

Risky Intraday Order Flow and Option Liquidity*

Hitesh Doshi[†] Paola Pederzoli[‡] Saim Ayberk Sert[§]

July 28, 2025

Abstract

This paper provides a novel analysis of marketwide and exchange-specific trading costs for short- and ultra-short-maturity options. We focus on inventory risk proxies related to order flow distribution and delta-hedging costs. Intraday order flow volatility emerges as the primary driver of spreads, while delta-hedging needs play a secondary role. Leveraging cross-exchange variation, we isolate this effect from broader factors that may jointly affect order flow volatility and spreads. The findings support models of active inventory management and suggest that, contrary to standard views, option liquidity providers rely more on trade matching than on delta-hedging to manage inventory risk.

Keywords: Bid-ask spread, order flow, order flow volatility, exchanges, big data

*We thank Rohit Allena, Pedro Garcia Ares, Vincent Bogousslavsky (discussant), Jean-Philippe Bouchaud, Kevin Crotty, David De Angelis, Kris Jacobs, Dmitriy Muravyev, Chay Ornthanalai (discussant), James Upson (discussant), Vipin Kerala Varma, and the participants to the 2024 Lonestar Conference, 2024 UQAM Seminar Series in Canada, 2024 CBOE-FMA Conference on Derivatives and Volatility, 2025 Midwest Finance Association Conference, and seminar at Capital Fund Management and the University of Houston for comments and suggestions. We thank SpiderRock Data & Analytics for providing data. This work was completed in part with resources provided by the Research Computing Data Core at the University of Houston. All errors are our own.

[†]C.T. Bauer College of Business, University of Houston. Email: hdoshi@bauer.uh.edu

[‡]C.T. Bauer College of Business, University of Houston. Email: ppederzoli@bauer.uh.edu

[§]C.T. Bauer College of Business, University of Houston. Email: sasert@bauer.uh.edu

1 Introduction

Investors are turning into short-maturity options, changing the standard trading dynamics in both the SPX options market and the market for individual stock options. In 2023, an impressive 80% of SPX options trading focused on options with expiration less than a month (Dim, Eraker, and Vilkov 2024; Bandi, Fusari, and Reno 2024). The surge in short-term options volumes has spurred new research exploring this novel market, its characteristics, and its implications for market stability. The main focus of these papers is, however, on the prices and returns of the options (Bandi, Fusari, and Reno 2024; Almeida, Freire, and Hizmeri 2024; Beckmeyer, Branger, and Gayda 2023), or on the impact of option trading on the underlying market (Dim, Eraker, and Vilkov 2024; Adams, Fontaine, and Ornthalalai 2024; Brogaard, Han, and Won 2023), with limited analysis on the quality of the market itself. Yet trading costs in this segment of the option market are exceptionally high. Even for liquid, at-the-money short-maturity options, bid ask spreads can reach up to 10% of the option price.

This paper contributes to the literature by analyzing effective trading costs in short- and ultra-short-maturity U.S. options and their relationship with order flow, with the goal of identifying trading patterns that may pose liquidity risk in a market characterized by exceptionally high volumes. We focus on two categories of order flow measures. First, we examine how the intraday distribution of order flow relates to bid-ask spreads. Intuitively, one-sided or highly volatile order flow may be risky and costly for liquidity providers, as it pushes them away from their inventory targets (see, e.g., Stoll 1978; Ho and Stoll 1983; Bogousslavsky and Collin-Dufresne 2023). Second, drawing on the literature on option market-making models (Jameson and Wilhelm 1992; Stoikov and Sağlam 2009), we consider measures that proxy for the delta-hedging costs potentially faced by liquidity providers who

manage their option exposure by trading in the underlying equity market.

Our empirical analysis encompasses the sample of options with a maximum maturity of seven weeks, including S&P500 index options (SPX options) and options on the constituents of the S&P500 index, spanning the period from 2004 to 2021. While SPX options are exclusively traded on the Chicago Board Options Exchange (CBOE), individual stock options are traded across sixteen exchanges, which enables us to conduct an additional analysis across venues.

We begin by analyzing the order flow of SPX options and document that, although buy and sell orders are generally balanced throughout the trading day, resulting in small daily order imbalances in absolute value (as shown by Dim, Eraker, and Vilkov 2024), the intraday pattern of order flow reveals much greater variation. To capture this, we divide the trading day into equispaced intraday intervals and compute the order imbalance within each interval. The resulting series of intraday order imbalances reflects the intraday distribution of order flow. We find that this distribution has exhibited high volatility since the financial crisis and in the years thereafter. Comparing days with high trading costs to those with low trading costs reveals that illiquidity is closely associated with elevated intraday order flow volatility. This suggests that more dispersed trading activity poses risks for liquidity providers and contributes to reduced market liquidity. Formal time series and panel regressions confirm this relationship: higher intraday order flow volatility on a given day is positively associated with higher trading costs. The effect is present across all option maturities but is particularly strong for ultra-short maturities, including options with zero days to maturity (0DTE) and options with one to six days to expiration.¹ This underscores the sensitivity of bid-ask spreads for these ultra-short maturity options to risky intraday trading patterns. Intuitively, for very short-term options, like 0DTE, liquidity providers do not have time to earn a premium on

¹Formally, 0DTE options are contracts traded on day t that expire at the end of the same day t .

their inventory (as in e.g., Fournier and Jacobs 2020) and must quickly adjust to order flow, incorporating the required premium immediately into the spread.

Overall, our analysis shows that the positive relationship between order flow volatility and illiquidity is highly robust and statistically significant. It applies to both the SPX options market and the market for individual stock options (in both the time-series and cross-sectional dimensions) and survives the inclusion of numerous controls, including daily measures of volume, order imbalance, volatility, option Greeks, stock characteristics and past spread levels. Importantly, all our regressions include time-fixed effects, such as day-of-the-week, month-of-the-year, and year dummies, to account for strong seasonalities in the spread.

To help isolate the effect of order flow volatility on spreads from that of common macro variables that may affect both, we leverage a unique feature of the U.S. equity options market: individual stock options are traded across sixteen exchanges. Trading activity is relatively well distributed, with no single venue dominating the market. For example, CBOE, the largest exchange, accounts for only about 20% of total volume. We begin by examining the cross-sectional dynamics of how liquidity adjusts following trades across the different exchanges. By tracking quoted spreads at the time of each trade, we find that transactions typically occur on the exchange offering the most favorable quoted spread, indicating that order-routing mechanisms help minimize transaction costs. In addition, we observe that after a trade is executed, the exchange that absorbs the trade tends to widen its quoted spread, while competing exchanges narrow theirs.² This behavior suggests an effort to attract order flow and points to a good degree of competition across U.S. option exchanges. The presence of multiple exchanges trading the same stock options creates a unique setting to examine

²In this analysis, we exclude trades whose size is equal to or larger than the quoted size at the best bid or ask, in order to avoid mechanical effects from trades that sweep the entire top-of-book depth.

how the relationship between trading costs and order flow volatility varies across venues for the same underlying asset on the same day. We exploit this setting and estimate panel regressions of exchange-specific effective spreads on exchange-specific order flow volatility, controlling for stock-by-day fixed effects. This approach helps us to isolate the relationship between intraday order flow volatility and illiquidity by absorbing all common shocks at the stock-day level. The results consistently confirm a positive and robust relationship, reinforcing our earlier findings.

Taken together, the results from both the marketwide and exchange-level analyses establish intraday order flow volatility as a key determinant of trading costs in short-maturity options. These findings raise a natural question: how do more traditional measures of inventory risk, particularly those associated with delta-hedging, compare in explanatory power? To explore this, we turn to a second set of inventory risk variables that specifically capture the hedging needs of liquidity providers. Using data from the CBOE Open-Close database, we analyze the positions of market-makers in the SPX options market, who are key liquidity providers and the participants most likely to delta-hedge their inventory. Following Dim, Eraker, and Vilkov (2024) and Ni, Pearson, Potoshman, and White (2021), we measure, for each day t in the sample, the intraday gamma of their inventory and the intraday delta of the new order flow they absorb on day t . These variables capture the extent of delta-hedge rebalancing for prior positions and the delta-hedging of new positions, respectively. We find a positive relationship between intraday gamma and trading costs for 0DTE options, and between the intraday delta of new order flow and trading costs for medium-maturity options, consistent with options theory (Jameson and Wilhelm 1992). However, when intraday order flow volatility is included in the regression, the significance of these variables weakens considerably or disappears altogether, suggesting that their explanatory power is overshadowed by the stronger relationship between volatile order flow and trading costs.

Interpreting our results through the lens of market microstructure theory (Glosten and Milgrom 1985; Stoll 1978), trading patterns that increase risk and cost for liquidity providers are naturally reflected in the bid-ask spread.³ A priori, our core variable, the order flow volatility, could reflect either informed trading or inventory risk. However, our findings suggest that it is more closely associated with inventory-related liquidity shocks. Specifically, we observe that order flow volatility is lower on days with a higher likelihood of informed trading, such as earnings announcements and the preceding day. In contrast, it spikes for very short-term options on the third Friday of each month, when many contracts expire, and for medium-term options on the first and last trading days of the month, patterns more consistent with inventory rebalancing. Supporting this interpretation, we also find that order flow volatility tends to be lower when stock volatility is high, further suggesting that it is not driven by market uncertainty either, and it has to be better interpreted as a measure of inventory risk. Given this, our analysis naturally aligns more closely with inventory risk models.

Inventory-based theory in the stock market that explicitly account for the stochastic nature of the order flow (e.g., Bogousslavsky and Collin-Dufresne 2023) emphasize the critical role of unbalanced order flow distribution and its positive relationship with illiquidity, consistent with our findings. In this model, liquidity providers actively rebalance inventory throughout the day, aiming to balance buy and sell orders and maintain a small inventory. Intraday volatile order flow increases the inventory risk they face while awaiting offsetting trades, leading to wider spreads for investors.

While this framework is well established in equity markets, inventory models in the options market typically take a different approach. They emphasize the role of delta-hedging,

³See Foucault, Pagano, and Röell (2013) for a review of models where bid-ask spreads endogenously compensate for inventory risk, asymmetric information, or order processing costs.

whereby liquidity providers manage inventory risk by offsetting their option positions with trades in the underlying asset (Jameson and Wilhelm 1992; Stoikov and Sağlam 2009; Cho and Engle 1999). These models predict that trading costs should reflect the risks and frictions associated with discrete delta-hedge rebalancing. However, recent empirical evidence from Hu, Kirilova, Muravyev, and Ryu (2024), based on detailed account-level data from market makers in the Korean options and futures markets, indicates that liquidity providers rely primarily on active inventory management and trade matching, with delta-hedging serving only as a secondary risk management tool. Our findings support this perspective: intraday order flow volatility emerges as the dominant determinant of trading costs, outweighing the role of delta-hedging variables. This suggests that option liquidity providers primarily manage inventory risk through active inventory management, using delta-hedging as a secondary tool.

Our paper contributes to different strands of literature. It is primarily related to the recent literature that studies the novel market of short-maturity options and ultra-short maturity options (Almeida, Freire, and Hizmeri 2024; Bandi, Fusari, and Reno 2024; Dim, Eraker, and Vilkov 2024; Beckmeyer, Branger, and Gayda 2023; Adams, Fontaine, and Ornthanalai 2024). The novel aspect of our investigation is the focus on option liquidity and its relation to the order flow.

In a closely related paper on option liquidity, Christoffersen, Goyenko, Jacobs, and Karoui (2018) document a substantial illiquidity premium in the option market for longer maturity options. They also analyze the determinants of the bid-ask spread in the cross-section of options and find that the daily absolute value of the order imbalances from non market-makers are positively related to illiquidity. Goyenko, Ornthanalai, and Tang (2015) also examine the link between spreads, order imbalance, and delta-hedging costs in the equity option market, while Cao, Jacobs, and Ke (2024) explore how this relationship varies throughout the trading

day in SPX options. Unlike these studies, we focus on the second moment of the order flow distribution and examine its relationship with illiquidity in short-maturity options, both at the market level and across exchanges.⁴

In this sense, our paper is closely related to Bogousslavsky and Collin-Dufresne (2023), who also examine the second moment of order flow distribution, though in equity markets. Our contribution is distinct along two key dimensions. First, while their focus is on the volume-spread puzzle, we take a broader approach to understanding illiquidity in short-maturity options by comparing the explanatory power of theoretically motivated variables. This allows us to test competing mechanisms underlying option trading costs. Second, we exploit unique features of the option market, its maturity structure and multi-exchange environment, to provide new economic insights not available in the equity setting.

2 Theoretical Framework and Literature Review

According to standard models of inventory management (e.g., Ho and Stoll 1983; Grossman and Miller 1988), liquidity providers set the bid-ask spread in the market to maximize their utility based on their final wealth. This wealth is determined by the cash earned from the bid-ask spread and their inventory position. These models typically assume utility functions that reflect an aversion to inventory variance, leading to wider bid-ask spreads as inventory

⁴Additional research on option bid-ask spread includes early studies such as George and Longstaff (1993) and Cho and Engle (1999), as well as more recent work on order-routing mechanisms across venues (Huang, Jorion, and Schwarz 2024) and the role of payment for order flow (Ernst and Spatt 2022). Other related studies has examined order-flow measures of trading and their impact on option prices, e.g., Bollen and Whaley (2004); Garleanu, Pedersen, and Potoshman (2008); Muravyev (2016); and Fournier and Jacobs (2020). Unlike these studies, which focus on the first moment of the order flow distribution, our research centers on the second moment and its impact on liquidity. There is also an extensive literature investigating the impact of the order flow on stock returns. A non exhaustive list is e.g., Chordia and Subrahmanyam (2004); Kelley and Tetlock (2013); Brogaard, Hendershott, and Riordan (2014); Chordia, Hu, Subrahmanyam, and Tong (2019).

risk increases. This aversion stems from the preference of liquidity providers to maintain a minimal and balanced inventory throughout the day, as holding a non-zero position between offsetting trades is risky.

This is best formalized in the model proposed by Bogousslavsky and Collin-Dufresne (2023). In their model, liquidity providers actively manage their inventory in the stock market by matching buy and sell orders to minimize imbalances. The bid-ask spread compensates these risk-averse liquidity providers for the inventory risk they face while awaiting offsetting order flow. If the arrival rates of buy and sell orders temporarily diverge during the day, it creates an unbalanced order flow, increasing inventory risk for liquidity providers. Consequently, holding trade volume constant, the equilibrium bid-ask spread rises with intraday volatility in order flow. Intuitively, in this framework, a large volume of shares bought in one period and sold in another entails more risk than smaller, continuous transactions spread evenly throughout the day.

In the options market, liquidity providers can manage inventory risk in two primary (non mutually exclusive) ways: i) by actively managing inventory and matching trades, as described above, or ii) by delta-hedging their option positions in the underlying stock market. Hu, Kirilova, Muravyev, and Ryu (2024), who analyze account-level data for market-makers in options and futures on the Korean Composite Stock Price Index (KOSPI 200), find that most market-makers do not delta-hedge their option inventory. Instead, they rely on active inventory reversal strategies as in (i), eliminating undesired positions within minutes. In such scenarios, intraday order flow volatility is expected to be a primary determinant of bid-ask spreads.

The first hypothesis tested in our analysis is whether the volatility of order flow is positively related to the effective spread. Acceptance of this hypothesis would suggest active inventory management by liquidity providers in the U.S. options market. However, this

would not preclude the possibility that liquidity providers also utilize other inventory management tools, such as delta-hedging, as in (ii). In such cases, additional factors are expected to influence the bid-ask spread, particularly when perfect inventory hedging is unattainable.⁵ The risk associated with discrete delta-hedge rebalancing is best represented by the inventory’s gamma, which measures the sensitivity of the delta position to price changes, thereby reflecting the rebalancing needs of liquidity providers (Jameson and Wilhelm 1992; Ni, Pearson, Potoshman, and White 2021; Dim, Eraker, and Vilkov 2024). Gamma also accounts for errors in discrete rebalancing caused by changes in the delta position. We hypothesize that if market-makers actively engage in inventory management, as observed in Hu, Kirilova, Muravyev, and Ryu (2024), inventory gamma will play a secondary role compared to order flow volatility in driving bid-ask spreads.

3 Data

We obtain options trade data from the CBOE’s LiveVol, including timestamp down to milliseconds, trade price and size in contracts, the prevailing NBBO prices, and the contemporaneous best bid and offer prices of underlying security for each trade reported by the Options Price Reporting Authority (OPRA). The dataset spans the intraday trading activity of all equity and index options from January 01, 2004, to July 16, 2021. We merge the LiveVol data with the Center for Research in Securities Prices (CRSP), from which we obtain daily stock returns, trading volumes, prices, and the number of outstanding shares. Additionally, we combine the intraday trade data with OptionMetrics, allowing us to access daily implied volatility and Greeks for option series. For each day, option series are required to be present

⁵Traditional models of liquidity providers in the option market (e.g., Stoikov and Sağlam 2009; Cho and Engle 1999) suggest that order flow imbalances and inventory should not impact spreads if perfect delta-hedging is achievable.

in all three data sources.

We focus on S&P500 index options and options on individual stocks which are the constituents of the S&P500 index. We track S&P 500 constituents on a monthly basis following the historical components file from CRSP. A stock is included in our cross-sectional sample for a given month if it was part of the S&P500 index in the previous month.

Our focus lies on short-term options with maturities of up to one month, as these have seen the most significant growth in trading activity over time (Almeida, Freire, and Hizmeri 2024), raising questions about the stability of the option market. Among these options, at-the-money (ATM) options are of special interest, as they have the highest decline in value as maturity approaches, and the highest value of gamma, which is particularly relevant for delta-hedgers liquidity providers (see Ni, Pearson, Potoshman, and White 2021). Moreover, the prices and spreads of ATM options are less affected by market microstructure noise than out-of-the-money (OTM) options (Duarte, Jones, and Wang 2024).⁶ Our main sample is thus composed by ATM options, defined by an absolute delta between 0.375 and 0.625, with up to 48 days to maturity. The delta of each option series is assessed at the close of the preceding business day; for example, an option on day t is considered at-the-money if its absolute delta, as recorded by OptionMetrics at the close of day $t - 1$, falls between 0.375 and 0.625.

We examine all options trades recorded by OPRA between 9:30 a.m. and 4:00 p.m. US Eastern time. The OPRA database encompasses trades occurring across the sixteen exchanges where investors can trade options. SPX index options are specifically traded only on the CBOE exchange, with regular trading hours concluding at 4:15 p.m. Additionally, they are available for trading during global trading hours before the market opens and after

⁶In the robustness Section 6, we analyze out-of-the-money options and find results consistent with those observed for the at-the-money options in the baseline analysis.

it closes, with this time frame gradually expanding over time.⁷ To ensure consistency in coverage across various securities and over time, we concentrate on the standard trading hours of 9:30 a.m. to 4:00 p.m. for all underlying stocks and the S&P500 index.

Following the literature, we apply filters to the intraday trade data to clean obvious errors and outlying records. We filter out the following observations: (1) cancelled trades; (2) trades with zero or negative price, size, and/or bid-ask spread; (3) trades whose sizes are higher than 100,000 contracts; (4) trades whose prices are below bid minus spread or above ask plus spread; and (5) trades whose prices are below \$0.10.

4 Empirical Results

This section explores the characteristics of intraday order flow distribution in short-term at-the-money options and its relation with illiquidity. Sections 4.1, 4.2, 4.3, and 4.4 focus on the order flow of SPX options, while Section 4.5 examines the cross-section of options on individual stocks. Section 4.6 performs the analysis at the exchange level.

4.1 Order Flow and Daily Statistics

Our primary focus is on analyzing the distribution of intraday order flow. To achieve this, we first need to flag every trade as buy (i.e., buyer-initiated) versus sell (i.e. seller-initiated), since the OPRA data does not explicitly provide this information.

Following the literature on high-frequency data of trades and quotes of stocks (Lee and Ready 1991; Bogousslavsky and Collin-Dufresne 2023), trades are categorized as buys or sells based on the quote rule and tick rule. Specifically, if a trade price is closer to the National

⁷In 2015, CBOE extended trading hours for SPX options to include 3 a.m. to 9:15 a.m. In 2021, the start time moved to 8 p.m. of the day before, and in 2022, CBOE added the ‘Curb’ session from 4:15 p.m. to 5 p.m.

Best Offer, it is classified as a buy; otherwise, it is classified as a sell. If a trade price falls at the NBBO quote midpoint, we follow Bryzgalova, Pavlova, and Sikorskaya (2023), and apply the quote rule to the Best Bid and Offer (BBO) prices from the exchange where the trade was executed. In cases where the trade price equals the BBO mid price, the tick rule is applied: if the current trade price exceeds the price of the last trade in the same option, the current trade is classified as a buy; conversely, it is classified as a sell.

In the stock market, it is well-known that the quote rule effectively classify trades that occur without any price improvements, resulting in buyer-initiated (seller-initiated) trade prices that are very close to the quoted ask (bid) prices. However, when a trade receives significant price improvement, the trade classification may be prone to misclassification (Ellis, Michaely, and O’Hara 2000). To validate our quote rule on this critical sample, we obtain a sample of about one million option trades executed on 2024-02-02 through auctions.⁸ These trades are mostly retail orders which have been automatically routed into auctions to receive the best price improvement. Within the auction database, we have access to the actual trade direction (buy versus sell) along with the prevailing bid and ask quotes of the exchange where the trade occurred. Analysis reveals that, in this sample, the quote rule successfully classifies approximately 85% of the trades.⁹

We then partition the trading day into equispaced time-intervals, and calculate the option order flow on day t in each interval d by subtracting the trade size of seller-initiated trades

⁸We thank SpiderRock Data & Analytics for providing this auction data.

⁹Another potential source of misclassification could occur with trades that are components of multi-leg strategies. Li et al. (2020) propose an heuristic approach to classify such trades. However, this methodology, relying on manual trade matching, cannot be verified without a sample containing the actual trade direction. Additionally, Li et al. (2020) find that in their sample, 70% of vertical spreads and 60% of straddles can be classified using the quote rule. Therefore, we opt to adhere to the standard quote rule for trade classification.

of all options i from that of buyer-initiated trades:

$$\text{Order Flow}_{t,d} = \sum_i \text{Trade Size of Buys}_{i,t,d} - \sum_i \text{Trade Size of Sells}_{i,t,d}. \quad (1)$$

Several choices for the length of the time intervals are possible. The optimal choice balances the need for high frequency data and option liquidity; if the intervals are too short, we risk having many empty intervals due to insufficient trading activity. While this might not be an issue for SPX options, it could be problematic for some individual stock tickers. Therefore, we opt for a 5-minute interval, which provides a suitable balance as an intermediate high frequency. The first interval spans from 9:30 am to 9:35 am, while the final interval spans from 3:55 pm to 4:00 pm, and in total we have 78 intervals per day.

To obtain the daily order flow, which we label order imbalance and denote it with the variable OI_t , we sum the order flows across the intra-day intervals:

$$OI_t = \sum_d \text{Order Flow}_{t,d}. \quad (2)$$

The order flow measures the buy versus sell pressure in the market. It is positive when investors are, overall, buying more options than selling them, and negative otherwise.

Finally, we calculate the daily options volume by summing the number of contracts traded across all option series:

$$\text{Volume}_t = \sum_i \text{Trade Size}_{i,t}. \quad (3)$$

[Figure 1 here]

Figure 1 displays the average daily volume and order imbalance for at-the-money put and call options in each year of the sample period. Panel A1 confirms the well-known upward trend in SPX option volumes since the years 2012-2013, observed in both call and put

options. Panel A2 documents some important characteristics of the daily order imbalances. On average, the order flow is positive for SPX put options and negative for SPX call options, displaying some variability across the years; this trend corresponds with findings from Chen, Joslin, and Ni (2019) and Jacobs, Mai, and Pederzoli (2024), among others. In the aftermath of the financial crisis, the order flow size surged, reaching an average of 2000 contracts as net order flow per day in 2010 (positive for put options and negative for call options). Post-crisis, the daily order flow size remained relatively stable with occasional deviations. For instance, during the years 2015 or 2018, we observe a modest average daily order flow in both call and put options. Particularly noteworthy are the last two years of our sample, 2020 and 2021, where we document an average negative order flow for both call and put options, with a magnitude around 2000.

Overall, the graph illustrates that, despite the surge in option volumes, buy and sell orders remain relatively balanced throughout the day, resulting in no significant increase in the overall size of the daily net order flow, consistent with results documented by exchange analysts and recent literature.¹⁰ The next section will offer a new perspective on order flow patterns by analyzing the intraday distribution, revealing that even when the daily order flow is small, there can be substantial intraday variation.

4.2 Intraday Order Flow Distribution

In this section, we start our novel analyzes of the intraday distribution of the order flow. Every day we calculate mean, standard deviation, skewness, and quartiles ($q_{0.25}$, $q_{0.5}$, and $q_{0.75}$) of the seventy-eight 5-minute intervals order flows calculated according to Equation 1.

[Table 1 here]

¹⁰See, for example, <https://www.cboe.com/insights/posts/volatility-insights-evaluating-the-market-impact-of-spx-0-dte-options/>

Panels A1 and B1 of Table 1 present the average of the daily statistics over the years for ATM SPX call and put options. Figure 2 complements Table 1 by illustrating the time-series of the average 5-minute order flow with intraday confidence intervals.¹¹

[Figure 2 here]

The intraday buy and sell orders are largely balanced over the sample period, with the average 5-minute order flow across years being -6 for ATM call options and 9 for ATM put options. These averages vary across years, ranging from a minimum of -28.8 (recorded in 2020 for ATM calls) to a maximum of 32.7 (recorded in 2016 for ATM puts). However, as shown in Figure 2, the mean 5-minute order flow does not exhibit any discernible time trend. Low skewness estimates across all years further highlight the overall symmetry of the intraday order flow distribution, which is confirmed by the median and 0.25–0.75 quartiles. Standard deviations, in contrast, are quite large, ranging from 229.7 (in 2004 for ATM calls) to 1764.7 (in 2011 for ATM puts). This results in wide confidence intervals for the average 5-minute order flow. For example, in 2011, the average 5-minute order flow for put options is 6.2 contracts, but with a standard deviation of 1764.7, the confidence interval spans $[-385, 398]$ contracts, reflecting substantial variability in intraday order flow. Examining the time-series of the average standard deviation by year, depicted in Figure 2, we find that the distribution initially exhibited a higher degree of concentration in the early years of the sample. Subsequently, it became more dispersed during the financial crisis in 2007, and, for ATM call options, it then stabilizes with some notable spikes around 2018. For ATM puts, the pattern is similar, with notable spikes in 2011 (concurrent to the European financial crisis), and 2018 (concurrent with the Volmageddon incident).

¹¹Specifically, for every day in the sample, we compute the average intraday 5-minute order flow, μ_t , with its confidence interval $\mu_t \pm Z \frac{\sigma_t}{\sqrt{n}}$, where σ_t is the standard deviation of the intraday 5-minute order flows. The figure displays the monthly averages of these daily quantities.

In summary, this analysis shows that, beginning with the financial crisis in 2007, the distribution of intraday order flow has remained stable over the years. It exhibits high symmetry but also a very high level of standard deviation. Notably, in the ATM put market, this standard deviation peaks during years marked by significant turbulence in volatility markets.

To gain a preliminary insight into the relationship between intraday order flow distribution and option market quality, we compare the distribution of intraday order flow during days characterized by high transaction costs with those characterized by low transaction costs. Our goal is to identify the distribution characteristics that are significant for liquidity.

In accordance with Christoffersen, Goyenko, Jacobs, and Karoui (2018) and Bogouslavsky and Collin-Dufresne (2023), we measure the cost of trading options with the effective spread incurred by option traders. Specifically, for each trade i on day t , we define the percent effective spread as:

$$\text{Effective Spread}_i = 2|\ln P_i - \ln M_i| \quad (4)$$

where P_i is the price of the trade i and M_i is the prevailing midpoint of the NBBO. For each day, the daily effective spread is the volume-weighted average of effective spreads across trades within the same option category (ATM calls and puts).

[Figure 3 here]

Panel A of Figure 3 displays the time series of the daily effective spread (ES_t) across the entire sample period for our samples of ATM SPX call and put options. The graph illustrates a downward trend in the spread throughout the sample period, along with recurrent spikes that may suggest seasonal patterns in both the spread and the daily changes in the spread.

We will account for seasonalities and time-trends in the regression analysis of Sections 4.3, 4.4, and 4.5.

We compare the intraday distribution of order flow on days characterized by low and high trading costs as follows: for each year in the sample, we identify the days falling in the bottom 10% and top 10% based on their ES_t values.¹² We then calculate the summary statistics (mean, standard deviation, skewness, and quartiles) shown in Table 1 for each of these subsamples. Panels A2 and B2 of Table 1 present the difference in these statistics between days with low and high transaction costs, segmented by year.

The results are qualitatively similar across the years for both call and put options markets. Days with low transaction costs have a distribution of intraday order flow that consistently shows lower standard deviation and smaller interquartile range compared to days with high transaction costs. Meanwhile, the distribution remains symmetric and with a small mean in both subsamples, as evidenced by the minimal change in skewness and mean values. The table also reports the results of testing whether the differences reported are statistically significant within each year. We find that for the majority of the years, the differences in standard deviations and first and third quartiles are statistically significant. None of the other statistics show the same consistent pattern. The table also reveals no time-trend in the difference between the standard deviation of order flow on days with low and high trading costs, indicating that extreme distribution days have not become more pronounced over time. However, the current high levels of volumes in the option market represent a mass of traders which could potentially generate a very volatile order flow. This underscores the importance of understanding the implications of volatile intraday order flow distributions.

In summary, the findings of this section suggest that the distribution of intraday order

¹²Similar results are obtained when splitting the sample according to ΔES_t instead of ES_t . Results are provided in Table IA.1 in the Online Appendix.

flow holds significant economic implications for market liquidity. Specifically, days in which the average 5-minute order flow is more volatile, as measured by the standard deviation of the distribution and the interquartile range, appear to coincide with days with low option market liquidity. Next section formally tests this pattern through a regression analysis.

4.3 Volatile Order Flow and Option Market Liquidity

In this section, we conduct a formal examination of the relationship between option market liquidity and order-flow. We estimate separate time-series regressions for SPX call and put options using the following specification:

$$ES_t = \alpha + \beta_1 \log(SD_t) + \beta_2 \log(\text{Volume}_t) + \beta_3 |OI_t| + \text{Time Controls} + \text{Other Controls} + \epsilon_t, \quad (5)$$

where ES_t measures the daily effective spread paid by investors for trading options on day t ,¹³ $\log(SD_t)$ denotes the logarithm of the standard deviation of the intraday order flow distribution on day t , $\log(\text{Volume}_t)$ is the logarithm of the daily volume calculated according to Equation 3, and $|OI_t|$ is the absolute value of the daily order imbalance calculated according to Equation 2.¹⁴ Time controls include day-of-the-week, month-of-year, and year dummies, while other controls include the market return and VIX level on day t ,¹⁵ the absolute value of the average delta, vega and gamma of the options on day t , and one-day and two-day

¹³An alternative measure of trading costs commonly used in the literature is the absolute spread, defined as the spread in dollar terms rather than as a percentage of the mid-price. The robustness section 6.1 presents the results using the absolute spread, which are qualitatively similar to those from the baseline analysis.

¹⁴We use the absolute value of the order imbalance, following the findings of Christoffersen, Goyenko, Jacobs, and Karoui (2018), who demonstrated that this measure is strongly related to illiquidity through a market-maker inventory channel.

¹⁵Qualitatively similar results are obtained when using maturity-specific implied volatility in place of the VIX index. Results are available upon request.

lags of ES_t .¹⁶ We further segment call and put option samples into maturity buckets with one-week intervals, ranging from options expiring on the same day (zero days to maturity or 0DTE), to options expiring in one week (1-6 days), and up to options expiring in seven weeks (42-48 days to maturity). All variables are calculated separately for ATM call and put options in each maturity bucket on day t ¹⁷, and standard errors are calculated using Newey-West with the optimal lag suggested by Andrews and Monahan (1992).

[Table 2 here]

Table 2 presents the summary statistics of the dependent and independent variables included in the regressions. The average spread is very high, ranging from a maximum of 9% for 0DTE options to a minimum of 2% for options with 21-48 days to maturity. Trading volume decreases with maturity, while order imbalance increases with maturity, indicating that the higher trading activity in ultra-short-term options is, on average, less directional compared to longer-term options.

[Table 3 here]

Panels A1 and B1 of Table 3 present the regression results segmented by option maturity buckets. The results consistently reveal a positive and statistically significant relationship between the intraday volatility of order flow $\log(SD_t)$ and the effective cost of trading, indicating that days characterized by greater volatility of intraday order flow correspond to lower liquidity. This result holds across various maturity buckets and put call samples, and remains robust after accounting for numerous controls. The breakdown of results into maturity buckets reveals a trend in the coefficient of $\log(SD_t)$: the coefficient is higher for

¹⁶Table IA.2 of the Online Appendix reports the results using the spread in changes rather than in level and Table IA.7 reports the results using $\log(SD_t)$ scaled by volumes. The findings are qualitatively similar to those from our baseline specification.

¹⁷For 0DTE options we considered the greeks recorded on day $t - 1$.

short-term options and decreases almost monotonically with option maturity. We formally test for differences in coefficients between the ultra-short maturity sample, including 0DTE options, and other maturities, by performing a pooled regression of ES_t on $\log(SD_t)$, with dummies identifying each maturity bucket. Specifically, we introduce seven dummies, D_{1-6} , D_{7-13} , D_{14-20} , D_{21-27} , D_{28-34} , D_{35-41} , and D_{42-48} , representing each maturity bucket except 0DTE. The coefficient of $\log(SD_t)$ measures the sensitivity of illiquidity to volatile order flow in 0DTE options, while interactions of $\log(SD_t)$ with these dummies assess whether the coefficient differs in other maturity buckets compared to the 0DTE bucket. Panels A2 and B2 of Table 3 present the results. The $\log(SD_t)$ coefficient is positive and significant, with a magnitude consistent with the estimate for the 0DTE sample alone. The interaction term coefficients are all negative and significant, confirming the lower sensitivity to order flow volatility in options with longer maturities. To further illustrate this pattern, Panels A3 and B3 display the standardized coefficients of $\log(SD_t)$ from Panels A1 and B1. Standardized coefficients are computed by multiplying the raw coefficients by the standard deviation of $\log(SD_t)$ and dividing by the standard deviation of ES_t , each computed within maturity buckets. The standardized coefficients also decline with maturity, though not strictly monotonically, with the first two buckets (0 and 1–6 days) showing values roughly twice as large as those for the longest maturities (35–41 and 42–48 days).

The coefficients in Table 3 related to the absolute value of order imbalance also offer important insights and connection with the literature. The measure has been utilized in the literature as a measure of demand pressure (Bollen and Whaley 2004; Garleanu, Pedersen, and Potoshman 2008) or as an indicator of changes in option market-maker positions and their associated inventory risk (Muravyev 2016; Christoffersen, Goyenko, Jacobs, and Karoui 2018). It may also serve, to some extent, as a proxy for order flow volatility. While it shows strong significance in univariate regressions, its significance weakens considerably in the full

specifications, remaining significant only in a few subsamples.¹⁸ Importantly, it does not overshadow the significance of order flow volatility. These results suggest that order flow risk is more effectively captured by its intraday volatility, as measured by $\log(SD_t)$, rather than by the aggregate order imbalance, $|OI_t|$.

While the preceding analysis focuses on daily measures of spread, it is well-established that bid-ask spread follows a U-shaped pattern over the trading day, widening at the open and close (see, e.g., Cao, Jacobs, and Ke 2024). This raises the question of whether the previously documented relationship between spreads and order flow volatility is concentrated during specific periods of the trading day. To investigate this, we estimate separate daily regressions using only trades executed within each individual trading hour. The regression specification is as follows:

$$ES_{h,t} = \alpha + \beta_1 \log(SD_{h,t}) + \beta_2 \log(\text{Volume}_{h,t}) + \beta_3 |OI_{h,t}| + \text{Time Controls} + \text{Other Controls} + \epsilon_{h,t}, \quad (6)$$

where $ES_{h,t}$ denotes the effective spread computed using trades from hour h ; $\log(SD_{h,t})$ is the volatility of the order flow, calculated from the 5-minute order imbalances within that hour; $\log(\text{Volume}_{h,t})$ is the logarithm of the trading volume during hour h ; and $|OI_{h,t}|$ is the absolute value of the order imbalance during that hour.

[Table 4 here]

Table 4 reports the regression results by maturity bucket, using broader maturity groupings than in the main specification to ensure sufficient coverage of options with varying maturities within each trading hour. All regressions include the same set of daily fixed ef-

¹⁸Table IA.3 in the Online Appendix presents the regression results of illiquidity on the absolute value of order imbalance, both in the univariate regression and together with $\log(SD)$.

fects as in Table 3, along with control variables computed over hour h .¹⁹ The relationship between $\log(SD_{h,t})$ and $ES_{h,t}$ remains positive and statistically significant across all trading hours and maturity groups, with only one exception (put options with 35–48 days to maturity during the 12 p.m.–1 p.m. interval).

Taken together, the results in this section show that among order-flow variables, order flow volatility stands out as the key driver, exhibiting a robust positive relationship with trading costs across maturities and throughout the trading day.

An empirical challenge in this setting is the potential influence of omitted factors that may simultaneously affect both variables. Although we control for a comprehensive set of variables, it is not possible to fully eliminate all sources of confounding variation in the SPX option setting. To address this concern, Section 4.6 leverages the structure of the U.S. equity options market, where individual stock options are traded across sixteen exchanges with relatively balanced volume. This structure enables us to estimate panel regressions with different combinations of exchange, time, and stock fixed effects, thereby effectively controlling for common shocks at the exchange, time, and stock levels.

4.4 Delta-Hedge Rebalancing Costs and Option Market Liquidity

An important potential determinant of option spreads is the cost of delta-hedging faced by liquidity providers. To manage inventory risk, they might hedge by trading the underlying asset in proportion to the option inventory’s delta (Δ). This strategy involves transaction costs and requires frequent rebalancing, as the delta of the position changes over time. These changes, captured by gamma (Γ), become especially pronounced as options approach

¹⁹Specifically, we control for day-of-the-week, month-of-year, and year fixed effects; one- and two-day lags of the dependent variable; SPY returns over the interval; and the average implied volatility, delta, gamma, and vega of the options during the hour.

expiration. Inventory models in the options market, which assume that liquidity providers provide liquidity while maintaining a discretely rebalanced delta-hedged inventory, suggest that spreads should reflect the associated delta-hedging costs (Jameson and Wilhelm 1992).

In this section, we examine the relationship between two proxies for liquidity providers' delta-hedging costs in the SPX options market and transaction costs. As a first proxy, we compute the delta-hedging cost generated by new order flow absorbed by liquidity providers during each intraday interval d on day t . This is calculated as the absolute value of the sum of order imbalances across all option series (puts and calls with all moneyness and maturities up to seven weeks), weighted by the options' delta and expressed in dollars, as in Dim, Eraker, and Vilkov (2024).²⁰ The total daily delta-hedging cost from new order flow, $DeltaOI_t$, is then obtained by summing across all intraday intervals:

$$DeltaOI_t = \sum_d DeltaOI_{d,t} = \sum_d \left| \sum_j OI_{d,t,j} \Delta_{d,t,j} S_{d,t} \right|$$

The second proxy for delta-hedging costs captures the gamma exposure of liquidity providers' inventory, measured at each intraday interval d . To construct this measure, we first approximate the starting inventory on day t using the market-maker inventory at the end of day $t - 1$, calculated from the CBOE Open-Close database.²¹ For each option series j , we then update this starting value by the cumulative net order flow absorbed by liquidity providers from the beginning of day t up to the end of interval d , yielding the inventory position $Inv_{d,t,j}$. We compute the gamma exposure in dollar terms at each interval d by

²⁰The delta of each option series and the price of the underlying asset are recorded at the end of each 5-minute interval.

²¹The CBOE Open-Close database reports the daily number of buy and sell orders by end-users (non-market-makers) in the SPX options market. The cumulative net order flow from end-users, calculated from options inception, provides a proxy for (minus) the inventory held by market-makers. We merge this data with OptionMetrics. See Jacobs, Mai, and Pederzoli (2024) for details on the filtering and merging procedures.

weighting these inventory positions by the option's gamma and the square of the underlying price. The daily gamma-based hedging cost is obtained by summing across all intraday intervals:

$$\text{GammaInv}_t = \sum_d \text{GammaInv}_{d,t} = \sum_d \left| \sum_j \text{Inv}_{d,t,j} \Gamma_{d,t,j} S_{d,t}^2 \right|.$$

This measure reflects the intensity of delta-hedging adjustments required throughout the day. These costs are higher when gamma is elevated, which typically occurs for near-expiration options.

[Tables 5(a) and 5(b) here]

Tables 5(a) and 5(b) report regressions of ES_t on GammaInv_t (Panel A) and DeltaOI_t (Panel B) for call and put options, respectively. For readability, only results for the first four maturity buckets are shown; results for longer maturities, which are qualitatively similar, are provided in Tables IA.4 and IA.5 of the Online Appendix. In the univariate specifications, GammaInv_t is positively associated with the spread for 0DTE calls and puts, as well as for puts with 7–13 days to maturity. In contrast, DeltaOI_t shows a positive relationship with the spread for all maturities except 0DTE. These findings align with inventory-based option pricing models (Jameson and Wilhelm 1992) and indicate that bid-ask spreads partially reflect liquidity providers' delta-hedging costs, with effects varying across maturities. However, once we control for order flow volatility, most coefficients on GammaInv_t and DeltaOI_t lose significance or even change sign. In the full specification, including $\log(SD)$ and all controls from the baseline analysis, the delta-hedging proxies remain insignificant or negative, while $\log(SD)$ consistently remains positive and significant.

Altogether, these results confirm that the volatility of the order flow is a key determinant of the spread and suggest that liquidity providers actively manage their inventories

throughout the day, with delta-hedging costs playing a secondary role in influencing trading costs.

4.5 Volatile Order Flow and Liquidity in Individual Stock Options

This section analyzes the relationship between trading costs and order-flow in the market for options on individual stocks. We consider the constituents of the S&P 500, tracking them monthly from the beginning of our sample. A stock-day is included in our sample if the stock was part of the S&P 500 index in the preceding month.²² Panel B of Figure 1 displays the average daily volume and order imbalance of at-the-money equity options with up to 48 days to maturity. Unlike SPX options, we find that investors trade more call options than put options on individual stocks, with the difference in volumes significantly increasing from 2020 onwards. Panel B2 indicates that the daily order imbalance is, on average, positive for both call and put options. Our findings are novel but qualitatively align with the summary statistics provided by Bryzgalova, Pavlova, and Sikorskaya (2023) and Bogousslavsky and Muravyev (2024) on retail trading, which accounts for a substantial portion of volume in options on individual stocks in recent years.

We construct order-flow variables separately for each stock s and day t using the same procedure as for SPX outlined in section 4.1. Specifically, daily volume, order imbalance, and effective cost of trading are constructed separately for each stock-day option-type following equations 3, 2, and 4, respectively, and $\log(SD_{s,t})$, is the logarithm of the daily standard deviation of the seventy-eight 5-minute order imbalances.²³

²²Figure IA.1 in the Online Appendix shows the number of individual stocks in our sample over time. In the early years, the count ranges from 100 to 200, eventually stabilizing between 300 and 400 from 2009 onward. The sample size aligns with that used by Christoffersen, Goyenko, Jacobs, and Karoui (2018). As a robustness check, we verified that our findings remain robust even when excluding the early years of the sample. Results are available upon request.

²³Since options on individual stocks may not be traded as frequently as SPX options, we include a stock-

We perform a panel regression of $ES_{s,t}$ on $\log(SD_{s,t})$, $\log(volume_{s,t})$, and $|OI_{s,t}|$ with stock-fixed effect.²⁴ Other controls include the average implied volatility of options on stock s and day t , $IV_{s,t}$, and the average of the options greeks, i.e., gamma, vega, and the absolute value of delta on stock s and day t .²⁵ We also control for stock characteristics, as stock return, firm size and stock volume. Time fixed-effects include day-of-the-week, month-of-the-year, and year controls. Standard errors are double clustered at the day and stock level.

[Table 6 here]

Table 6 presents the results for call options (Panel A) and put options (Panel B), with samples further divided into two maturity groups: options with up to 24 days to expiration and those with 25 to 48 days. The results are robust and align with our earlier findings in SPX options. First, we observe a strong positive relationship between the standard deviation of order flow and illiquidity for both call and put samples. Moreover, the coefficient is larger for very short-maturity options (up to 24 days), confirming that trading costs in shorter-dated contracts are more sensitive to intraday order flow volatility. Finally, the results show that $|OI_{s,t}|$ is positively related to the spread only in the sample of medium-term options (25–48 days to expiration).

4.6 Evidence from Exchange-Level Analysis

In this section, we analyze the relationship between spreads and order flow volatility at the exchange level. We begin by documenting a substantial degree of heterogeneity in volume

day option type (call/put) in the sample if the option group has at least ten non-empty intervals out of the seventy-eight.

²⁴Qualitatively similar results are obtained using a cross-sectional Fama-MacBeth regression instead of the panel regression, and they are presented in Table IA.6 in the Online Appendix.

²⁵For 0DTE options, we use the greeks recorded on day $t - 1$.

absorption across exchanges, with no single exchange dominating the market. The immediate cross-exchange liquidity response following trades also suggests a good level of competition among exchanges. Therefore, this setting allows us to estimate panel regressions of exchange-specific illiquidity on exchange-specific order flow volatility, while controlling for various sets of fixed effects.

Individual stock options trade simultaneously across sixteen exchanges, and the OPRA database reports the exchange identifier for each trade, along with the contemporaneous best bid and offer quotes across all exchanges. For these analyses at the exchange level we focus on the constituents of the Dow Jones which have been part of the index since the start of our sample period, January 2004. Our sample includes the following sixteen tickers: AXP, BA, CAT, DIS, DOW, HD, IBM, INTC, JNJ, JPM, KO, MMM, MRK, MSFT, PG, and WMT. These options have been actively traded throughout the sample period, with trading volumes well distributed across exchanges. Consistent with our main analysis, we examine one-month (up to 48 days to maturity) at-the-money call and put options, where moneyness is based on the delta reported by OptionMetrics at the close of the previous day.

[Figure 4 here]

Panel A of Figure 4 shows the percentage of trading volume in our sample absorbed by each of the sixteen exchanges. The exchange with the highest share is CBOE, which accounts for 19.8% of the volume, followed by ISE (15%), PHLX (13.5%), ARCA (9.69%), NASDAQ (8.86%), AMEX (8.08%), BZX (6.87%), and BOX (4.21%). These are the largest exchanges, and with the exception of BZX, they have operated since 2004. The remaining exchanges also absorb a non-negligible share of trading activity, although less than the major venues listed above. Panel B of Figure 4 reports the percentage of trading volume by exchange for each year in our sample period. It shows that cross-sectional heterogeneity in volume has

increased over time, with the leading exchanges, CBOE and ISE, gradually losing market share to other venues. Overall, the graphs indicate that trading activity in our sample is relatively well distributed, with no single exchange dominating the market.

To investigate the cross-exchange relationship between spreads and order flow, we begin with a trade-by-trade analysis aimed at answering two key questions: (i) Are trades, on average, executed on the exchange quoting the lowest spread? And (ii) Following a trade, how do exchanges revise their quoted spreads depending on whether they absorbed the trade or not?

To this end, we track, for each ticker and trade, the level of the quoted spread and its subsequent change across all exchanges.²⁶ For this analysis only, we exclude trades that sweep an entire layer of the order book, that is, those with a size greater than or equal to the corresponding bid or ask quantity on the executing exchange. All other filters are consistent with those used in the daily baseline analysis. We then estimate the two following pooled regressions specification, separately for each stock:

$$Illiq_{i,j,\tau} = \alpha + \beta Dummy_{i,j,\tau} + \epsilon_{\tau}, \quad (7)$$

where $Illiq_{i,j,\tau}$ is either $Spread_{i,j,\tau}$, the quoted spread in exchange i for option j at time τ , or $\Delta Spread_{i,j,\tau+1}$, the change in the quoted spread in exchange i for option j from the trade time τ to the next trade time $\tau + 1$. We measure the quoted spread as the difference between the quoted ask and bid prices on the exchange, divided by the exchange mid price. $Dummy_{i,j,\tau}$ is a dummy variable which equals one for exchange i where the trade in option j occurred at time τ . It is zero for all other exchanges. The dummy variable thus captures

²⁶For simultaneous trades occurring on the same option and exchange, we consolidate them into a single observation. This aggregated observation has a trade size equal to the signed sum of the individual trade sizes and a trade price that is the average of the individual trade prices.

the differential liquidity between the exchange that absorbed the trade versus the others.

[Table 7 here]

The results are presented by stock and option type in Panel A of Table 7, with all coefficients multiplied by 100.²⁷ We find consistent and robust patterns across both individual stocks and option types (calls and puts). The first specification, which uses the level of the quoted spread as the dependent variable, shows that trades tend to occur on exchanges quoting spreads approximately 6% lower than the average quoted spread across other exchanges. This result is in line with the expectation that investors and brokers use routing strategies to minimize transaction costs. The second specification, which examines the change in spreads following a trade, reveals that quoted spreads tend to decline by ten to twenty basis points on exchanges that did not absorb the trade, while increasing on the exchange where the trade occurred. The net change in the trading exchange's spread is approximately thirty basis points. These findings suggest that exchanges respond immediately to order flow: trading exchanges widen their spreads to incorporate the cost of liquidity provision, while non-trading exchanges tighten theirs, potentially to attract subsequent volume. Overall, the results point to a good degree of competition across exchanges.

We next delve deeper into the relationship between spreads and order flow volatility documented in the main analysis, leveraging this rich cross-sectional structure of trading across exchanges. Our setting, in which the same stock options trade simultaneously across multiple venues, offers a unique opportunity to help isolate the impact of order flow volatility on effective spreads from the influence of other macro-variables that may affect both. This design helps mitigate concerns about omitted variables that may confound the baseline

²⁷All regressions include day fixed effect, and the standard errors are clustered at the day and exchange levels. $Spread_{i,j,\tau}$ and $\Delta Spread_{i,j,\tau+1}$ are also winsorized at the 1% and 99% levels to eliminate instances of apparently unrealistic quotes reported by OPRA.

results. Specifically, we estimate a panel regression of the effective spread for stock s on exchange i and day t , $ES_{i,s,t}$, on $\log(SD_{i,s,t})$, the intraday order flow volatility calculated exclusively from trades absorbed by exchange i . We begin with a specification that includes exchange and day fixed effects, which control for structural differences across venues as well as macro-level variation over time. We then estimate a second specification with stock-by-day fixed effects, which absorb all variation across both stocks and time, thereby accounting for any stock-specific or macroeconomic factors. Panel B of Table 7 displays the results. To further strengthen the analysis, we address the possibility that results are driven by small versus large exchanges by estimating two separate regressions: one using data from all exchanges (specification labeled “All Exchanges”) and another using only the six major venues that have been active throughout the sample period (specification labeled “Six Major Exchanges”). We also control for exchange-specific volume, $\log(Volume_{i,s,t})$, and for exchange-specific order imbalance, $|OI_{i,s,t}|$. The results document that the coefficient on $\log(SD_{i,s,t})$ is positive across all specifications, reinforcing the core finding that greater order flow volatility is associated with higher transaction costs.

While this cross-exchange approach enhances identification, it is not without limitations. In an ideal setting, exchanges would differ only in the order flow they absorb. In practice, however, trading venues may vary in unobservable ways that influence both order flow volatility and liquidity provision. We have addressed these concerns in two ways. First, we have restricted part of the analysis to the six largest and most liquid exchanges, which are more homogeneous in their structure and function. Second, we have included exchange fixed effects in some specifications to control for time-invariant differences across venues. The results remain strong and consistent across specifications, alleviating concerns that they are driven by structural differences across exchanges. We conclude that the cross-sectional exchange design meaningfully enhances identification and supports the robustness of our core

findings.

5 When is the Order Flow More Volatile?

Order flow can stem from two primary sources: informational trading and liquidity-driven activity. Disentangling these channels has long been a central challenge in market microstructure research. In this section, we aim to shed light on the underlying drivers of our key variable, the order flow volatility. While informed trading may generate order flow volatility due to asymmetric information, inventory and liquidity needs can create temporary but substantial intraday imbalances as investors rebalance their positions. To evaluate which mechanism better explains the behavior of order flow volatility, we estimate panel regressions of $\log(SD_{s,t})$ on a set of event-day indicators designed to capture either information-driven activity or inventory-related rebalancing. The results, presented in Table 8, are reported separately for ATM calls and puts across two maturity buckets: short-term (0–24 days) and medium-term (25–48 days).

[Table 8 here]

The table documents that order flow volatility drops significantly on earnings announcement days for all maturities, both on the day of the announcement and slightly on the preceding day. This pattern is hard to reconcile with the interpretation that order flow volatility reflects informed trading, as earnings days are typically associated with heightened informational asymmetry.

In contrast, we observe spikes in order flow volatility on calendar dates more closely linked to inventory rebalancing. Specifically, $\log(SD_{i,s,t})$ increases for medium-maturity options at the end of the month and shows a slight rise at the beginning of the month, consistent with

institutional portfolio rebalancing cycles. For short-term options, order flow volatility is slightly lower on these days, which can be consistent with the idea that end-of-month and beginning-of-month rebalancing activity is concentrated in longer-dated contracts. For short-term options, order flow volatility is higher on the third Friday of each month, coinciding with standard monthly option expirations. In contrast, medium-maturity options exhibit a slightly lower values of $\log(SD_{i,s,t})$ on these Fridays, which may reflect the greater flexibility that investors have in managing longer-dated positions without needing to act precisely on expiration day. Finally, we find that order flow volatility decreases across all maturity buckets when underlying stock volatility is high, supporting the view that it is not driven by market uncertainty either.

Taken together, these results suggest that order flow volatility is more closely tied to inventory risk than to adverse selection risk.

6 Additional Analysis and Robustness

This section presents the results of several robustness checks. Section 6.1 shows that the main findings hold when using the dollar spread instead of the relative spread. Section 6.2 incorporates option return volatility as an additional control. Section 6.3 uses order flow volatility scaled by volume and confirms the robustness of the results in the out-of-the-money options sample. It also provides additional analyses showing that the relationship between volatile order flow and illiquidity is not driven by (i) retail trading, (ii) market opening and closing sessions, and that it remains robust even when time fixed effects are excluded. The corresponding tables for Section 6.3 are provided in the Online Appendix.

6.1 Dollar Spread

The main measure of trading costs used in our analysis is the effective spread, calculated according to Equation 4. This measure expresses the spread in log terms, providing a relative measure of trading costs with respect to the option price. It captures the reduction in option returns that traders incur due to transaction costs. Our choice aligns with the existing literature (Bogousslavsky and Collin-Dufresne 2023; Christoffersen, Goyenko, Jacobs, and Karoui 2018) and reflects our goal of understanding how trading costs faced by traders are influenced by potentially risky patterns in the order flow distribution.

As a robustness check, we test our results using an alternative measure of trading costs: the spread expressed in dollar terms. This measure quantifies the dollar gain a liquidity provider earns by supplying liquidity in a trade and immediately reversing the position with another trade of the opposite sign. Specifically, for each trade i , the dollar spread is defined as:

$$\text{Dollar Spread}_i = |P_i - M_i|$$

where P_i is the trade price and M_i is the prevailing midpoint of the NBBO. On each day, the daily dollar spread is calculated as the volume-weighted average of the dollar spreads across trades, scaled by the value of the underlying asset on day t , S_t .

[Table 9 here]

[Table 10 here]

Panel A of Tables 9 and 10 present the regression results for SPX options and equity options, respectively. In these regressions, we use the dollar spread as the dependent variable instead of our baseline measure of spread, ES_t . All other controls remain consistent with

those used in Tables 3 and 6. The results consistently show a positive and robust relationship between the volatility of the order flow and the dollar spread for both call and put options, as well as for SPX and equity options. Panel A of Table 9 documents that for SPX options, the magnitude of the coefficients is particularly high for ultra-short-term options (0DTE and 1-6 days to maturity), confirming that liquidity is more sensitive to volatile order flow for these very short maturity categories. For equity options, Panel A of Table 10 shows that the difference in coefficients between short and medium maturity options is positive but smaller. Nonetheless, even a small difference in the absolute spread can result in a substantial difference in the relative spread and trading costs for investors (as documented in Table 6), given the lower option prices for shorter maturities options.

6.2 Relation with Realized Option Volatility

All market microstructure models of inventory and asymmetric information (see Foucault, Pagano, and Röell 2013 for a review) predict that transaction costs should be positively related to asset volatility. This section formally tests this hypothesis and assesses the robustness of our results to the inclusion of a variable that measures the realized intraday volatility of options. This addresses the potential concern that order flow volatility might act as a proxy for the underlying volatility of the options themselves.

We compute the realized option volatility (ORV_t) for day t by summing the squared average 5-minute option returns across the seventy-eight 5-minute intervals throughout the day. Panel B of Tables 9 and 10 present the regression results for SPX options and equity options, respectively. The results confirm a robust positive relationship between ORV_t and illiquidity, consistent with theoretical predictions. However, the significance of the volatility of the order flow ($\log(SD)$) is not subsumed by ORV_t , as shown in the second specification

of each subsample. This indicates that, while the two variables are generally correlated (ranging from a minimum of -10% to a maximum of 30%, depending on the sample), they convey distinct information about transaction costs and liquidity.

6.3 Additional Robustness

In this section, we perform a series of additional checks to validate the robustness of our main findings. We begin by scaling both the volatility of order flow and order imbalance by daily volume, resulting in the variables $\log(SD/volume)_t$ and $|OI/volume|_t$. Table IA.7 reports the time-series regression for SPX options using these scaled variables. The findings confirm a positive relationship between the intraday volatility of order flow $\log(SD/volume)_t$ and illiquidity, consistent with our baseline results.

Next, we assess whether the relationship holds when changing the moneyness criterion of the sample. Specifically, we replace at-the-money (ATM) options with out-of-the-money (OTM) options, defined as those with absolute delta values between 0.125 and 0.375 (based on OptionMetrics data at the close of day $t - 1$). Table IA.8 reports the results. Consistent with our baseline analysis, we find a positive and statistically significant relationship between intraday order flow volatility and trading costs. The effect is even more pronounced than in the ATM sample, with coefficients declining as option maturity increases.

We then examine the sensitivity of our results to the exclusion of time fixed effects, such as day-of-week, month-of-year, and year dummies. This check addresses the concern raised by Jennings, Kim, Lee, and Taylor (2024) that controlling for calendar effects may inadvertently bias regression estimates. Table IA.9 report the results obtained excluding day-of-week and month-of-year controls. While the adjusted R^2 values mildly decrease after removing these time controls, the primary results and inferences remain consistent.

Finally, we investigate whether retail trading or specific trading sessions contribute disproportionately to the observed effects. To do so, we re-estimate the main regression from Table 3, first excluding trades identified as retail (using the ‘SLAN’ flag, following Bryzgalova, Pavlova, and Sikorskaya 2023), and then excluding trades executed during the first and last 30 minutes of the trading day. Tables IA.10 and IA.11 report the corresponding results. In both cases, the positive relationship between order flow volatility and illiquidity remains robust, suggesting that the main effect is not driven by retail activity or by market open and close dynamics.

Taken together, these robustness checks reinforce our main conclusion: intraday order flow volatility is a key determinant of trading costs in the options market, and the relationship is stable across a range of specifications, subsamples, and alternative controls.

7 Conclusion

The recent surge in volumes in option contracts with increasingly shorter expirations has raised concerns among academics and regulators about the stability of this expanding market. A key characteristic of the options market is its high level of transaction costs, leaving an open question as to how effectively liquidity providers can further absorb large, potentially imbalanced order flows while maintaining an efficient and well-functioning market.

Our analysis documents economically and statistically significant positive relationship between intraday order flow volatility and illiquidity in options market, particularly for ultra-short term options. The effect is pervasive: it holds in the time-series and cross-sectional dimension, and it outweighs the significance of more traditional daily first-moment measures of order flow dynamics, such as volumes or absolute order imbalances. Furthermore, it also outweighs the significance of traditional measures capturing the delta-hedging needs

of market makers. These findings suggest that liquidity providers rely primarily on active inventory rebalancing and trade matching throughout the day, with the main source of inventory risk arising from providing liquidity to unbalanced order flows. An exchange-level analysis that includes stock-by-day fixed effects further confirms the robustness of this relationship and helps mitigate concerns related to omitted variable bias.

Our findings underscore the potential risks posed by high volumes in short-term option contracts, which can amplify intraday order flow volatility and challenge market stability. We show that as intraday order flow volatility rises, liquidity providers widen bid-ask spreads to manage the elevated risk, resulting in higher hedging costs for investors increasingly dependent on short-term rollover strategies over long-term hedges. This spread widening, in turn, can impair market efficiency by reducing liquidity and price discovery, which may in turn elevate systemic risk. These dynamics highlight critical aspects that regulators should consider to maintain stability and market quality in financial markets. An interesting direction for future research would be to explore the broader implications of unbalanced order flow in the options market, including its impact on investors' portfolios, hedging strategies, and, more generally, on risk premia in financial markets.

References

Adams, Greg, Jean-Sebastien Fontaine, and Chayawat Ornthanalai, 2024, The market for 0-days-to-expiration: The role of liquidity providers in volatility attenuation, *Working Paper*.

Almeida, Caio, Gustavo Freire, and Rodrigo Hizmeri, 2024, 0DTE Asset Pricing, *Working Paper*.

Andrews, Donald WK, and J Christopher Monahan, 1992, An improved heteroskedasticity and autocorrelation consistent covariance matrix estimator, *Econometrica: Journal of the Econometric Society* pp. 953–966.

Bandi, Federico, Nicola Fusari, and Roberto Reno, 2024, 0DTE Option Pricing, *Working Paper*.

Beckmeyer, Heiner, Nicole Branger, and Leander Gayda, 2023, Retail Traders Love 0DTE Options... But Should They?, *Working Paper*.

Bogousslavsky, Vincent, and Pierre Collin-Dufresne, 2023, Liquidity, volume, and order imbalance volatility, *The Journal of Finance* 78, 2189–2232.

Bogousslavsky, Vincent, and Dmitriy Muravyev, 2024, An Anatomy of Retail Option Trading, *Working Paper*.

Bollen, Nicolas PB, and Robert E Whaley, 2004, Does net buying pressure affect the shape of implied volatility functions?, *The Journal of Finance* 59, 711–753.

Brogaard, Jonathan, Jaehee Han, and Peter Y Won, 2023, How does zero-day-to-expiry options Trading affect the volatility of underlying assets?, *Working Paper*.

Brogaard, Jonathan, Terrence Hendershott, and Ryan Riordan, 2014, High-frequency trading and price discovery, *The Review of Financial Studies* 27, 2267–2306.

Bryzgalova, Svetlana, Anna Pavlova, and Taisiya Sikorskaya, 2023, Retail trading in options and the rise of the big three wholesalers, *The Journal of Finance* 78, 3465–3514.

Cao, Jie, Kris Jacobs, and Sai Ke, 2024, Derivative Spreads: Evidence from SPX Options, *Working Paper*.

Chen, Hui, Scott Joslin, and Sophie Xiaoyan Ni, 2019, Demand for crash insurance, intermediary constraints, and risk premia in financial markets, *The Review of Financial Studies* 32, 228–265.

Cho, Young-Hye, and Robert F Engle, 1999, Modeling the impacts of market activity on bid-ask spreads in the option market, *National Bureau of Economic Research, Cambridge, Massachusetts, USA*.

Chordia, Tarun, Jianfeng Hu, Avanidhar Subrahmanyam, and Qing Tong, 2019, Order flow volatility and equity costs of capital, *Management Science* 65, 1520–1551.

Chordia, Tarun, and Avanidhar Subrahmanyam, 2004, Order imbalance and individual stock returns: Theory and evidence, *Journal of Financial Economics* 72, 485–518.

Christoffersen, Peter, Ruslan Goyenko, Kris Jacobs, and Mehdi Karoui, 2018, Illiquidity premia in the equity options market, *The Review of Financial Studies* 31, 811–851.

Dim, Chukwuma, Bjorn Eraker, and Grigory Vilkov, 2024, 0DTEs: Trading, Gamma Risk and Volatility Propagation, *Working Paper*.

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

Ellis, Katrina, Roni Michaely, and Maureen O’Hara, 2000, The accuracy of trade classification rules: Evidence from Nasdaq, *Journal of Financial and Quantitative Analysis* 35, 529–551.

Ernst, Thomas, and Chester S Spatt, 2022, Payment for order flow and option internalization, *Available at SSRN 4056512*.

Foucault, Thierry, Marco Pagano, and Ailsa Röell, 2013, *Market liquidity: theory, evidence, and policy*. (Oxford University Press, USA).

Fournier, Mathieu, and Kris Jacobs, 2020, A tractable framework for option pricing with dynamic market maker inventory and wealth, *Journal of Financial and Quantitative Analysis* 55, 1117–1162.

Garleanu, Nicolae, Lasse Heje Pedersen, and Allen M Poteshman, 2008, Demand-based option pricing, *The Review of Financial Studies* 22, 4259–4299.

George, Thomas J, and Francis A Longstaff, 1993, Bid-ask spreads and trading activity in the S&P 100 index options market, *Journal of Financial and Quantitative Analysis* 28, 381–397.

Glosten, Lawrence R, and Paul R Milgrom, 1985, Bid, ask and transaction prices in a specialist market with heterogeneously informed traders, *Journal of Financial Economics* 14, 71–100.

Goyenko, Ruslan, Chayawat Ornthanalai, and Shengzhe Tang, 2015, Options illiquidity: Determinants and implications for stock returns, *Rotman School of Management Working Paper*.

Grossman, Sanford J, and Merton H Miller, 1988, Liquidity and market structure, *the Journal of Finance* 43, 617–633.

Ho, Thomas SY, and Hans R Stoll, 1983, The dynamics of dealer markets under competition, *The Journal of Finance* 38, 1053–1074.

Hu, Jianfeng, Antonia Kirilova, Dmitriy Muravyev, and Doojin Ryu, 2024, Options Market Makers, *Working Paper*.

Huang, Xing, Philippe Jorion, and Christopher Schwarz, 2024, Some Anonymous Options Trades Are More Equal than Others, *Available at SSRN 4951825*.

Jacobs, Kris, Anh Thu Mai, and Paola Pederzoli, 2024, Identifying Demand Curves in Index Option Markets, *Management Science*.

Jameson, Mel, and William Wilhelm, 1992, Market making in the options markets and the costs of discrete hedge rebalancing, *The Journal of Finance* 47, 765–779.

Jennings, Jared, Jung Min Kim, Joshua Lee, and Daniel Taylor, 2024, Measurement error, fixed effects, and false positives in accounting research, *Review of Accounting Studies* 29, 959–995.

Kelley, Eric K, and Paul C Tetlock, 2013, How wise are crowds? Insights from retail orders and stock returns, *The Journal of Finance* 68, 1229–1265.

Lee, Charles MC, and Mark J Ready, 1991, Inferring trade direction from intraday data, *The Journal of Finance* 46, 733–746.

Muravyev, Dmitriy, 2016, Order flow and expected option returns, *The Journal of Finance* 71, 673–708.

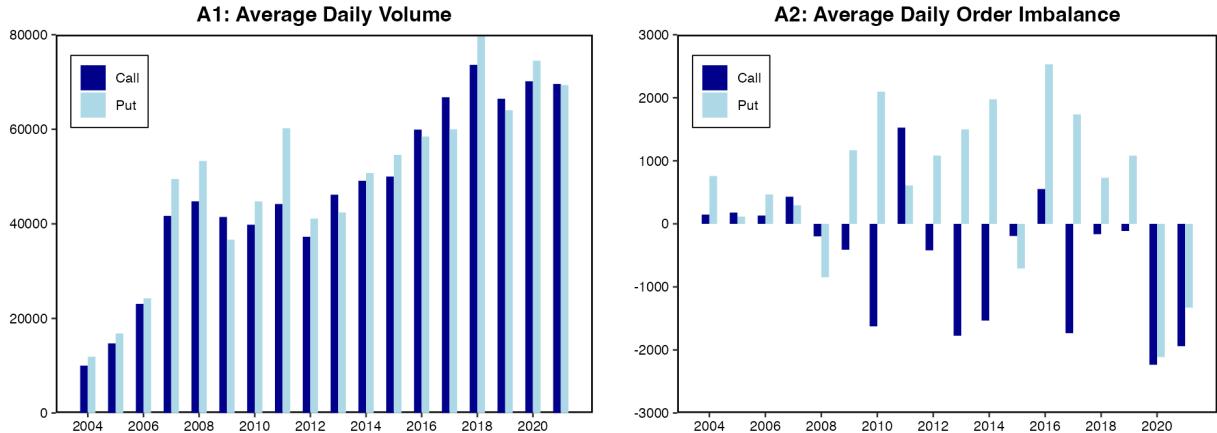
Ni, Sophie X, Neil D Pearson, Allen M Potoshman, and Joshua White, 2021, Does option trading have a pervasive impact on underlying stock prices?, *The Review of Financial Studies* 34, 1952–1986.

Stoikov, Sasha, and Mehmet Sağlam, 2009, Option market making under inventory risk, *Review of Derivatives Research* 12, 55–79.

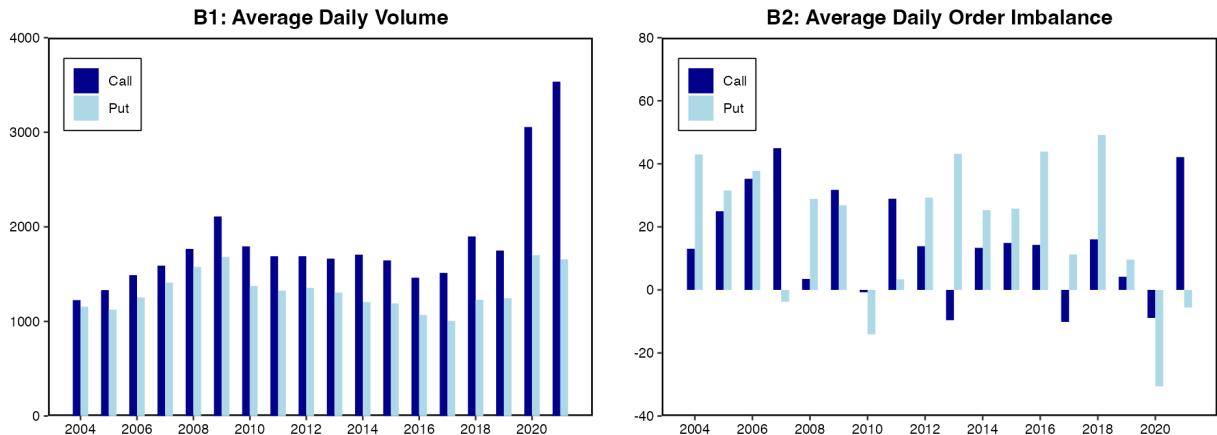
Stoll, Hans R, 1978, The supply of dealer services in securities markets, *The Journal of Finance* 33, 1133–1151.

Figure 1: Daily Volumes and Order Imbalances

Panel A: SPX options

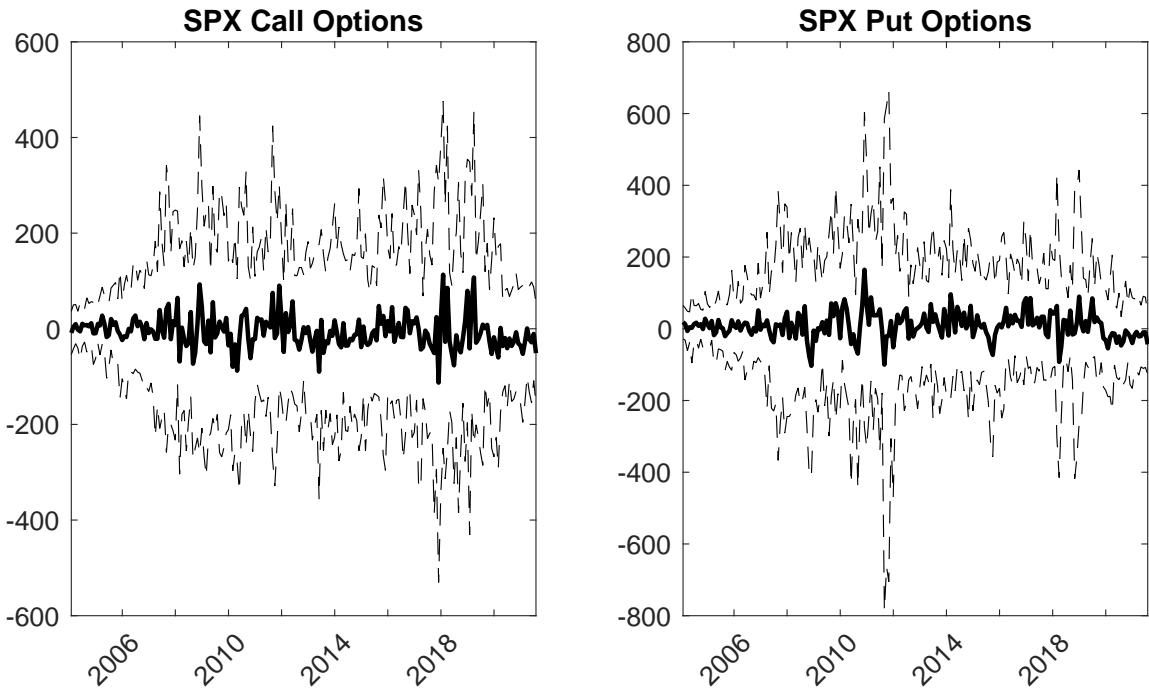


Panel B: Individual Stock Options



This figure displays the average daily volume and order imbalance for at-the-money (ATM) options with maturities up to one month (48 days), across each year in our sample period. Daily volume is the total number of contracts traded, and daily order imbalance is the difference between buy and sell initiated trades. Panel A displays the average daily volume (A1) and order imbalance (A2) for SPX call and put options. Panel B plots for call and put options written on the stocks which are part of the S&P500 index, where we compute average daily volume and order imbalance for each stock-year, and then we take the cross-sectional averages for each year.

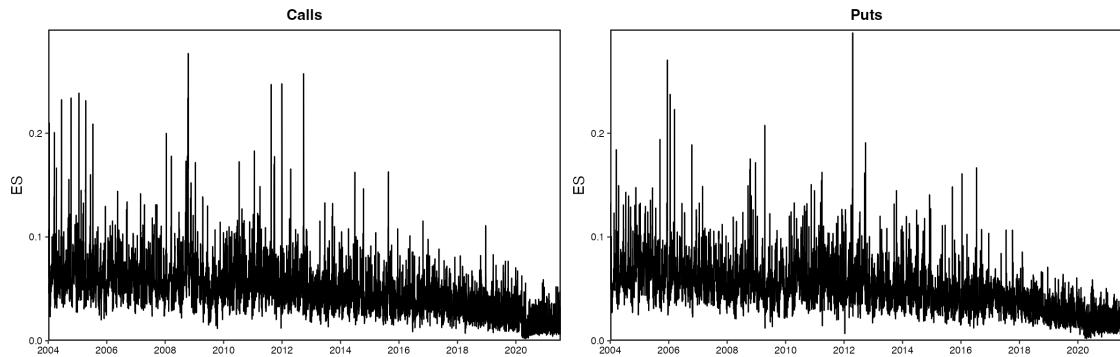
Figure 2: Intraday Order Flow Distribution Over the Years



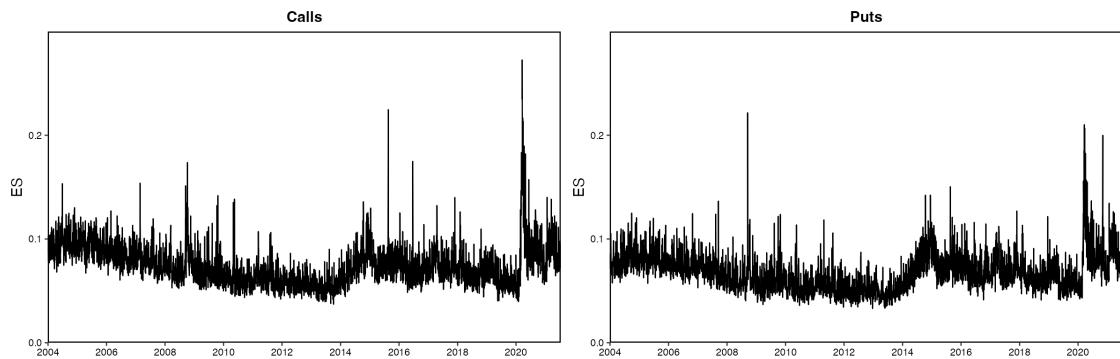
This figure displays the time-series of the average intraday 5-minute order flow for SPX ATM call and put options with confidence intervals. The graph is obtained by dividing each trading day into seventy-eight equal intervals, each covering five minutes, and calculating the order flow (buys minus sells) of put and call options within each interval. The solid lines display the daily average of these 5-minute order flows, μ_t , while the dotted lines depict the 95% confidence intervals, calculated as $\mu_t \pm \frac{Z\sigma_t}{\sqrt{n}}$, where σ_t is the intraday standard deviation of the seventy-eight order flows. For readability, the graph displays the monthly averages of these daily quantities.

Figure 3: Time-Series of ES_t

Panel A: SPX Options



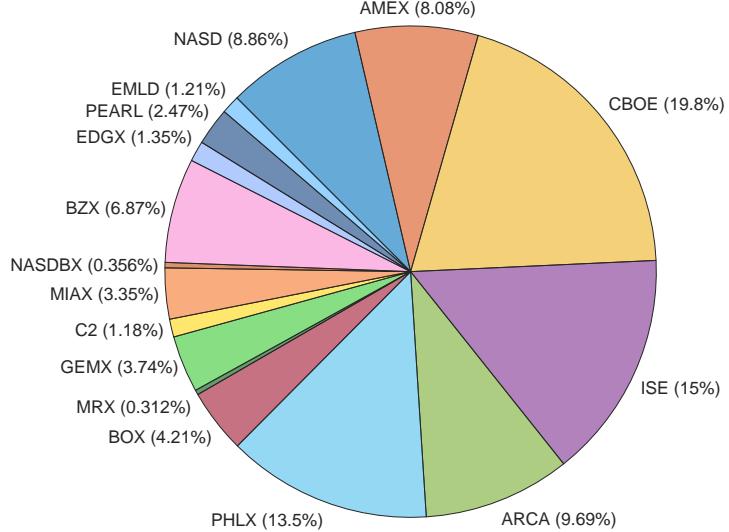
Panel B: Individual Stock Options



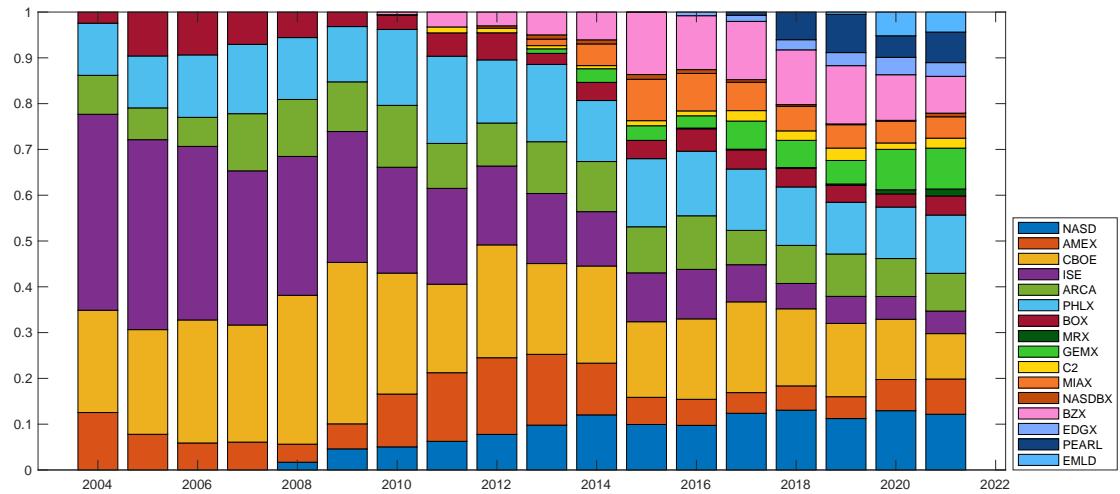
The figure presents the time-series of the daily effective spread for ATM call and put options. Panel A presents the graph for SPX options while Panel B presents the graphs for individual stock options, where a stock-day is included in our sample if the stock was part of the S&P 500 index in the preceding month.

Figure 4: Percentage of Volume Absorbed by Exchanges

Panel A: Full Sample



Panel B: Annual Time-Series



Panel A reports the percentage of at-the-money call and put option volume (with maturities up to seven weeks) absorbed by each exchange for the sixteen tickers analyzed in Section 4.6 over the 2004–2021 sample period. Panel B shows the yearly breakdown of these exchange-level volume shares.

Table 1: Intraday Order Flow Distribution Over the Years

Panel A: SPX Calls												
Year	A1: Five-minute Order Flow Summary Statistics						A2: Difference in Distribution Between Low and High ES Days					
	Mean	Std	Skewness	Q25	Q50	Q75	ΔMean	ΔStd	ΔSkew	ΔQ25	ΔQ50	ΔQ75
2004	2.264	229.677	-0.175	-31.429	2.998	41.916	0.068	-161.011***	1.263*	27.712***	-3.904***	-47.394***
2005	3.001	350.244	0.166	-48.560	-0.181	48.458	5.256	-192.330***	1.471*	40.721***	-0.731*	-44.673***
2006	1.217	510.297	-0.172	-68.341	1.978	82.333	-13.141	-194.946**	-1.176	64.615***	-8.481	-79.904***
2007	5.031	917.111	-0.174	-129.776	3.199	148.560	-11.598	-576.354***	-0.546	181.260***	-9.250	-209.769***
2008	-3.349	969.571	-0.479	-127.992	3.079	153.907	-4.778	-839.847***	0.334	193.269***	-11.250	-237.298***
2009	-6.875	1047.319	-0.086	-113.103	1.246	111.810	-65.604	-688.481***	0.230	110.087***	-3.673	-140.462***
2010	-21.329	971.170	-0.593	-119.851	1.587	125.033	-98.646*	-433.039**	-0.022	41.356	-8.288	-164.808***
2011	18.689	922.002	-0.193	-127.552	6.863	151.914	-46.192	-504.000***	0.721	131.096***	-19.404*	-232.990***
2012	-6.925	723.389	-0.190	-121.593	3.850	124.954	-20.333	-319.204***	0.924	54.680**	-12.800*	-98.560**
2013	-23.995	875.890	-0.070	-152.160	-6.550	121.287	-1.147	-235.637	-1.075	111.615***	0.654	-85.404***
2014	-19.301	846.771	-0.209	-152.172	-4.938	135.027	-51.440	-331.513***	-0.988	111.413***	11.077	-119.740***
2015	-3.070	792.471	-0.312	-133.594	3.210	140.449	-42.433	-295.734**	-0.667	78.231***	-14.788***	-150.106***
2016	6.924	852.823	-0.080	-164.524	2.062	179.867	-83.110**	-213.614	-0.300	97.106***	-18.154*	-238.606***
2017	-22.156	1157.599	-0.274	-179.111	-4.414	160.257	86.379**	-432.589**	1.334	142.875***	4.000	-110.856***
2018	-1.567	1121.613	-0.042	-196.344	-1.600	189.645	-84.126*	-247.488	-1.019	170.817***	7.462	-212.990***
2019	-2.056	1020.217	0.064	-160.959	-2.187	148.254	10.189	-431.078	0.253	98.981***	-1.269	-74.798***
2020	-28.763	597.248	-0.354	-155.767	-12.866	124.071	92.753***	-407.609***	0.411	199.413***	64.904***	-48.010

Panel B: SPX Puts												
Year	B1: Five-minute Order Flow Summary Statistics						B2: Difference in Distribution Between Low and High ES Days					
	Mean	Std	Skewness	Q25	Q50	Q75	ΔMean	ΔStd	ΔSkew	ΔQ25	ΔQ50	ΔQ75
2004	10.890	267.683	0.394	-34.179	3.432	50.474	-16.995	-212.596***	0.045	37.644***	-2.577	-58.538***
2005	0.827	385.758	-0.082	-54.116	2.437	62.690	-10.693	-176.646**	-0.718	48.769***	0.231	-47.260***
2006	6.737	486.335	0.097	-76.521	1.319	87.596	20.745	-379.076***	0.687	136.962***	-0.212	-119.625***
2007	3.936	1011.164	0.007	-163.017	2.255	178.691	28.269	-626.536**	0.198	221.077***	0.000	-176.558***
2008	-12.104	1070.962	-0.249	-186.361	-1.318	180.029	31.167	-812.081***	0.191	240.269***	19.115	-193.548***
2009	15.571	980.570	0.074	-90.806	4.847	111.010	28.412	216.043	0.312	94.702***	-4.827	-104.769***
2010	28.771	1279.395	0.445	-111.428	9.440	143.686	-25.708	-882.587***	-0.426	134.298***	-24.058	-186.240***
2011	6.218	1764.674	-0.096	-143.393	10.411	183.536	-18.898	621.615	-0.691	52.683	-13.577	-99.519
2012	13.453	864.470	0.031	-122.614	3.780	140.172	-25.652	-48.268	0.434	63.590*	-4.460	-85.960**
2013	19.639	811.893	0.216	-118.472	7.240	145.224	-13.845	-411.241***	-0.851	64.644**	-6.885	-106.548***
2014	25.477	858.392	0.278	-131.870	11.296	179.445	-64.608	-226.117	-0.796	67.269***	-12.635	-117.029***
2015	-8.600	803.362	-0.187	-153.960	4.496	161.197	-35.684	-277.738*	-0.745	113.058***	-8.808	-149.885***
2016	32.651	756.723	-0.080	-145.812	14.167	214.905	-79.378***	-370.184***	-0.029	110.067***	-18.500***	-266.760***
2017	22.620	678.144	-0.074	-155.832	8.472	194.332	-39.823	-276.824***	-0.405	63.298*	-13.058	-130.875***
2018	9.680	1027.553	0.087	-197.475	7.886	232.516	16.172	-606.494*	0.962	120.971**	-17.250	-218.548**
2019	14.007	662.457	0.079	-158.144	2.615	170.965	-20.837	-291.757***	-0.975	95.875***	-2.846	-105.404***
2020	-27.125	563.053	-0.103	-180.096	-22.883	129.011	70.447***	-290.051***	-0.867	171.837***	66.885***	0.106

This table displays averages of intraday order flow distribution statistics for SPX ATM call (Panel A) and put options (Panel B). We divide each trading day into seventy-eight equal intervals, each covering five minutes, and we compute the order flow (buy minus sell orders) within each interval. Panels A1 and B1 display the daily mean, standard deviation (*Std*), skewness, first quartile (*Q25*), median (*Q50*), and third quartile (*Q75*) of the five-minute order flow distribution. Panels A2 and B2 display differences in these average statistics (mean, std, skewness, and quantiles) between high and low liquidity days, classified annually into low liquidity (top 10%) and high liquidity (bottom 10%) days based on *ES* values. Significance levels are denoted by *, **, and ***, representing the 10%, 5%, and 1% levels, respectively.

Table 2: Descriptive Statistics of ES, $\log(\text{SD})$, Volume and Order Imbalance. SPX Options

Panel A: Calls																
DTM	ES								log(volume)							
	0	1-6	7-13	14-20	21-27	28-34	35-41	42-48	0	1-6	7-13	14-20	21-27	28-34	35-41	42-48
Mean	0.087	0.062	0.035	0.028	0.025	0.023	0.021	0.020	9.067	8.623	8.771	8.404	8.206	8.364	7.860	7.620
Std	0.067	0.082	0.031	0.024	0.021	0.020	0.019	0.019	1.144	1.589	1.026	1.255	1.376	1.579	1.921	1.976
Skewness	3.184	14.963	1.894	1.569	1.324	1.635	1.476	1.744	-0.717	-1.175	-0.677	-0.327	-0.146	-0.151	-0.064	-0.135
Kurtosis	26.060	459.054	6.569	3.022	1.676	5.170	2.907	5.931	0.821	1.028	1.614	0.305	-0.351	-0.542	-1.077	-1.093
ρ	0.285	0.257	0.601	0.633	0.671	0.609	0.685	0.651	0.211	0.561	0.424	0.535	0.558	0.585	0.643	0.604
N	1040	2992	2901	2799	2737	2598	2364	2167	1038	2990	2899	2797	2735	2596	2362	2165
log(SD)									OI							
DTM	0	1-6	7-13	14-20	21-27	28-34	35-41	42-48	0	1-6	7-13	14-20	21-27	28-34	35-41	42-48
Mean	4.786	4.895	5.183	5.047	4.996	5.219	4.863	4.703	1.521	1.870	2.239	2.237	2.137	2.929	2.883	2.372
Std	0.855	1.210	0.999	1.143	1.198	1.292	1.630	1.708	2.013	2.947	3.644	3.817	3.798	5.183	5.372	4.352
Skewness	-0.425	-0.599	-0.175	-0.189	-0.182	-0.335	-0.224	-0.316	3.016	3.969	5.246	4.860	4.221	3.655	3.429	3.476
Kurtosis	0.508	0.470	0.385	-0.108	-0.319	-0.121	-0.717	-0.743	13.679	24.015	46.830	41.300	24.836	17.151	15.824	16.331
ρ	0.064	0.408	0.315	0.414	0.357	0.414	0.501	0.470	0.026	0.166	0.181	0.231	0.200	0.258	0.272	0.250
N	1038	2990	2899	2797	2735	2596	2362	2165	1038	2990	2899	2797	2735	2596	2362	2165
Panel B: Puts									log(volume)							
DTM	0	1-6	7-13	14-20	21-27	28-34	35-41	42-48	0	1-6	7-13	14-20	21-27	28-34	35-41	42-48
Mean	0.095	0.063	0.036	0.029	0.025	0.023	0.021	0.019	9.227	8.687	8.782	8.407	8.185	8.269	7.883	7.694
Std	0.060	0.061	0.029	0.024	0.021	0.019	0.019	0.020	1.121	1.539	1.078	1.283	1.398	1.574	1.861	1.920
Skewness	1.970	2.666	1.472	1.577	1.482	1.380	1.491	1.445	-0.818	-1.105	-0.639	-0.202	0.057	0.019	-0.043	-0.042
Kurtosis	8.147	11.526	2.595	4.584	2.982	2.180	3.571	1.880	1.028	1.000	1.622	0.167	-0.447	-0.565	-0.944	-0.976
ρ	0.348	0.439	0.625	0.657	0.657	0.693	0.683	0.712	0.321	0.565	0.415	0.546	0.612	0.656	0.668	0.634
N	1040	2962	2905	2793	2740	2620	2376	2218	1038	2960	2903	2791	2738	2618	2374	2216
log(SD)									OI							
DTM	0	1-6	7-13	14-20	21-27	28-34	35-41	42-48	0	1-6	7-13	14-20	21-27	28-34	35-41	42-48
Mean	4.880	4.941	5.186	5.063	4.965	5.109	4.819	4.665	1.776	1.906	2.224	2.196	2.201	2.772	2.472	2.426
Std	0.843	1.165	1.003	1.124	1.210	1.299	1.553	1.675	2.172	3.152	3.842	4.666	4.589	6.001	5.200	5.795
Skewness	-0.728	-0.605	-0.199	-0.123	-0.117	-0.152	-0.308	-0.192	2.336	5.770	5.949	7.604	5.903	5.878	5.352	6.209
Kurtosis	1.228	0.604	0.759	0.129	-0.266	-0.075	-0.577	-0.725	7.352	67.028	53.343	85.581	52.169	50.961	42.339	56.230
ρ	0.118	0.374	0.329	0.422	0.448	0.500	0.554	0.542	0.157	0.162	0.209	0.339	0.243	0.334	0.373	0.257
N	1038	2960	2903	2791	2738	2618	2374	2216	1038	2960	2903	2791	2738	2618	2374	2216

The table reports the time-series mean, standard deviation, skewness, excess kurtosis, AR(1) coefficient (ρ), and total number of observations (N) of the effective spread (ES), logarithm of daily volume ($\log(\text{volume})$), logarithm of volatility of order-flow ($\log(\text{SD})$), and absolute value of daily order flow ($|\text{OI}|$) across option maturity buckets. Panel A presents the results for SPX call options while Panel B presents the results for SPX put options. Absolute value of daily order-flow is divided by 1000.

Table 3: Time-series Regressions of ES_t on $\log(SD_t)$ for SPX Options

DTM	A1: Subsample Regressions					A2: Pooled Regression with Maturity Dummies			A3: Standardized Coeff	
	0	1-6	7-13	14-20	21-27	28-34	35-41	42-48	DTM	
$\log(SD_t)$	0.020*** (4.098)	0.020*** (7.302)	0.007*** (7.018)	0.003*** (3.807)	0.004*** (5.484)	0.003*** (6.694)	0.001*** (3.084)	0.001*** (1.336)	$\log(SD_t)$	0.025*** 0
$\log(\text{volume}_t)$	0.001 (0.251)	-0.013*** (-4.894)	0.002** (2.290)	0.003*** (4.954)	0.002*** (4.125)	-0.001 (-0.442)	0.002*** (4.395)	0.002*** (5.644)	$\log(SD_t)$ D_{1-6}	-0.018*** 1-6
$ OI_t $	-0.016** (-2.067)	0.002 (0.511)	0.001 (0.812)	0.001 (1.144)	0.002 (1.465)	0.001 (1.735)	0.002*** (3.490)	0.004*** (4.908)	$\log(SD_t)$ D_{7-13}	-0.016*** 7-13
R_{Mt}	-1.920*** (-3.753)	-0.007 (-0.039)	-0.030 (-0.301)	-0.013 (-0.285)	-0.002 (-0.055)	-0.084 (-1.148)	0.022 (0.749)	-0.031 (-0.946)	$\log(SD_t)$ D_{14-20}	-0.018*** 14-20
VIX_t	-0.028 (-0.911)	0.039 (1.621)	-0.011 (-0.858)	-0.005 (-0.695)	-0.005 (-0.051)	-0.001 (-0.994)	0.021** (2.157)	0.046*** (2.157)	$\log(SD_t)$ D_{21-27}	-0.020*** 21-27
Time Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Maturity Dummies	Yes
Other Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Time Controls	Yes
N	1035	2990	2899	2797	2735	2596	2362	2165	Other Controls	Yes
Adj. R ²	0.484	0.304	0.640	0.664	0.681	0.629	0.703	0.675	N	19579
									Adj. R ²	0.451
Panel A: Calls										
DTM	B1: Subsample Regressions					B2: Pooled Regression with Maturity Dummies			B3: Standardized Coeff	
	0	1-6	7-13	14-20	21-27	28-34	35-41	42-48	DTM	
$\log(SD_t)$	0.023*** (6.192)	0.019*** (7.004)	0.008*** (7.862)	0.004*** (5.164)	0.003*** (4.339)	0.002*** (3.863)	0.002*** (3.638)	0.002*** (4.292)	$\log(SD_t)$	0.021*** 0
$\log(\text{volume}_t)$	-0.010*** (-2.588)	-0.012*** (-4.686)	0.001 (0.466)	0.003*** (3.926)	0.003*** (4.647)	0.002*** (4.970)	0.002*** (3.909)	0.002*** (3.514)	$\log(SD_t)$ D_{1-6}	-0.012*** 1-6
$ OI_t $	-0.004 (-0.532)	0.008** (2.540)	0.002 (0.980)	0.000 (0.258)	0.004** (2.360)	0.001 (0.925)	-0.000 (-0.241)	0.001 (1.550)	$\log(SD_t)$ D_{7-13}	-0.011*** 7-13
R_{Mt}	1.439*** (3.670)	-0.070 (-0.615)	0.083 (1.077)	0.015 (0.400)	-0.034 (-1.035)	0.058 (-0.773)	0.120 (1.327)	-0.029 (-0.717)	$\log(SD_t)$ D_{14-20}	-0.013*** 14-20
VIX_t	-0.054 (-1.641)	-0.015 (-0.740)	-0.015 (-1.237)	-0.010 (-1.478)	0.014 (1.443)	0.015 (1.492)	0.032** (2.551)	0.017* (1.781)	$\log(SD_t)$ D_{21-27}	-0.014*** 21-27
Time Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Maturity Dummies	Yes
Other Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Time Controls	Yes
N	1035	2960	2903	2791	2738	2618	2374	2216	Other Controls	Yes
Adj. R ²	0.513	0.528	0.633	0.666	0.674	0.685	0.685	0.727	N	19635
									Adj. R ²	0.611
Panel B: Puts										

Panels A1 and B1 present the time series regressions of ES_t on $\log(SD_t)$ for SPX call (A1) and put options (B1), performed separately for different maturity buckets. ES_t is the effective spread on day t , $\log(SD_t)$ is the logarithm of the standard deviation of the intraday order flow distribution on day t , $\log(\text{volume}_t)$ is the logarithm of the daily options volume, and $|OI_t|$ is the absolute value of the daily order imbalance divided by 10,000. R_{Mt} is daily return on SPX on day t and VIX_t is the level of VIX divided by 100 on day t . Panels A2 and B2 present the results of pooled regressions of ES_t on $\log(SD_t)$ and the interaction of $\log(SD_t)$ with dummies that identify the different maturity buckets. Time controls include day-of-the-week, month-of-year, and year dummies. Other controls contain one-day and two-day lags of ES_t , absolute value of the average delta, vega, and gamma of the options on day t . Standard errors are computed using Newey-West with the optimal lag suggested by Andrews and Monahan (1992), and corresponding t-statistics are presented in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level. Panels A3 and B3 present the coefficients of $\log(SD_t)$ of Panels A1 and B1 in standardized units.

Table 4: Time-Series Regressions of $ES_{h,t}$ by Intraday Hour

Maturity		Panel A: SPX Calls Options								
		0-6			7-34			35-48		
Interval		$\log(SD_{h,t})$	$\log(volume_{h,t})$	$ OI_{h,t} $	$\log(SD_{h,t})$	$\log(volume_{h,t})$	$ OI_{h,t} $	$\log(SD_{h,t})$	$\log(volume_{h,t})$	$ OI_{h,t} $
9:30-10:00		0.011*** (4.669)	-0.005* (-1.934)	0.059*** (3.408)	0.006*** (6.654)	0.001 (0.368)	0.003 (0.527)	0.002*** (3.441)	0.001 (1.446)	0.011** (2.383)
10:00-11:00		0.012*** (4.493)	-0.003 (-1.084)	0.026** (2.065)	0.003*** (3.418)	0.003*** (3.247)	0.003 (1.233)	0.002*** (2.937)	0.002*** (3.591)	0.004*** (2.722)
11:00-12:00		0.015*** (7.841)	-0.005*** (-2.687)	0.002 (0.210)	0.003*** (3.394)	0.003*** (3.366)	-0.002 (-1.197)	0.001* (1.883)	0.002*** (4.191)	0.006*** (4.151)
12:00-13:00		0.020*** (3.108)	-0.012 (-1.516)	0.005 (0.296)	0.005** (2.430)	0.005*** (3.651)	0.003 (1.067)	0.001 (1.478)	0.002*** (4.675)	0.004* (1.668)
13:00-14:00		0.008*** (4.447)	-0.003 (-1.229)	0.044** (2.454)	0.002** (2.218)	0.003*** (4.123)	0.001 (0.176)	0.001*** (2.822)	0.001*** (2.935)	0.004 (1.555)
14:00-15:00		0.008*** (3.583)	-0.002 (-0.933)	0.057*** (2.620)	0.003*** (3.061)	0.003*** (2.970)	-0.001 (-0.134)	0.001* (1.681)	0.002*** (4.300)	0.002*** (1.003)
15:00-16:00		0.005** (2.029)	0.006** (2.490)	-0.001 (-0.080)	0.005*** (4.526)	0.003*** (3.245)	-0.007*** (-2.826)	0.001** (2.573)	0.002*** (4.465)	0.003 (1.151)
Maturity		0-6			7-34			35-48		
		$\log(SD_{h,t})$	$\log(volume_{h,t})$	$ OI_{h,t} $	$\log(SD_{h,t})$	$\log(volume_{h,t})$	$ OI_{h,t} $	$\log(SD_{h,t})$	$\log(volume_{h,t})$	$ OI_{h,t} $
Interval										
9:30-10:00		0.013*** (7.913)	-0.006*** (-3.514)	0.066*** (4.916)	0.006*** (7.294)	-0.001 (-0.935)	0.005 (1.404)	0.004*** (3.525)	-0.001 (-0.442)	0.003 (1.213)
10:00-11:00		0.014*** (6.202)	-0.004** (-2.108)	0.050*** (4.130)	0.004*** (4.372)	0.003*** (2.881)	0.003 (1.150)	0.003*** (4.435)	0.001** (2.259)	0.005* (1.712)
11:00-12:00		0.009*** (4.334)	-0.002 (-0.883)	0.050*** (3.800)	0.005*** (5.003)	0.001 (0.844)	0.001 (0.591)	0.001*** (2.932)	0.002*** (5.296)	-0.000 (-0.007)
12:00-13:00		0.007*** (3.569)	-0.000 (-0.113)	0.011 (1.144)	0.004*** (4.859)	0.002** (1.998)	0.001 (0.776)	0.001*** (3.195)	0.002*** (3.659)	0.003 (1.015)
13:00-14:00		0.008*** (3.906)	0.002 (1.078)	0.006 (0.601)	0.003*** (3.405)	0.002*** (2.802)	-0.000 (-0.057)	0.002*** (4.800)	0.001** (2.487)	0.001 (0.230)
14:00-15:00		0.011*** (4.361)	-0.001 (-0.491)	0.027* (1.836)	0.002** (2.241)	0.0037*** (4.101)	0.001 (0.346)	0.001*** (2.676)	0.002*** (4.350)	0.004 (1.187)
15:00-16:00		0.006* (1.771)	0.004 (1.416)	0.032* (1.809)	0.002** (2.140)	0.003*** (3.437)	-0.001 (-0.420)	0.002*** (3.230)	0.002*** (3.557)	0.007** (1.972)

The table reports daily time-series regressions of $ES_{h,t}$ on $\log(SD_{h,t})$, $\log(volume_{h,t})$, and $|OI_{h,t}|$, where all variables are computed for specific intraday hour indicated in the first column. The dependent variable, $ES_{h,t}$, is the effective spread calculated using only trades executed during hour h . The key independent variable, $\log(SD_{h,t})$, measures the volatility of order flow using 5-minute order imbalances within hour h . The controls include $\log(volume_{h,t})$, the logarithm of trading volume during that hour, and $|OI_{h,t}|$, the absolute value of the order imbalance, scaled by 10,000. All regressions control for day-of-the-week, month-of-year, and year fixed effects, as well as the same set of additional controls used in Table 3, computed on a daily basis over hour h . Standard errors are calculated using the Newey-West estimator with the optimal lag length selected according to Andrews and Monahan (1992). t-statistics are reported in parentheses. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively.

Table 5(a): Regressions of ES_t on Liquidity Providers Delta-Hedging Variables for Call Options

Maturity	0		1-6		7-13		14-20	
	Panel A: Gamma Inventory	Panel B: Inventory Change	Panel A: Gamma Inventory	Panel B: Inventory Change	Panel A: Gamma Inventory	Panel B: Inventory Change	Panel A: Gamma Inventory	Panel B: Inventory Change
$log(SD)_t$	0.026*** (12.390)	0.019*** (3.880)	0.010*** (8.805)	0.018*** (6.412)	0.010*** (20.306)	0.006*** (6.524)	0.007*** (24.020)	0.003*** (3.655)
$GammaInv_t$	0.522** (2.240)	-0.036 (-0.157)	0.040 (0.187)	-0.279 (-1.557)	-0.345** (-2.020)	-0.207 (-1.324)	0.086 (1.424)	0.057 (1.023)
Time Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Other Controls	No	No	Yes	No	No	No	No	No
N	1035	1035	1035	2990	2990	2899	2899	2797
Adj. R ²	0.256	0.369	0.472	0.233	0.253	0.283	0.487	0.533
				0.253	0.283	0.487	0.589	0.629
				0.472	0.523	0.623	0.660	0.660

The table reports time-series regressions of ES_t on the liquidity providers' inventory variables $GammaInv_t$ (Panel A) and $DeltaOI_t$ (Panel B) for at-the-money S&P 500 call options. Regressions are estimated separately across maturity buckets. The dependent variable, ES_t , is the daily effective spread, while $GammaInv_t$ and $DeltaOI_t$ are the two proxies introduced in Section 4.4 to capture the delta-hedging costs incurred by liquidity providers on day t . Specifically, $GammaInv_t$ reflects the cost of rebalancing the delta-hedge of pre-existing positions, and $DeltaOI_t$ captures the cost associated with hedging new positions absorbed during the day. The remaining control variables are analogous to those used in the baseline regressions in Table 3. Standard errors are computed using the Newey-West estimator with the optimal lag length selected according to Andrews and Monahan (1992). t-statistics are reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 5(b): Regressions of ES_t on Liquidity Providers Delta-Hedging Variables for Put Options

Maturity		0		1-6		7-13		14-20	
Panel A: Gamma Inventory									
$log(SD)_t$		0.021*** (12.755)	0.024*** (6.354)	0.010*** (9.869)	0.018*** (6.608)	0.009*** (18.420)	0.008*** (7.781)	0.007*** (19.296)	0.004*** (5.404)
$GammaInv_t$	0.525* (1.702)	0.368 (1.196)	0.197 (0.598)	0.010 (0.065)	-0.033 (-0.219)	0.125* (1.765)	0.131* (1.745)	0.015 (0.236)	0.035 (0.625)
Time Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Other Controls	No	No	Yes	No	No	No	No	No	Yes
N	1035	1035	1035	2960	2960	2903	2903	2791	2791
Adj. R ²	0.357	0.435	0.511	0.433	0.472	0.499	0.511	0.605	0.634
Panel B: Inventory Change									
Maturity		0		1-6		7-13		14-20	
$log(SD)_t$		0.021*** (12.779)	0.024*** (6.374)	0.010*** (9.118)	0.018*** (6.588)	0.009*** (17.720)	0.008*** (7.786)	0.008*** (18.940)	0.004*** (5.497)
$DeltaOI_t$	-0.013 (-0.428)	-0.044 (-1.446)	-0.052* (-1.931)	0.115*** (6.618)	0.040** (2.338)	0.054*** (3.322)	0.062*** (4.966)	0.003 (0.279)	0.030*** (0.149)
Time Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Other Controls	No	No	Yes	No	No	No	No	No	Yes
N	1035	1035	1035	2960	2960	2903	2903	2791	2791
Adj. R ²	0.355	0.435	0.512	0.440	0.473	0.500	0.518	0.604	0.634

The table reports time-series regressions of ES_t on the liquidity providers' inventory variables $GammaInv_t$ (Panel A) and $DeltaOI_t$ (Panel B) for at-the-money S&P put options. Regressions are estimated separately across maturity buckets. The dependent variable, ES_t , is the daily effective spread, while $GammaInv_t$ and $DeltaOI_t$ are the two proxies introduced in Section 4.4 to capture the delta-hedging costs incurred by liquidity providers on day t . Specifically, $GammaInv_t$ reflects the cost of rebalancing the delta-hedge of pre-existing positions, and $DeltaOI_t$ captures the cost associated with hedging new positions absorbed during the day. The remaining control variables are analogous to those used in the baseline regressions in Table 3. Standard errors are computed using the Newey-West estimator with the optimal lag length selected according to Andrews and Monahan (1992). t-statistics are reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 6: Panel Regressions of $ES_{s,t}$ on $\log(SD_{s,t})$ for Individual Stock Options

Maturity	Panel A: Calls		Panel B: Puts	
	0-24	25-48	0-24	25-48
$\log(SD_{s,t})$	0.015*** (29.842)	0.007*** (28.965)	0.014*** (24.001)	0.006*** (23.930)
$\log(volume_{s,t})$	-0.015*** (-24.181)	-0.007*** (-26.652)	-0.013*** (-18.412)	-0.006*** (-21.834)
$ OI_{s,t} $	-0.001 (-0.67)	0.003** (2.014)	-0.001 (-0.349)	0.002 (1.613)
$Return_{s,t}$	-0.129*** (-4.883)	-0.006 (-0.460)	0.095*** (4.548)	0.021** (2.137)
$IV_{s,t}$	0.036*** (9.856)	0.010*** (4.405)	0.023*** (8.478)	-0.001 (-0.674)
Stock FE	Yes	Yes	Yes	Yes
Time Controls	Yes	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes	Yes
Adj. R ²	0.237	0.281	0.190	0.232

This table presents the results of panel regressions of $ES_{s,t}$ on $\log(SD_{s,t})$ for ATM call options (Panel A) and put options (Panel B) written on the stocks that are the constituents of the S&P500. The results are presented for two maturity buckets: 0-24 days to maturity and 25-48 days to maturity. $ES_{s,t}$ is the daily effective spread on day t for options on stock s . $\log(SD_{s,t})$ is the logarithm of the standard deviation of the intraday order flow distribution on day t for options on stock s , $\log(volume_{s,t})$ is the logarithm of the daily options volume, and $|OI_{s,t}|$ is the absolute value of the daily order imbalance (scaled by 10,000). $Return_{s,t}$ is the return of underlying stock on day t , and $IV_{s,t}$ is the average implied volatility of the options on stock s on day t . Other controls include firm size, stock volume, one-day and two-day lags of $ES_{s,t}$, and absolute values of the average delta, vega and gamma of the options on day t . Time controls include day-of-the-week, month-of-year, and year dummies. Standard errors are clustered at the day and stock level. The corresponding t-statistics are presented in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level.

Table 7: Exchange-Level Analysis

Panel A: Trade Absorption and Liquidity

	Call Options				Put Options			
	Regression of $Spread_{i,j,\tau}$		Regression of $\Delta Spread_{i,j,\tau+1}$		Regression of $Spread_{i,j,\tau}$		Regression of $\Delta Spread_{i,j,\tau+1}$	
	α	$D_{i,j,\tau}$	α	$D_{i,j,\tau}$	α	$D_{i,j,\tau}$	α	$D_{i,j,\tau}$
AXP	17.10***	-5.77***	-0.12***	0.28***	17.44***	-5.91***	-0.15***	0.37***
BA	22.05***	-10.06***	-0.03	0.83*	22.65***	-10.05***	-0.03	0.91**
CAT	13.17***	-3.83**	-0.07**	0.22***	13.19***	-3.90**	-0.08***	0.24***
DIS	19.00***	-9.19***	-0.01	0.52*	22.20***	-11.50***	-0.03	0.67**
DOW	16.46***	-4.20***	-0.20***	0.25**	16.31***	-4.67***	-0.32***	0.27***
HD	17.71***	-6.85***	-0.07	0.44**	17.94***	-6.65***	-0.09**	0.44***
IBM	14.88***	-4.33***	-0.04	0.25***	15.02***	-4.41***	-0.05*	0.25***
INTC	11.24***	-3.39***	-0.06***	0.29***	11.06***	-3.27***	-0.07***	0.30***
JNJ	18.45***	-6.43***	-0.10*	0.46***	19.42***	-6.51***	-0.13***	0.47***
JPM	10.19***	-3.40***	-0.05**	0.17***	10.48***	-3.29***	-0.05***	0.17***
KO	16.03***	-4.95***	-0.12**	0.32*	15.97***	-4.97***	-0.16***	0.34***
MMM	18.50***	-5.71***	-0.20***	0.41**	18.09***	-5.46***	-0.25***	0.40***
MRK	21.05***	-5.72***	-0.09**	0.60***	21.28***	-5.88***	-0.10**	0.62***
MSFT	18.04***	-7.91**	-0.01	0.37***	17.51***	-7.50**	0.00	0.40***
PG	16.88***	-5.04***	-0.10***	0.33***	16.18***	-4.99***	-0.15***	0.32***
WMT	16.18***	-6.44***	-0.06	0.35*	15.81***	-5.60***	-0.09***	0.32***

Panel B: Daily Regressions of Exchange-Specific $ES_{i,s,t}$

	Call Options				Put Options			
	All Exchanges		Six Major Exchanges		All Exchanges		Six Major Exchanges	
$\log(SD_{i,s,t})$	0.0079*** (9.18)	0.0038*** (8.73)	0.0081*** (10.21)	0.0036*** (7.70)	0.0072*** (11.06)	0.0041*** (8.44)	0.0076*** (10.99)	0.0044*** (10.15)
$\log(Volume_{i,s,t})$	-0.0075*** (-7.24)	-0.0040*** (-7.18)	-0.0070*** (-7.61)	-0.0034*** (-5.40)	-0.0067*** (-7.62)	-0.0043*** (-7.86)	-0.0062*** (-7.95)	-0.0042*** (-7.15)
$ OI _{i,s,t}$	0.0033 (1.49)	0.0029* (1.73)	0.0010 (0.03)	0.0013 (1.26)	0.0030 (0.72)	0.0012 (0.66)	-0.002 (-0.61)	-0.0002 (-0.20)
Exchange FE	Yes	No	Yes	No	Yes	No	Yes	No
Day FE	Yes	No	Yes	No	Yes	No	Yes	No
Stock x Day FE	No	Yes	No	Yes	No	Yes	No	Yes
Adj. R ²	0.22	0.55	0.24	0.63	0.15	0.45	0.20	0.59

Panel A reports the results of panel regressions of the quoted spread ($Spread_{i,j,\tau}$) and its change following each trade ($\Delta Spread_{i,j,\tau+1}$), estimated across exchanges. The regressions include a constant (coefficient α) and a dummy variable $D_{i,j,\tau}$, which equals one if the trade of option j at time τ was executed on exchange i . Regressions are run separately by ticker and by option type (calls and puts), and include day fixed effects. Standard errors are clustered at the day and exchange level. All coefficients are multiplied by 100. Panel B presents results from panel regressions of the exchange-specific effective spread, $ES_{i,s,t}$, on $\log(SD_{s,i,t})$, the logarithm of option order flow volatility based on trades executed on exchange i . The regressions also include $\log(Volume_{i,s,t})$, the logarithm of daily option volume for stock s on exchange i , and $|OI_{s,i,t}|$, the absolute value of daily order imbalance (scaled by 10,000) absorbed by exchange i . Standard errors are clustered at the day, stock, and exchange level.

Table 8: Log(SD) on Information and Inventory Days

Maturity	0-24		25-48	
	Calls	Puts	Calls	Puts
EA _{s,t}	-0.229*** (-20.391)	-0.207*** (-16.954)	-0.189*** (-20.217)	-0.173*** (-16.163)
EA _{s,t-1}	-0.141*** (-13.918)	-0.110*** (-9.480)	-0.099*** (-12.248)	-0.046*** (-5.084)
Third Friday	0.007 (1.067)	0.022*** (3.305)	-0.054*** (-7.582)	-0.052*** (-7.221)
BoM	-0.025*** (-4.972)	-0.017*** (-3.338)	0.063*** (7.056)	0.047*** (4.893)
EoM	-0.010* (-1.822)	-0.012** (-2.202)	0.131*** (9.845)	0.109*** (8.861)
RV _{s,t}	-8.818*** (-3.091)	-9.619*** (-3.050)	-6.974*** (-4.296)	-8.998*** (-4.509)
Stock FE	Yes	Yes	Yes	Yes
Time Controls	Yes	Yes	Yes	Yes
Additional Controls	Yes	Yes	Yes	Yes
Adj. R ²	0.749	0.753	0.776	0.787

This table presents the results of panel regressions of $\log(SD_{s,t})$ on a set of event-day dummies for at-the-money (ATM) call and put options written on S&P 500 constituent stocks. Results are reported separately for two maturity buckets: 0–24 days and 25–48 days to maturity. The dependent variable, $\log(SD_{s,t})$, is the logarithm of the standard deviation of the intraday order flow distribution for options on stock s on day t . The regression includes several event indicators: $EA_{s,t}$ ($EA_{s,t-1}$) equals one on the earnings announcement day (or the day before) for stock s ; *Third Friday* equals one on the third Friday of each month; *BoM* and *EoM* are dummies for the first and last trading days of each month, respectively. $RV_{s,t}$ denotes the realized volatility of stock s on day t , computed as the sum of squared five-minute intraday returns. Control variables include day-of-the-week, month-of-year, and year fixed effects, as well as $\log(Volume_{s,t})$, the logarithm of the daily option volume, and $|OI_{s,t}|$, the absolute value of the daily option order imbalance for stock s . Standard errors are clustered at the day and stock level. t-statistics are shown in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 9: Robustness: Dollar Spread and Realized Option Volatility for SPX Options

Panel A: Dollar Spread and $\log(\text{SD})_t$										
Maturity	0	1-6	7-13	14-20	21-27	28-34	35-41	42-48		
$\log(\text{SD})_t$	0.035*** (8.706)	0.046*** (7.508)	0.020*** (4.741)	0.009** (2.168)	0.016*** (3.994)	0.025*** (5.590)	0.006 (1.500)	-0.005 (-1.122)		
Controls	Yes									
N	1035	2090	2899	2797	2735	2596	2362	2165		
Adj. R ²	0.550	0.688	0.743	0.767	0.770	0.766	0.792	0.760		
Maturity	0	1-6	7-13	14-20	21-27	28-34	35-41	42-48		
$\log(\text{SD})_t$	0.043*** (5.720)	0.066*** (7.277)	0.037*** (4.815)	0.024*** (4.350)	0.016*** (2.956)	0.023*** (3.268)	0.016*** (3.324)	0.020*** (3.072)		
Controls	Yes									
N	1035	2060	2903	2791	2738	2618	2374	2216		
Adj. R ²	0.517	0.652	0.713	0.749	0.746	0.742	0.765	0.760		

Panel B: ES_t and ORV										
Maturity	0	1-6	7-13	14-20	21-27	28-34	35-41	42-48		
$\log(\text{SD})_t$	0.027*** (4.868)	0.022*** (8.187)	0.007*** (6.926)	0.003*** (3.815)	0.003*** (5.557)	0.004*** (6.563)	0.001*** (3.314)	0.001*** (3.314)		
ORV _t	0.010*** (13.756)	0.053*** (5.520)	0.046*** (6.050)	0.029*** (3.621)	0.019*** (2.845)	0.012*** (5.401)	0.010*** (4.446)	0.016*** (4.889)	0.009*** (5.004)	0.016*** (4.506)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	1035	2090	2899	2899	2797	2735	2596	2362	2165	2165
Adj. R ²	0.374	0.514	0.328	0.356	0.510	0.647	0.568	0.684	0.533	0.604
Maturity	0	1-6	7-13	14-20	21-27	28-34	35-41	42-48		
$\log(\text{SD})_t$	0.026*** (5.949)	0.020*** (7.585)	0.008*** (8.002)	0.004*** (5.328)	0.003*** (4.407)	0.002*** (3.849)	0.002*** (3.839)	0.002*** (4.515)		
ORV _t	0.004*** (5.161)	0.003*** (2.657)	0.038*** (4.340)	0.034*** (4.100)	0.011*** (2.927)	0.017*** (5.406)	0.004* (1.216)	0.007*** (1.649)	0.009*** (4.747)	0.009*** (6.667)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	1035	2060	2903	2903	2791	2738	2618	2374	2216	2216
Adj. R ²	0.373	0.518	0.496	0.574	0.515	0.641	0.535	0.668	0.555	0.677

Panel A of the table presents time-series regressions of Dollar Spread_t on the volatility of the order flow $\log(\text{SD})_t$ for SPX call and put options. Dollar Spread_t is the daily effective dollar spread. Panel B presents time-series regressions of ES_t on the realized option volatility ORV_t , which is calculated as the sum of squared 5-minute average option returns on day t . The other variables are analogous to those analyzed in the baseline regression in Table 3. Dollar Spread_t is winsorized at 99.5% and 0.5% levels. The coefficients of $\log(\text{SD})_t$ are multiplied by 1,000 in Panel A. Standard errors are computed using Newey-West with the optimal lag suggested by Andrews and Monahan (1992), and corresponding t-statistics are presented in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level.

Table 10: Robustness: Dollar Spread and Realized Option Volatility for Individual Stock Options

Panel A: Dollar Spread _{s,t} and log(SD) _{s,t}					
	Calls		Puts		
	0-24	25-48	0-24	25-48	
log(SD) _{s,t}	0.114*** (28.583)	0.107*** (27.950)	0.114*** (26.861)	0.110*** (22.727)	
Stock FE	Yes	Yes	Yes	Yes	Yes
Time Controls	Yes	Yes	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes	Yes	Yes
Adj. R ²	0.541	0.564	0.530	0.559	

Panel B: $ES_{s,t}$ and $ORV_{s,t}$								
	Calls		Puts					
	0-24	25-48	0-24	25-48				
log(SD) _{s,t}	0.016*** (31.704)	0.007*** (29.301)	0.015*** (27.005)	0.006*** (24.448)				
ORV _{s,t}	0.018*** (9.316)	0.014*** (8.932)	0.009*** (5.118)	0.006*** (4.941)	0.026*** (19.724)	0.024*** (19.120)	0.011*** (4.474)	0.009*** (4.525)
Stock FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Other Controls	No	Yes	No	Yes	No	Yes	No	Yes
Adj. R ²	0.113	0.275	0.087	0.287	0.136	0.258	0.097	0.241

Panel A of the table presents panel regressions of Dollar Spread_{s,t} on the volatility of the order flow log(SD)_{s,t} for individual stock options. Dollar Spread_{s,t} is the daily effective dollar spread for options on stock s . Panel B presents panel regressions of $ES_{s,t}$ on the realized option volatility $ORV_{s,t}$, which is calculated as the sum of squared 5-minute average option returns for stock s on day t . The other variables are analogous to those analyzed in the baseline regression in Table 6. The coefficients of log(SD)_{s,t} are multiplied by 1,000. Dollar Spread_{s,t} is winsorized at 99.5% and 0.5% levels, and standard errors are clustered at the day and stock level.