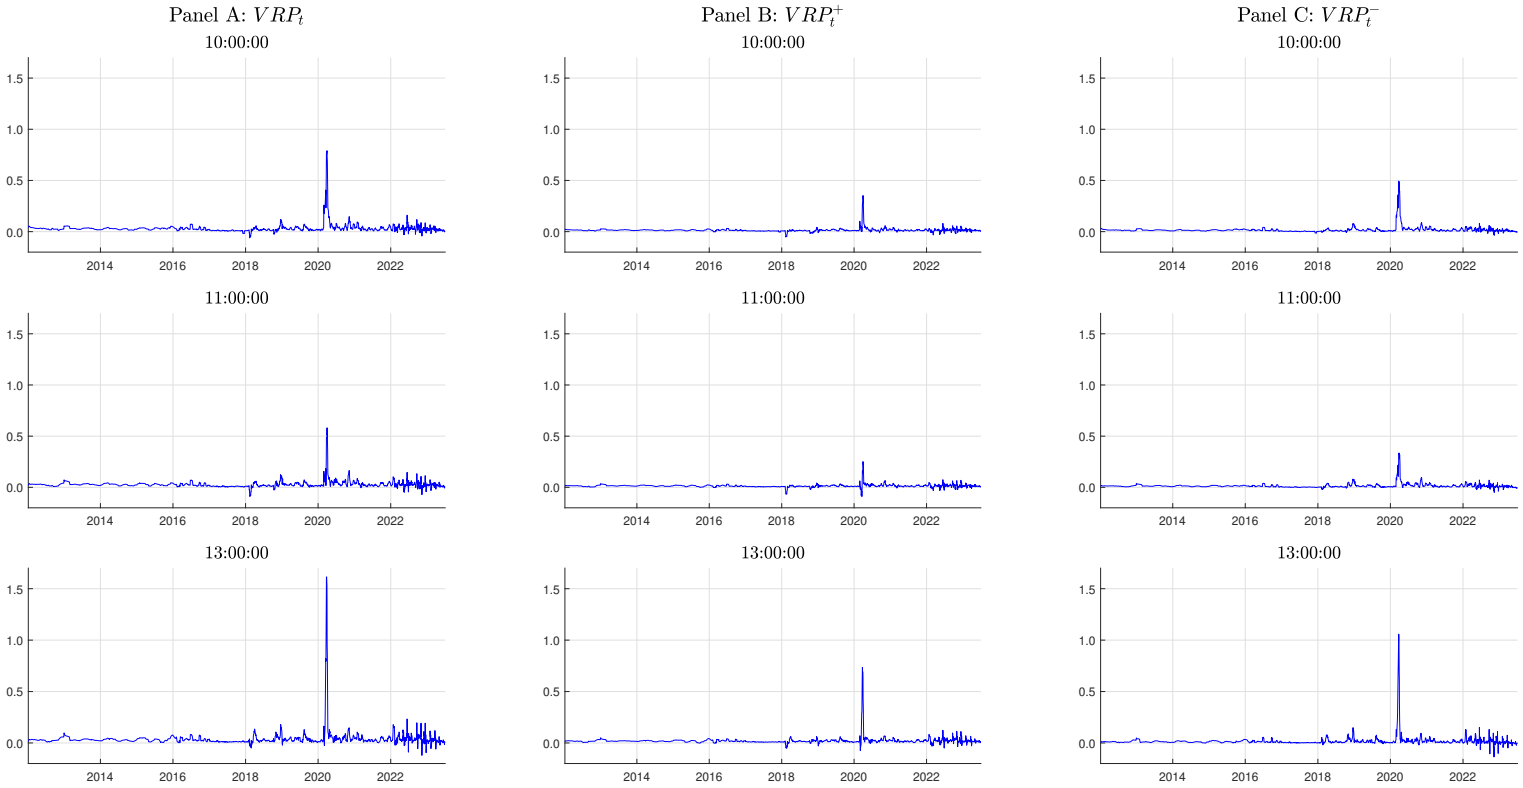

Note: The figure plots, for different times of the day, the one-week moving average (for ease of visualization) of the $VRP_{t,T}$, $VRP_{t,T}^+$ and $VRP_{t,T}^-$ over time. The measures are annualized. The sample ranges from January 6 2012 to July 3 2023.

Electronic copy available at: https://ssrn.com/abstract=4701401

Figure 11: Average intra-day pricing kernels over return states

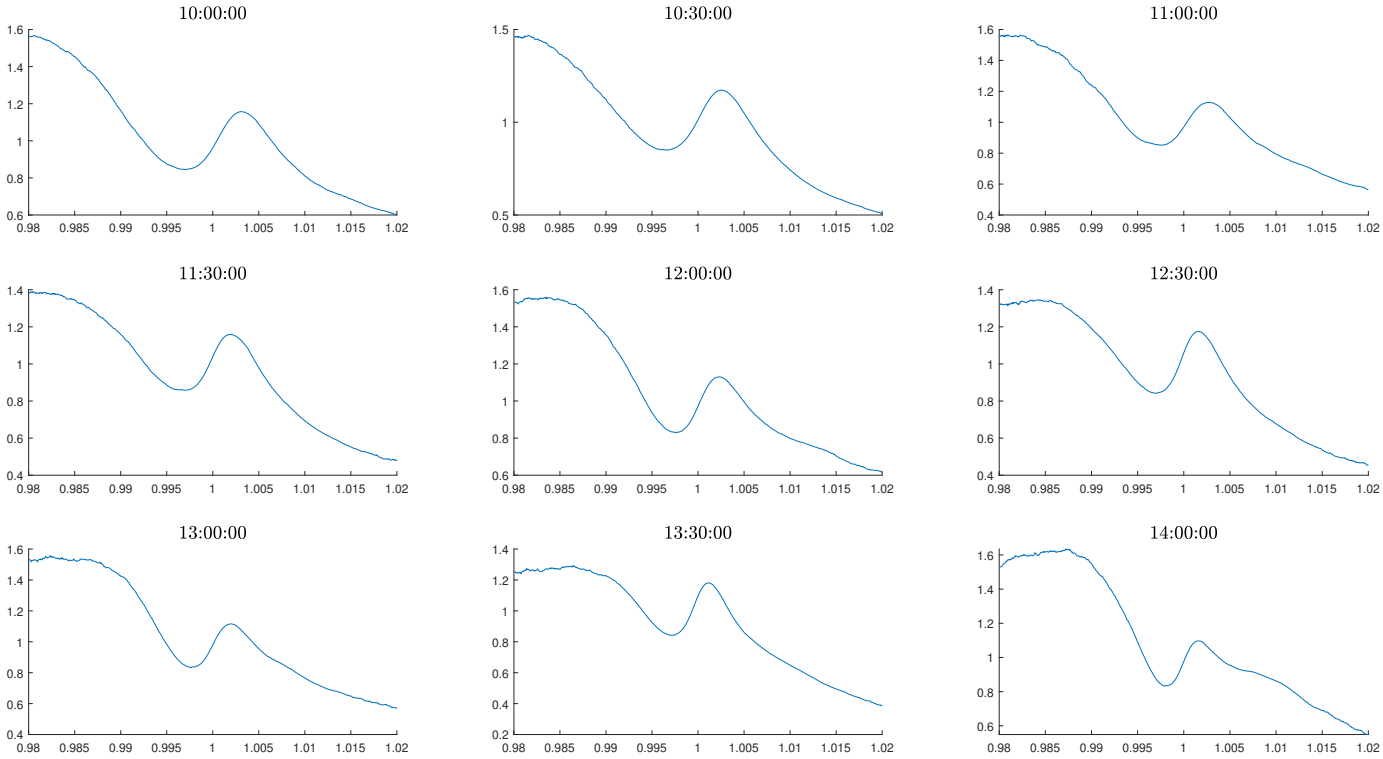

Note: The figure plots, for different times of the day, the average of the pricing kernel, as a function of market returns, over time. The sample ranges from January 6 2012 to July 3 2023.

Electronic copy available at: https://ssrn.com/abstract=4701401

Figure 12: Average intra-day pricing kernels over standardized log-moneyness

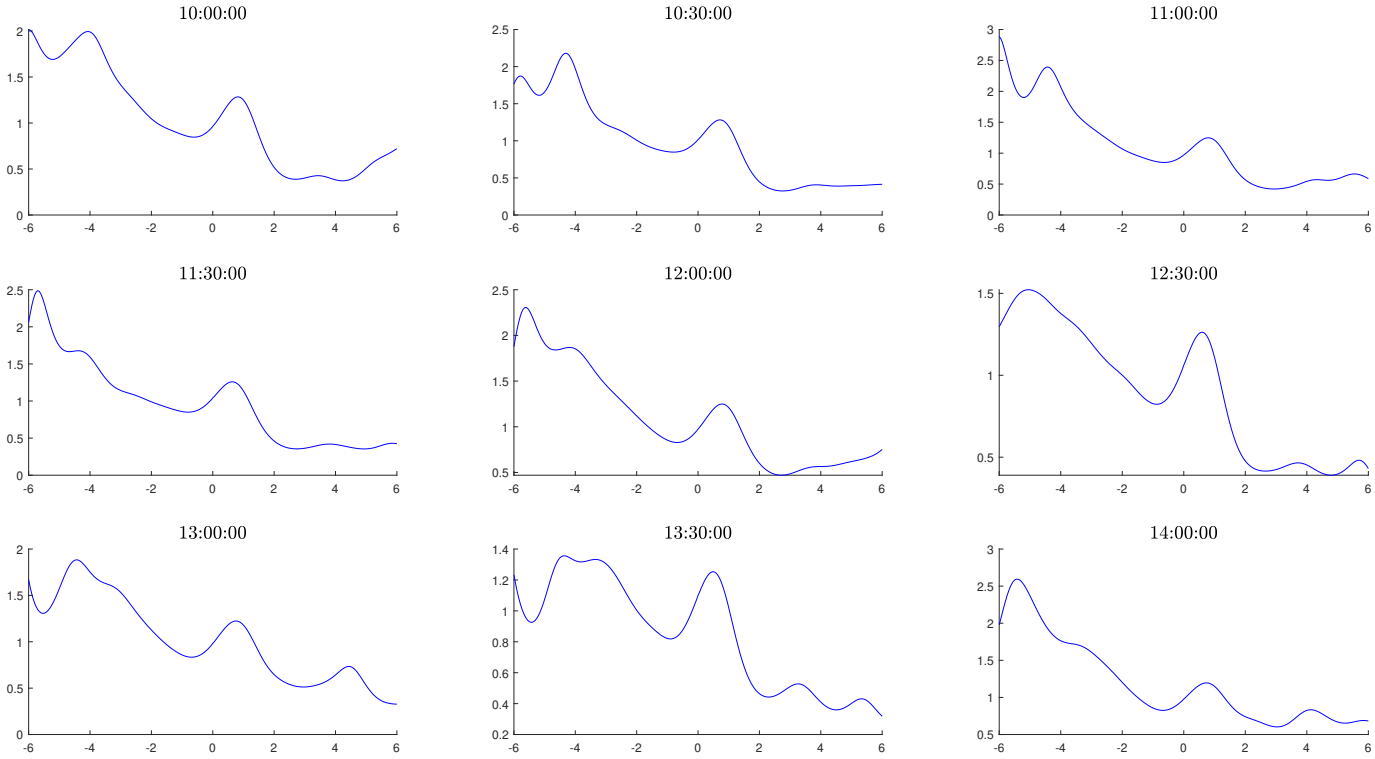

Note: The figure plots, for different times of the day, the average of the pricing kernel, as a function of standardized log-moneyness, over time. The sample ranges from January 6 2012 to July 3 2023.

Electronic copy available at: https://ssrn.com/abstract=4701401

Figure 13: Average SSD bounds and prices for 0DTE call options

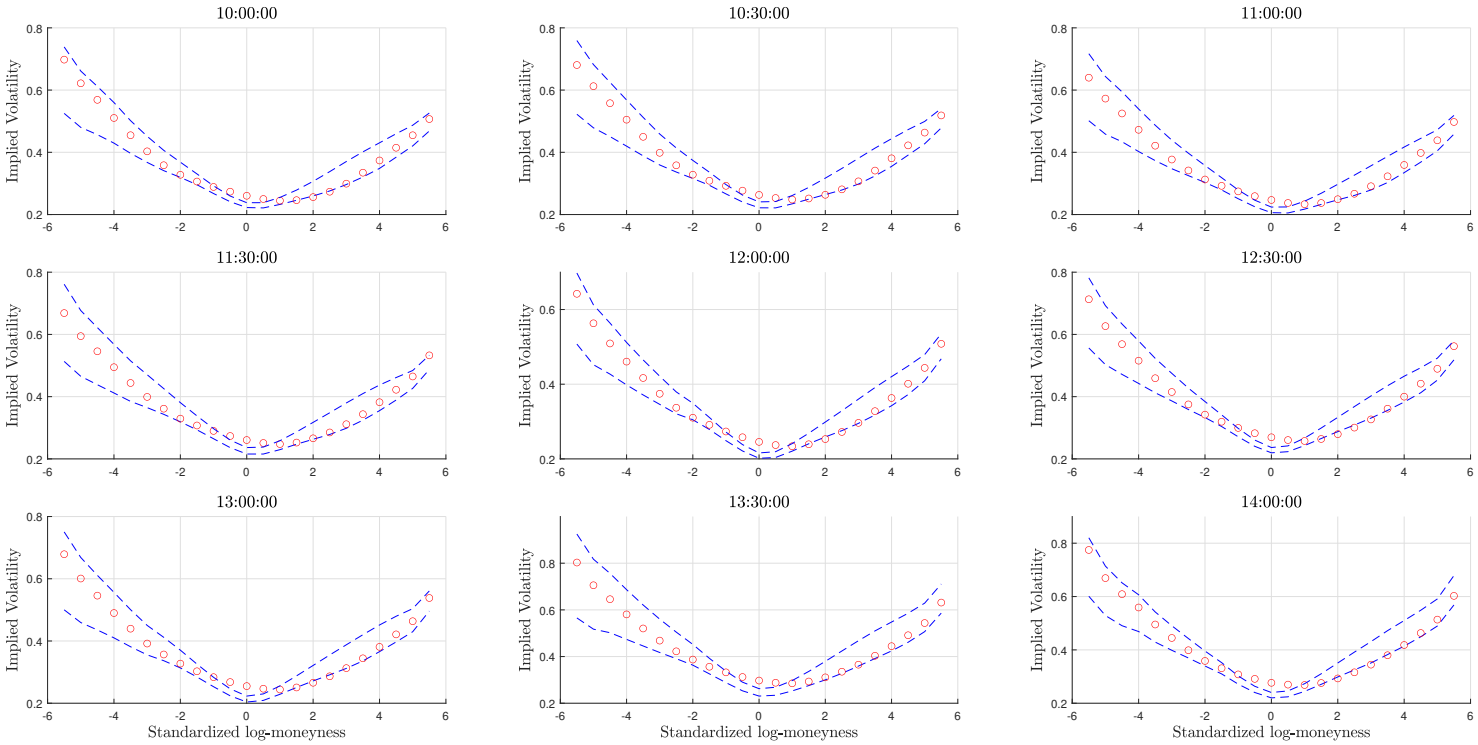

Note: The figure plots, for different times of the day, the average of the SSD option price bounds (in dashed blue) against the average option prices (in red circles) for calls, in implied volatility space, over standardized log-moneyness. The sample ranges from January 6 2012 to July 3 2023.

Electronic copy available at: https://ssrn.com/abstract=4701401

Figure 14: Average SSD bounds and prices for 0DTE put options

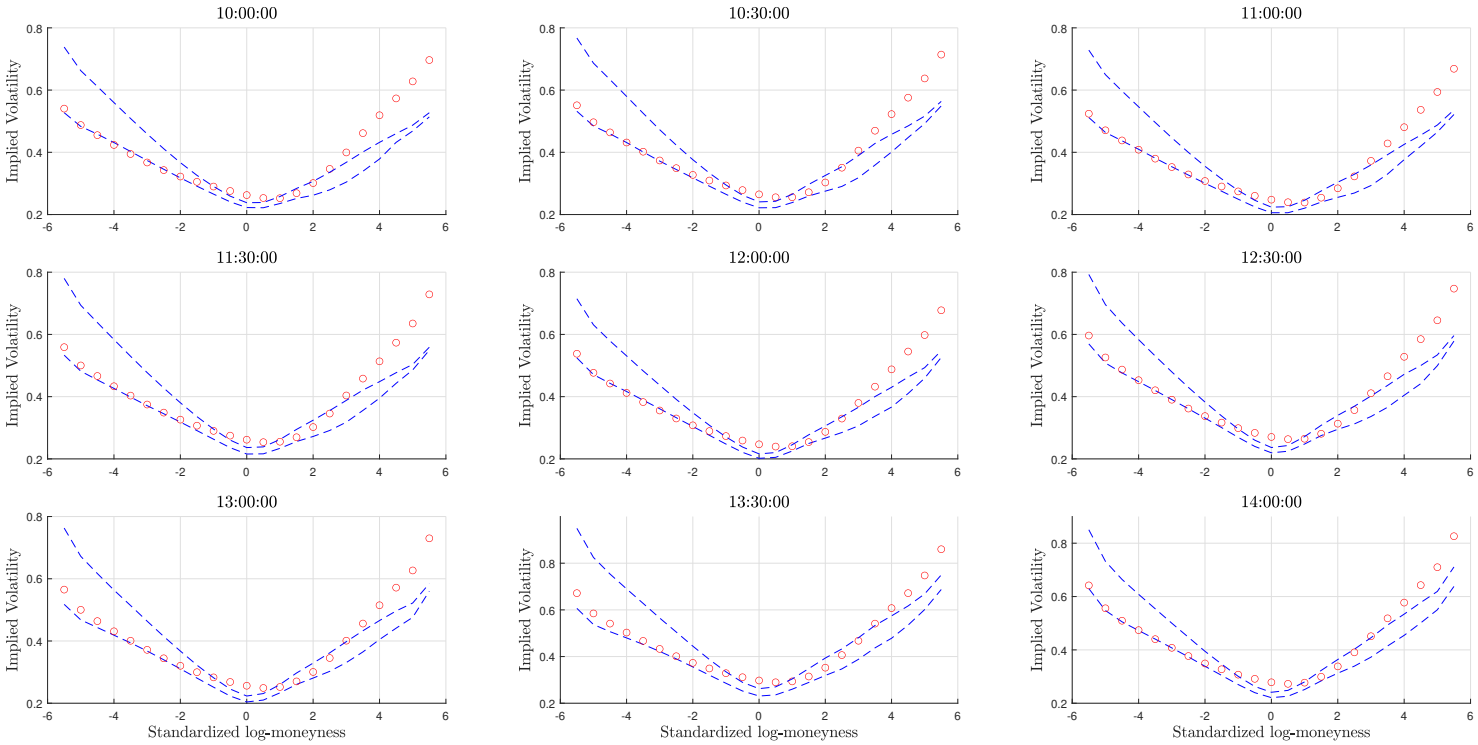

Note: The figure plots, for different times of the day, the average of the SSD option price bounds (in dashed blue) against the average option prices (in red circles) for puts, in implied volatility space, over standardized log-moneyness. The sample ranges from January 6 2012 to July 3 2023.

Electronic copy available at: https://ssrn.com/abstract=4701401

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: 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, ..., 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.

Electronic copy available at: https://ssrn.com/abstract=4701401

Table 3: 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										
	In	37.406	40.729	40.814	40.158	35.509	34.420	37.863	43.902	34.852
All Calls	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
	In	60.819	64.409	65.970	61.882	57.971	54.621	58.203	67.859	54.723
OTM	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
	In	7.715	8.378	8.557	9.925	7.126	7.709	8.862	9.200	5.204
ATM	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
	In	18.239	23.645	21.911	25.306	18.374	20.247	24.632	28.929	22.608
ITM	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										
	In	24.595	29.450	29.561	32.180	27.719	30.562	35.369	40.521	29.170
All Puts	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
	In	42.950	50.125	49.237	55.027	47.552	52.695	60.898	66.601	49.033
OTM	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
	In	7.296	8.024	7.881	8.904	5.550	6.344	8.338	9.676	4.812
ATM	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
	In	8.882	12.391	14.295	13.195	12.613	13.034	14.446	20.800	14.395
ITM	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.

Electronic copy available at: https://ssrn.com/abstract=4701401

Table 4: 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.

Electronic copy available at: https://ssrn.com/abstract=4701401

Table 5: 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
RV	0.022	0.050	0.036	0.046	0.025	0.028	0.036	0.060	0.052
t-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
$MFIS$	−0.005	0.000	−0.008	−0.013	−0.011	−0.004	−0.008	0.000	0.004
t-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
$MFIK$	0.003	0.009	0.009	0.007	0.003	−0.015	−0.011	−0.018	−0.013
t-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
$SVIX$	0.001	0.028	0.006	0.009	−0.010	−0.001	−0.034	0.010	0.005
t-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.

Electronic copy available at: https://ssrn.com/abstract=4701401

Table 6: 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^2_{adj} (%)	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.

Electronic copy available at: https://ssrn.com/abstract=4701401