The Rise of Algorithmic Retail Options Traders*

Diego Amaya[1], Pedro A. Garcia-Ares[2], Neil D. Pearson[3,4], and Aurelio Vasquez[5]

August 12, 2026

Abstract

Trading in SPX 0DTE options spikes at fixed times. Using OPRA transaction records from May 2022 through January 2026, we show that the bursts begin within seconds of each half-hour mark, finish within a minute, and grow each year. They consist of a small set of defined-risk strategies whose strikes, widths, and credits resemble templates of retail automation platforms, and can be identified from those characteristics alone, without the clock. At spike times quoted spreads narrow while effective spreads widen: predictability attracts liquidity, and simultaneity consumes it. The order flow is synchronized though the algorithms never learn, communicate, or interact.

JEL Classification: G12, G14, G18, G23

Keywords: Retail option trading; 0DTE; algorithmic trading; market microstructure; fintech

* We thank SpiderRock LLC for generously providing access to the options trade and quote data.

[1] Lazaridis School of Business & Economics, Wilfrid Laurier University, 75 University Avenue West, Waterloo, Ontario, Canada, N2L 3C5. Tel.: +1 (519) 886-9351 ext. 4428. E-mail: damaya@wlu.ca.

[2] Instituto Tecnológico Autónomo de México (ITAM), Río Hondo No. 1, Col. Progreso Tizapán, México, D.F., México, 01080. Tel.: +52 (55) 5628 4000 ext. 6518. E-mail: pedro.garcia@itam.mx.

[3] Department of Finance, University of Illinois at Urbana-Champaign, 1206 South Sixth Street, Champaign, Illinois, 61820. Tel.: +1 (217) 244-0490. E-mail: pearson2@illinois.edu.

[4] Canadian Derivatives Institute, 3000, Chemin de la Côte-Sainte-Catherine, Montréal, Québec H3T 2A7, Canada.

[5] Instituto Tecnológico Autónomo de México (ITAM), Río Hondo No. 1, Col. Progreso Tizapán, México, D.F., México, 01080. Tel.: +52 (55) 5628 4000 ext. 6518. aurelio.vasquez@itam.mx.

The Rise of Algorithmic Retail Option Traders

August 12, 2026

Abstract

Trading in SPX 0DTE options spikes at fixed times. Using OPRA transaction records from May 2022 through January 2026, we show that the bursts begin within seconds of each half-hour mark, finish within a minute, and grow each year. They consist of a small set of defined-risk strategies whose strikes, widths, and credits resemble templates of retail automation platforms, and can be identified from those characteristics alone, without the clock. At spike times quoted spreads narrow while effective spreads widen: predictability attracts liquidity, and simultaneity consumes it. The order flow is synchronized though the algorithms never learn, communicate, or interact.

JEL Classification: G12, G14, G18, G23

Keywords: Retail option trading; 0DTE; algorithmic trading; market microstructure; fintech

1. Introduction

A growing literature asks whether autonomous algorithms can produce coordinated outcomes without communication. In product markets, independently operating pricing algorithms can learn to sustain supra-competitive prices, both in simulation (Calvano, Calzolari, Denicolò, and Pastorello, 2020) and in the field (Brown and MacKay, 2023; Assad, Clark, Ershov, and Xu, 2024). In financial markets, the evidence on such coordination comes largely from simulation: algorithmic market makers can fail to learn competitive pricing (Colliard, Foucault, and Lovo, 2026), and speculators trained by reinforcement learning can sustain collusive profits (Dou, Goldstein, and Ji, 2025a) or plan across periods to create and then burst price bubbles (Dou, Goldstein, and Ji, 2025b). These mechanisms arise because algorithms learn from their environment or respond strategically to one another. Field evidence is difficult to obtain because deployed trading algorithms do not reveal their rules: an observer sees the orders they send but not the schedule, the strategy, or the risk limits behind them.

We provide evidence of a different mechanism. Algorithms can generate synchronized aggregate behavior without learning, communication, or strategic interaction. When many investors delegate trading decisions to automated rules built on common technological platforms, independently operating algorithms inherit common entry times, strategy templates, and risk constraints. Shared technological infrastructure alone can transform independent decisions into predictable, synchronized order flow. This is synchronization without coordination.

We identify this mechanism in S&P 500 index (SPX) options with zero days to expiration (0DTE). At predetermined clock times on every trading day, 0DTE trading surges. Trade arrivals jump within seconds of half-hour marks and return toward normal levels within 30 to 60 seconds. They are concentrated in a small set of, defined-risk complex strategies—short call and put vertical spreads and iron condors—with characteristics that closely resemble the strategies and parameter choices offered by retail automation platforms. The evidence points to a new form of retail trading: individually small and independently chosen, but algorithmically synchronized in aggregate.

The rules that generate this flow are visible in the data. The schedule appears in the timing of the trades, the template and the risk limits in the strikes, widths, and credits of the positions, and the spread of the technology in the arrival of activity at clock times where there was none previously. Institutional execution algorithms leave footprints in the tape, but their parameters are

proprietary and must be inferred. Here the synchronization, the rules behind it, and its consequences for market quality can be measured in the same data.

We study this phenomenon using transaction-level Option Price Reporting Authority (OPRA) records for SPX options traded on the Cboe C1 options exchange from May 2022 through January 2026. The sample begins when expirations on every trading day became available in May 2022 and spans the subsequent rapid growth of 0DTE trading. By 2025, 0DTE contracts accounted for 59% of SPX option volume, and their share subsequently exceeded 60%.[1] The large gamma and rapid time decay of at- and near-the-money 0DTE options make execution timing consequential, while defined-risk spreads can be implemented using a small number of parameters, making these contracts particularly suitable to automated trading. Retail participation is substantial: Cboe estimates attribute roughly one-half to three-fifths of SPX 0DTE volume to retail investors.

Our first result is the emergence and diffusion of deterministic periodicity in 0DTE trading. Trade counts and volume spike most prominently at 10:00, 10:30, and 11:00 and again at 13:00, 13:30, and 14:00. At these times, trade arrivals run roughly twice their level in surrounding minutes, while complex trade arrivals more than triple at the 13:30 spike. In second-by-second data, there is no buildup immediately beforehand: trading jumps within seconds of the clock time and dissipates within 30 to 60 seconds. The largest bursts occur in the early afternoon, when scheduled news is sparse and conventional information-based explanations for concentrated trading are least natural.

The evolution of these patterns over time is also consistent with the diffusion of automated trading. When daily SPX expirations were introduced in 2022, the spikes were relatively modest and concentrated at half-hour marks. They subsequently grew relative to background trading, while periodicity spread to additional clock times. Quarter-hour marks, which were largely indistinguishable from surrounding minutes in 2022, become clearly visible from 2023 onward, and smaller spikes eventually emerge at five-minute intervals. By 2026, the 13:45 spike reaches approximately twice the activity of its neighboring minutes. The gradual intensification of existing

[1] The 59% figure is from Cboe's full-year 2025 volume report (https://ir.cboe.com/news/news-details/2026/Cboe-Global-Markets-Reports-Trading-Volume-for-December-and-Full-Year-2025/default.aspx). The share continued to rise in 2026, reaching 63% in February and 66% in July (Cboe monthly volume reports, https://ir.cboe.com/news/news-details/2026/Cboe-Global-Markets-Reports-Trading-Volume-for-February-2026/default.aspx and https://ir.cboe.com/news/news-details/2026/Cboe-Global-Markets-Reports-Trading-Volume-for-July-2026/default.aspx).

spikes together with the appearance of new scheduling frequencies is more consistent with the diffusion of trading technology than with a fixed execution protocol.

Several features of the spikes point to automated execution rather than discretionary trading. Bursts that begin within seconds of an exact clock time, with no buildup beforehand, and dissipate within a minute are difficult to reconcile with independent manual order submission. The spikes occur at the same clock times every day, are largest in the early afternoon when scheduled news is sparse, and are not concentrated around macroeconomic announcements. They also invert the classic intraday pattern: trading in equities and options has long concentrated at the open and close, a U-shape rationalized by information and liquidity clustering (Admati and Pfleiderer, 1988; Foster and Viswanathan, 1990; Chan, Christie, and Schultz, 1995), whereas the largest spikes here arrive in the middle of the afternoon, when these explanations predict relatively little activity. Finally, the spikes are concentrated in multi-leg strategies and in 0DTE contracts, for which automated specification and precise execution timing are most natural.

OPRA records do not identify investors, so we cannot directly observe who submits these orders.[2] The observable characteristics of the flow, however, point to retail automation. The trades are small: more than four-fifths involve five contracts or fewer, and the fraction with total premium below $25 nearly doubles during the bursts, from 6.6% to 11.9%. The trades follow standardized patterns, with common spread widths, strike-selection rules, and defined-risk structures that resemble the strategy menus of retail automation platforms. Their execution also differs from the institutional algorithms emphasized in the literature, which typically seek to reduce market impact and information leakage by spreading execution through time; the flow we document instead arrives at predictable clock times in identifiable packages. Finally, Cboe observes account information that OPRA omits, and its own analyses attribute these bursts to automated systems used by retail traders.[3]

---

[2] OPRA transaction records report neither investor identity nor trade direction. We infer direction from quote midpoints, following Muravyev (2016) for single-leg trades and using midpoint-implied net premia for multi-leg packages (Section 3), and we identify the origin of the flow from its observable characteristics (Sections 4 through 6). Exchange open-close datasets classify volume by customer type, by opening versus closing transactions, and by buy and sell direction. We are in the process of obtaining these data, and a future version of the paper will use them to validate both the trade signing and the identification of the flow.

[3] Schwartz, Henry, *The State of the Options Industry:2025 Year in Review*, Cboe Global Markets, January 21, 2026. The slide deck and data visuals include references to retail algorithmic execution tools such as Trade Automation Toolbox (IBKR), Option Alpha (Tastytrade, Tradier, TradeStation), Trade Steward (TD), Option Omega and Brokerbotics and their association with intraday volume spikes.
https://www.cboe.com/insights/posts/the-state-of-the-options-industry-2025/

The institutional environment provides a mechanism for these patterns. Retail brokers provide APIs through which investors can submit orders programmatically, while no-code platforms allow users to deploy trading bots from pre-built entry, exit, and risk-management rules without writing software. These platforms package multi-leg strategies into templates—vertical spreads, iron condors, and butterflies—and provide menus of entry times, deltas, widths, position sizes, profit targets, and exit rules. Once an investor selects the rules, subsequent orders require no new decision. Each account acts independently, but many accounts drawing from the same menus and schedules can submit similar positions at precisely the same times. The decision maker remains human, but the execution process becomes algorithmic.

The trading inside the spikes is not simply more of the surrounding flow; its composition changes sharply. Over the last two years of the sample, the share of complex, multi-leg trades increases by 12.2 percentage points during spike minutes, from 37.1% to 49.4%. At 13:30, the complex share rises by 21.3 percentage points. The increase is concentrated in defined-risk, short-premium strategies. Short put vertical spreads nearly double their share of trading, from 3.6% to 6.7%, short call verticals increase from 3.1% to 5.5%, and short iron condors from 0.6% to 1.2%. Within these strategies, spread widths increase, short strikes move closer to the index, net credit per unit of width falls, and trade sizes decline. These characteristics move jointly in the direction implied by strategies built around target deltas and maximum-loss constraints. Scheduling therefore predicts not only when investors trade but also what they trade. Positions held to expiration earn positive average profits, but these estimates are imprecise. Regardless, we find no evidence that these strategies are systematically wealth-depleting.

The trades in the bursts can also be distinguished from surrounding trades without using the clock. Using 11.1 million vertical-spread transactions from January 2024 through January 2026, we compare trades executed in the first minute of each half-hour window with trades executed during quiet minutes later in the same half hour. Arrival rates in the first minute are 2.6 times those in the quiet minutes. A gradient-boosted classifier trained on trade characteristics and preceding market conditions, but given no information about the time of day, distinguishes burst-window from quiet-window trades out of sample with an ROC area of 0.685, and its performance remains stable over time. The synchronized flow therefore differs not only in when it arrives but also in the positions it creates.

Finally, we examine how markets respond when retail order flow becomes predictable and synchronized. Liquidity conditions change sharply at the spike times, and standard measures give opposite readings. Quoted spreads for the dominant strategies narrow by \$0.004 for short iron condors to \$0.009 for short call verticals, although this narrowing appears only in the afternoon; at the 10:00 spike, quoted spreads instead widen. Over the same intervals, effective spreads widen by approximately \$0.004 to \$0.006, so traders demanding immediacy pay more when many similar orders arrive within seconds. Realized spreads at the ten-minute horizon move in opposite directions for call-side and put-side verticals, consistent with differences in the hedging demands the two positions create. Displayed liquidity therefore improves at the same times that the cost of immediacy rises. Predictability attracts liquidity; simultaneity consumes it.

The liquidity results are related to sunshine trading (Admati and Pfleiderer, 1991), in which preannounced uninformed orders attract liquidity supply. Here, predictability arises unintentionally because investors delegate execution to algorithms that use common schedules. Liquidity suppliers quote more aggressively against the predictable flow, but the simultaneous arrival of similar orders increases the cost of immediacy. Automation therefore creates two opposing forces: a benefit from predictability and a cost from synchronization.

Our findings contribute to three literatures. First, we provide field evidence on a distinct mechanism through which algorithms can generate synchronized outcomes. Existing work on autonomous algorithms emphasizes learning and strategic interaction (Calvano et al., 2020; Colliard, Foucault, and Lovo, 2026; Dou, Goldstein, and Ji, 2025a, 2025b). We show that neither is necessary for synchronization. Common technological infrastructure can align independently chosen algorithms mechanically. In our setting, aggregate synchronization arises not because algorithms learn to cooperate or respond strategically to one another, but because they inherit common rules from the same technological environment. The traders we document are of the kind these models take as given rather than the kind that does the coordinating. They do not learn, plan, or respond to one another, and their flow is correlated because the rules generating it come from the same platforms.

Second, we document a new consequence of the automation of retail trading. The existing literature characterizes retail flow as largely discretionary and driven by attention, sentiment, gambling preferences, and information (Barber and Odean, 2000, 2008, 2013; Boehmer et al., 2021; Barber et al., 2022; Poteshman and Serbin, 2003; Lakonishok et al., 2007; Bryzgalova et al.,

2023; de Silva et al., 2025; Eaton et al., 2026). Automation changes the aggregate statistical properties of this flow. Individual transactions remain small and originate from independent accounts, but their timing and form become correlated. In aggregate, many retail investors can therefore begin to resemble a single systematic trader.

Third, we identify a new form of predictable liquidity demand and show how markets respond to it. Classic market-microstructure models study how information and liquidity motives determine when traders execute (Kyle, 1985; Admati and Pfleiderer, 1988, 1991). Here timing is determined partly outside the market, by rules selected in advance and implemented through common technological platforms. Liquidity suppliers can anticipate the resulting flow, but synchronization creates concentrated demand for immediacy.

More broadly, we find that fintech platforms do not simply reduce the cost of trading; they standardize decision-making. When many investors draw their strategies from the same templates, defaults, and schedules, platform design can determine not only individual behavior but also the correlation of behavior across users. The algorithms in our setting do not require sophisticated artificial intelligence, reinforcement learning, or communication with one another. Their economic importance comes instead from scale and commonality. Synchronization need not be programmed across algorithms or emerge through learning. It can arise because independent algorithms inherit common rules from the same technological infrastructure.

The remainder of the paper proceeds as follows. Section 2 describes the fintech environment and explains how retail options trading has become automated. Section 3 describes the data and variable construction. Section 4 documents the intraday trading patterns in SPX 0DTE options and how they have changed over the sample. Section 5 analyzes the composition, structure, and profitability of the trades executed at the spike times. Section 6 shows that these trades can be identified from their characteristics alone. Section 7 examines liquidity and execution costs around the spikes. Section 8 concludes.

2. Fintech and the Automation of Retail Options Trading

Retail investors are increasingly able to trade SPX 0DTE options through automated rules. The rules specify when to enter a position, which option strategy to trade, which strikes to use, and when to close the position. Brokers and fintech platforms supply this automation through APIs, trading bots, and predefined strategy templates. Automated trading leaves patterns in transaction data that discretionary trading does not. An investor who responds to news or to changing market

conditions trades at different times on different days. An investor whose rules specify an entry time trades at the same time every day, and an investor who trades from a strategy template submits the same combination of options each time the rule fires.

2.1. From discretionary retail flow to rules-based retail flow

As noted in the introduction, retail investors account for roughly one-half to three-fifths of SPX 0DTE trading activity.[4] An important change in this market is the availability of tools that allow retail investors to automate their trading. Broker APIs allow investors to write programs that monitor market conditions and submit orders when predetermined conditions are satisfied. No-code platforms provide the same functionality without requiring investors to write computer code.

These tools reduce the cost of implementing rule-based strategies. The investor specifies an entry time, an option strategy, the strikes to trade, the size of the position, and rules for closing it. Once the rules are set, the orders are submitted automatically. The decision is made once, when the rule is written, and executed on every subsequent day.

2.2. APIs, trading bots, and strategy templates

Retail automation relies on several related technologies. Brokers provide APIs that allow investors to code trading strategies and submit orders automatically. These programs monitor prices, volatility, and other market variables and submit orders when predetermined conditions are met. Brokers also provide execution algorithms that determine how and when orders are submitted. Interactive Brokers, for example, offers more than 100 order types and algorithms intended to improve execution, control timing, and manage risk.[5] Tools of this type were traditionally used primarily by professional investors.

A second form of automation removes the need to write code. Option Alpha, for example, allows users to convert a backtest into a bot that follows predetermined entry, exit, and position rules.[6] Fintech platforms also supply templates for multi-leg strategies such as iron condors, credit spreads, and butterflies, so that the investor does not construct each combination of options separately. The investor chooses a small number of parameters: the option deltas, the width of the spread, the entry time, a profit target, and an exit time.

---

[4] Mandy Xu, "0DTEs Decoded: Positioning, Trends, and Market Impact," *Cboe Insights*, May 2, 2025. https://www.cboe.com/insights/posts/0-dt-es-decoded-positioning-trends-and-market-impact/

[5] Interactive Brokers, "IBKR Order Types, Algos and Tools," product documentation (accessed 2026).

[6] Option Alpha, "Automatically Generate Bots From Any Backtest," October 29, 2024. https://optionalpha.com/blog/automatically-generate-bots-from-any-backtest

Consider the SPX 0DTE iron condor bot that Option Alpha provides as an example. [7] The iron condor is a credit trade (it is "sold") that is equivalent to a simultaneous sale of call and put vertical spreads in which all four options are OTM.[8] The investor selects a delta for the short strikes, a width for each spread, and an entry time, and the bot then submits the same four-leg order at that time on every trading day. The delta choice determines which strikes are traded, given the level of the index. The width determines the maximum loss and therefore the margin the position requires. The entry time determines when the order reaches the market. Two investors who make the same three choices submit orders at the same time, in the same strategy, at strikes that differ only by the level of the index when each order is sent, even though neither investor knows the other exists. The parameters are chosen from the same menus, and the menus offer a small number of round values.

Investors need not choose the parameters at all. Platforms publish ready-made bots, together with the backtests that motivate them, which an investor adopts with a single decision. Every investor who adopts the same published bot submits the same strategy, at the same strike selection rule, at the same time of day, for as long as the bot runs. Standardization is then complete: the trades differ only in size. [9]

2.3. Defined-risk strategies and common trading rules

Retail SPX 0DTE trading is concentrated in defined-risk positions. Cboe reports that more than 95% of SPX 0DTE trades have limited risk, either because investors buy options outright or because short option positions are combined with long options in spreads. Naked short option positions account for a only small fraction of trading.

This pattern follows from the restrictions retail investors face. Brokers impose margin requirements and risk controls, and platforms emphasize strategies with a predetermined

[7] Option Alpha, "Peg Research Based 0DTE SPX Bot," October 17, 2024. https://optionalpha.com/videos/peg-research-based-0dte-spx-bot

[8] Letting K_i denote the strike price of an option, an iron condor involves four strikes $K_1 < K_2 < K_3 < K_4$. Most commonly, the strikes are chosen so that the middle two strikes bracket the underlier. In a sold iron condor, the investor sells an OTM put vertical spread; that is, he sells a put with strike K_2 and buys a put with strike K_1 that limits the losses in the event that the underlier declines in price. The investor simultaneously sells an OTM call vertical spread by selling a call with strike K_3 and buying a call with strike K_4. Recognizing that the option deltas map to moneyness, that is, the option strikes, the investor specifies the short strikes K_2 and K_3 by specifying the deltas. The widths of the put and call vertical spreads are $K_2 - K_1$ and $K_4 - K_3$; these determine the maximum possible payout if the iron condor is held to expiration, and thus the required margin.

[9] The published parameters of these bots, and the backtests that accompany them, can in principle be matched to the trades we observe. We leave this comparison to future work.

maximum loss. Automated retail strategies therefore use vertical spreads, iron condors, and other multi-leg positions rather than naked short options.

These restrictions reinforce the standardization. Defined-risk strategies are submitted as complex orders, so investors who use similar delta targets and spread widths trade similar combinations of options. Exit rules work the same way: a bot closes a position at a profit target, at a loss limit, or at a set time before expiration, so common exit rules concentrate closing trades as well as opening trades.[10]

2.4. Synchronization without coordination

The mechanism requires no coordination among investors. Investors who schedule entries at the same clock time submit orders at that time, and investors who use the same templates submit orders with the same characteristics. Neither requires that they communicate, observe one another, or intend to trade together. As automation spreads, trading concentrates at more clock times: some investors schedule at half-hour intervals, others at quarter-hour or five-minute intervals, and each scheduling convention adds its own times to the pattern. Bloomberg reported in May 2025 that retail trading bots submit orders simultaneously and generate visible increases in options trading volume.[11]

This mechanism gives us four predictions. First, automated trading concentrates trades at fixed clock times. Second, the concentration is strongest for complex trades, because the templates are multi-leg strategies. Third, trades at these times share characteristics, including common deltas and spread widths. Fourth, entry and exit activity occurs at times consistent with common scheduling rules.

The mechanism also separates this flow from the trading that surrounds it. The timing is set in advance of the trade, so the flow cannot respond to information arriving at the moment it executes. The form is set by the template, so the flow is standardized. Because both are fixed before the order is sent, an observer who has seen the pattern can anticipate when the orders will come and what they will look like. We test these predictions in the empirical analysis below.

[10] Broker policies to liquidate positions customers lack the buying power to settle at expiration can also generate synchronized closing flow (Amaya, Garcia-Ares, Pearson, and Vasquez (2026)).

[11] Bernard Goyder, "At 10 AM, Stock Options Soar as Retail Traders Unleash New Bots," *Bloomberg*, May 31, 2025. https://www.bloomberg.com/news/articles/2025-05-31/at-10-am-stock-options-soar-as-retail-traders-unleash-new-bots

3. Data and Variable Construction

3.1. Data sources and sample selection

Our data consist of trade and quote records from the Options Price Reporting Authority (OPRA), obtained from SpiderRock. The data include all transactions in S&P 500 Index options (SPX) executed on the Cboe C1 exchange. Because SPX options are listed exclusively on Cboe C1, these records cover all trades in the contracts. Our sample spans May 19, 2022, through January 31, 2026. Daily expirations became available at the start of this window, so the sample covers zero-day-to-expiration (0DTE) trading from the time when all expiration days first became available through its growth into the largest segment of SPX volume.

We apply standard filters following Bryzgalova, Pavlova, and Sikorskaya (2023). We restrict the sample to regular trading hours (9:30 to 16:00 EST) and exclude canceled trades, transactions with non-positive sizes or prices, and quotes with crossed markets, in which the best bid exceeds the best ask. To limit the influence of reporting errors, we require trade prices to lie within the prevailing quote interval, defined as the best bid minus the spread and the best ask plus the spread.

For single-leg trades, we aggregate prints that share the same contract, quote time, exchange identifier, trade price, and condition code into a single observation. We infer trade direction following Muravyev (2016). We first compare the trade price to the midpoint of the BBO quotes attached to the trade record and classify the trade as buyer-initiated (seller-initiated) if the price is above (below) that midpoint. We estimate the trade direction of multi-leg trades as described below.

3.2. Identification of complex strategies

Complex (multi-leg) trades are executed in either the complex order book, the complex price improvement auction, or on the exchange floor. In each case, the trade is executed as a package at a single net price. The exchange then assigns prices to the constituent legs that are consistent with the net price, and disseminates the trades, including to OPRA. OPRA then disseminates the trades in the legs to its subscribers.

When OPRA disseminates trades, it includes the legs of complex trades in the same quote stream with the single-leg trades. The OPRA trade records include a flag that identifies whether a reported trade was a component of a complex trade. However, the records do not identify which other legs belong to the same order or which strategy the order represents. We therefore reconstruct

the parent strategy from the component execution records in two steps: we first assemble packages of legs that were reported in the same millisecond and with the same underlying price, and we then classify each package into a strategy type from its structure.

3.2.1. Constructing complex-order packages

We group execution records into packages by sorting the data by ticker (SPX or SPXW), date, execution time, print type, and underlying price, and assigning a sequential leg counter within each group. Legs that belong to the same complex order are executed at the same instant and reported with the same underlying price, so this grouping recovers the order as it was submitted. The number of legs in a package is the maximum leg index within the group.[12]

We describe the construction for the two dominant retail structures, two-leg vertical spreads and four-leg iron condors; the same procedure applies to the other structures we report. For each leg in a package, we retain the option type, strike price, expiration date, executed premium, trade size, and the contemporaneous NBBO midpoint.

3.2.2. Classification of vertical spreads and iron condors

We classify packages using the characteristics of their legs and the net premium implied by the market midpoints.

Vertical spreads. A package is a vertical spread if it contains exactly two legs that satisfy four conditions: both legs are of the same option type, both share the same expiration date, the legs have distinct strike prices, and the trade sizes match.

For each vertical spread, we compute the executed net premium as the absolute difference between the leg prices, and the midpoint-implied net premium as the absolute difference between the leg midpoints. We assign direction by comparing the two. A trade is long, or buyer-initiated, if the executed net premium exceeds the midpoint-implied net premium, so that the initiator paid more than the midpoint or received a smaller credit than the midpoint. A trade is short, or seller-initiated, if the executed net premium is below the midpoint-implied net premium. Trades at which the two are equal are classified as neutral. Combining direction with option type yields four main

[12] Complex-order reconstruction from leg-level prints follows Li, Musto, and Pearson (2023), who group legs within a short time window; Galati, Hendershott, Khan, and Riordan (2026) tighten this to a 3-millisecond cap on package span. We require exact equality of the execution timestamp and the underlying price. Legs of a genuine package that print microseconds apart are therefore split, so strategies with more legs are undercounted relative to those with fewer.

categories: long and short call verticals and long and short put verticals, along with a small number of neutral trades.

Iron condors. A package is an iron condor if it contains exactly four legs that partition into one call vertical spread and one put vertical spread, with matched sizes and call strikes above put strikes. We compute the net package premium and the midpoint-implied premium as the sums of the corresponding values of the component vertical spreads, and assign direction as for vertical spreads. We distinguish symmetric iron condors, in which the call-spread and put-spread widths are equal, from asymmetric ones, and retain both in the main analysis.

Packages that do not match any of these structures are retained in the sample as unclassified complex trades, so the classification does not affect the counts of complex trading we report.

3.3. Measures of liquidity and execution quality

We measure market quality around the trading bursts with three standard spread measures, each expressed as a half-spread so that the three are directly comparable.

The quoted spread is computed from quotes in the single-leg limit order book. It is half the difference between the best offer and the national best bid at the time of the trade. It measures the cost of a round trip at the displayed single-leg quotes, and it does not depend on where the trade executes.

The effective spread is the difference between the trade price and the prevailing NBBO midpoint, signed by trade direction. It measures the cost of demanding immediacy and captures price improvement and execution inside or outside the quotes, which the quoted spread does not.

The realized spread measures market maker revenue net of adverse selection. It is the difference between the trade price and the NBBO midpoint at a horizon after the trade, signed by trade direction; we report results for horizons of one minute and ten minutes. The difference between the effective and realized spreads is the price impact of the trade.

For complex trades, we compute each measure for the package by aggregating the corresponding values of its constituent legs. Because the vertical spreads and iron condors are executed in the complex LOB or price improvement auction, the quoted spread from the single-leg LOB is generally not directly relevant to them. An exception is that orders in the complex LOB that are marketable against the single-leg BBO interact with the single-leg BBO.

4. Intraday Periodicity in SPX 0DTE Trade Intensity

If order flow is generated by rules that specify entry times, its timing is deterministic. This section tests that prediction. We aggregate trades at the minute and second level and find sharp, recurring spikes in trade intensity that occur at the same clock times every day, grow over the sample period, and are concentrated in complex, ultra-short-maturity strategies.

4.1 Trades and volume through the day

Figure 1 plots the average number of trades per minute across the trading day (Panel A) and the average volume per minute (Panel B). Both panels display a regular periodic structure: sharp spikes occur at half-hour intervals, most prominently at 10:00, 10:30, 11:00, and 11:30, and continuing through the afternoon with particularly large spikes at 13:00, 13:30, and 14:00. At the spike minutes, trade counts run roughly twice their level in the surrounding minutes, and complex-trade counts more than three times. The spikes rise abruptly from a flat baseline, and smaller secondary spikes appear at quarter-hour marks, which we examine in Section 4.2. The spikes appear in both the number of trades and total volume, so they reflect many orders arriving together rather than isolated large trades.

The pattern differs from the U-shaped intraday profile documented in equities and options (Admati and Pfleiderer, 1988; Chan, Christie, and Schultz, 1995). Activity does not build or decay around these times; it appears at the clock time and disappears. The largest spikes also occur in the early afternoon, when trading is normally at its lowest and when these explanations predict the least activity.

4.2 Changes over the sample period

Figure 2 plots the average number of trades per minute separately by year. Three features stand out.

First, activity increases over the sample, reflecting the growth of SPX 0DTE trading over the period.

Second, the spikes grow relative to the rest of the day. In the early part of the sample, periodicity is present but muted; the spikes are visible but small relative to baseline activity. By the end of the sample, they dominate the intraday profile, and a large fraction of daily trades occur in the narrow set of minutes at the scheduled times.

Third, spikes appear at more clock times. In 2022, elevated activity is confined to the half-hour marks: the quarter-hour marks average 1.05 times the activity of the surrounding minutes,

which is indistinguishable from noise. From 2023 onward, the quarter-hour marks average between 1.3 and 1.5 times surrounding activity, and by the end of the sample smaller spikes appear at five-minute marks, rising from 1.02 times surrounding activity in 2022 to 1.12 times in 2026. The 13:45 mark shows this most clearly: indistinguishable from its neighbors in 2022, it reaches roughly twice their level by 2026. Half-hour spikes are the largest, quarter-hour spikes are smaller, and five-minute spikes are smaller still. Scheduling at round clock intervals produces exactly this, and information arrival has no reason to.

The evolution is also informative for identification. A fixed institutional execution protocol would be present from the start of the sample. Gradual intensification, together with the appearance of spikes at additional clock times, is what adoption of a new tool produces.

4.3 The spikes at the second level

Figure 3 aggregates trading activity second by second, centered on selected half-hour marks. Trades arrive abruptly at the mark and decay within 30 to 60 seconds. In recent years, the number of trades in the peak second reaches four to six times its level over the preceding minute. There is no buildup in the seconds before the mark.

A pattern that begins at the mark with no buildup and ends within a minute is inconsistent with human reaction times and with manual order submission, which would be spread over the surrounding seconds. The spike is not a window of elevated activity but a surge lasting seconds. Retail order flow that responds to news or to attention builds over minutes or hours (Da, Engelberg, and Gao, 2011; Cookson and Niessner, 2020); this does not.

4.4 Simple and complex trades

Figure 4 separates simple trades (Panel A) from complex trades (Panel B) using the OPRA transaction codes. The two panels differ sharply. Panel A (simple trades) shows spikes, but they are much less pronounced than the spikes in complex trades displayed in Panel B. Complex trades display the full set of half-hour spikes, with magnitudes resembling those in Figure 1. The spikes are therefore primarily driven by complex strategies rather than by single-leg trades.

This is what the mechanism predicts. Complex trades are the multi-leg positions that platform templates supply, and they require legs to be submitted together, which is straightforward for an automated system and difficult for a person entering orders by hand. Prior work often treats retail order flow as homogeneous (Barber et al., 2009); here the trades that occur at fixed times are concentrated in the subset that templates produce.

4.5 Trades by maturity

Figure 5 splits average trades by maturity. The half-hour spikes are concentrated in 0DTE contracts. The spikes weaken as maturity increases and are largely absent for longer-dated options.

Timing matters most at the shortest maturity. Gamma and theta are largest for 0DTE contracts, so the payoff of a position depends on when it is opened and closed in a way that it does not for longer-dated options. Retail option activity in prior work spans a range of maturities and does not show comparable clustering (Lakonishok et al., 2007; Eaton et al., 2026); in the 0DTE segment, the time of day at which a position is opened is itself a strategy parameter.

The evidence in this section is consistent with the first prediction. The spikes occur at fixed clock times, recur across days, are largest in the afternoon when scheduled news is sparse, and are concentrated in complex 0DTE strategies. They are difficult to attribute to information arrival, to macroeconomic announcements, or to market making. The next section examines what trades at these times.

5. The Composition of Trading at the Spike Times

The clock times show when the automated strategies trade. What trades at those times shows which rules they follow. This section examines the composition, structure, and performance of the option trades executed at the half-hour marks between January 2024 and January 2026, the most active years of the sample. Trading at these times shifts toward small, complex, short-premium positions: tight vertical spreads and iron condors, entered at fixed times of day and closed by rule. These are the positions that platform templates supply, and they are not the positions that discretionary institutional trading produces.

5.1. Option characteristics at the spike minutes

Table 1 compares option characteristics in the spike minute with those in the preceding five minutes. We compute averages across all half-hour spikes from 10:00 to 15:30 and compare them with the averages over the five minutes before each spike. We exclude the open and the close, where trading follows well-known patterns unrelated to the spikes we study. The comparison shows whether the spikes scale up the trading that surrounds them or change its composition.

The results show that the composition of trades in the spike differs from the preceding five minutes. The share of complex trades rises by 12.2 percentage points, from 37.1% in the pre-spike window to 49.4% during the spike, and the share of single-leg trades falls correspondingly. The spikes are therefore not periods of intensified background activity but moments when traders

choose multi-leg positions. In the intraday patterns studied in equity markets, predictable clustering of volume does not change what is traded (Admati and Pfleiderer, 1988; Foster and Viswanathan, 1990); here it does.

Trade size falls at the same time. The fraction of trades with total premium below $25 rises by 5.3 percentage points, from 6.6% to 11.9%. Small trades are commonly associated with retail participation (Barber et al., 2009; Boehmer, Jones, Zhang, and Zhang, 2021), and in earlier work on retail option trading they are single-leg or lottery-style positions. Here the small trades are the multi-leg ones.

Traders also shift toward out-of-the-money strikes, particularly on the short side of spreads. This is systematic positioning rather than the lottery demand documented by Kumar (2009), Bali, Cakici, and Whitelaw (2011), and Boyer and Vorkink (2014).

These effects vary over the day, and the averages understate them at the most active times. Spike by spike, the increase in the complex share ranges from 6.1 percentage points at 10:00 to 21.3 percentage points at 13:30 (*t*-statistic 48.7), stays above 18 percentage points through 14:30, and falls to 3.3 percentage points at 15:30 (Internet Appendix Table IA1). The share of trades with premium below $25 follows the same profile, rising by 9.6 percentage points at 13:30 against 2.2 at 10:00. The automation platforms described in Section 2 direct entries away from the open and into the early afternoon, and the largest volume spikes occur at the same times.

5.2. Strategies and trade direction

Table 2 splits trading activity by strategy type and trade direction. Simple strategies decline across categories, and the decline is largest in calls, the traditional instrument of retail speculation. The share of long calls falls by 2.0 percentage points, from 16.6% to 14.6%, and short calls fall by 1.6 percentage points, from 15.8% to 14.3%. The bursts are therefore not driven by demand for lottery-like options. Simple trades appear to be background order flow, tied to discretionary positioning or hedging.

Complex strategies increase, and the increase is almost entirely in defined-risk, short-premium positions. Short put verticals nearly double, rising by 3.1 percentage points from 3.6% to 6.7%. Short call verticals rise by 2.4 percentage points, from 3.1% to 5.5%. These two strategies alone account for a 5.5 percentage point increase in market share within the one-minute window, and both are net sales of option premium. The averages understate the effect at the most active times: the share of short put verticals rises by 7.8 percentage points at 13:30, and short call verticals

peak at 14:00 with 4.3 points, against 0.7 and 0.6 points at 10:00. By 15:30 both are close to zero, which is what entry schedules produce: no new short-premium position is opened half an hour before expiration. Market makers building inventory would not concentrate their trading at fixed clock times, so this imbalance is difficult to attribute to liquidity provision.

These patterns differ from recent evidence on retail option trading around volatility events. Eaton et al. (2026) find that retail traders are net buyers of short-dated options, and that brokerage outages reduce implied volatility for the contracts retail investors buy. de Silva, Smith, and So (2025) find that retail investors concentrate option purchases before earnings announcements and overpay for anticipated volatility. The traders we study sell volatility at fixed clock times rather than buying it before scheduled events. Their trading is timed by the clock, not by the arrival of information.

Participation also varies with volatility. Table 3 sorts the spike intervals into quintiles by realized volatility over the preceding thirty minutes and reports how often each strategy trades. Short-premium trading falls in the highest quintiles: the share of short put verticals declines from 8.10% in the lowest quintile to 4.69% in the highest. These strategies continue to trade when volatility is high but account for less of the trading at those times, which is what scheduled entries produce when discretionary trading expands around them.

5.3. The structure of sold spreads and iron condors

Table 4 examines the characteristics of the short vertical spreads and short iron condors traded at the spike minutes. If these trades express discretionary views, their strikes, widths, and sizes should vary across trades. If they come from templates, they should not.

The short strikes move closer to the underlying price at the spike minutes. For short call verticals the distance to spot falls by 0.6 points, from 23.2 to 22.6, and for short put verticals it falls by 1.6 points, from 28.7 to 27.2. Closer strikes raise the delta of the position, which is what a delta target produces. Spread widths increase at the same time: short call spreads widen by 13.5 points, from 28.8 to 42.4, and put spreads by 20.4 points, from 41.9 to 62.3. Wider spreads raise the maximum loss and hold the credit close to a target for a given delta. Net credit per unit of width falls by 5.9 percentage points for call spreads, from 17.2% to 11.3%, and by 4.2 points for put spreads, from 13.3% to 9.2%. Trade sizes fall: the average short call spread trades 1.3 fewer contracts, from 4.7 to 3.4, and the average dollar size falls by about $310, from $1,171.7 to $862.2.

These movements go together. A trader who fixes a delta target and a maximum loss per position must take wider spreads at closer strikes and fewer contracts, which is what the table shows. Strike placement, width, credit, and size are the parameters that platform templates ask the user to set, and they move together here in the way a template implies. Option trading by customers who trade through non-retail brokerages shows no comparable uniformity (Ni, Pan, and Poteshman, 2008).

Three features characterize the trades at the spikes: they occur at fixed clock times, they follow a small number of strategy templates, and they are small. Each is documented separately in the literature. Their joint occurrence is not. Discretionary traders do not concentrate execution at fixed clock times while holding to narrow templates, and market makers do not build short-volatility positions in retail-sized increments at fixed intervals. Automated retail execution produces all three: the strategy is specified in advance, the parameters come from a menu, and the order is sent when the clock reaches the scheduled time. The flow looks discretionary in each account and standardized in aggregate.

5.4. Profits by horizon

Table 5 reports profits at three horizons: five minutes after execution, thirty minutes after execution, and the end of the day, when the options expire.

Profits measured five and thirty minutes after execution are small and mostly negative. Averaging across days, short call verticals executed in the twelve spike windows lose $6,954 at the five-minute horizon, or $75.50 per complex trade. Short put verticals lose $2,541 and short iron condors $1,762, or $55.75 and $167.16 per complex trade, respectively. At thirty minutes the losses are smaller: short call verticals lose $859 or $4.19 per complex trade and short iron condors $1,742, or $237.03 per complex trade, while short put verticals earn $191, or $85.69 per complex trade.

Profits assuming positions held to expiration are greater. Short put verticals earn $35,535.5 per day in aggregate, or $392 per complex trade. Short call verticals earn $2,011 and short iron condors earn $3,089.8, or $71.46 and $345.71 per complex trade. These are the returns to selling premium and holding the position: the position pays as time decay accrues, not in the minutes after execution. The end-of-day averages are positive for all three strategies but are statistically indistinguishable from zero, with t-statistics of 0.1, 0.8, and 0.4, and they vary substantially across days. Thus, there is no evidence indicates that these strategies are systematically wealth-depleting.

However, the results do not establish that they are reliably profitable, as the positive average profits are not statistically different from zero. The strategies are defined-risk, so the maximum loss on each position is fixed when it is opened. The averages therefore do not hide large losses due to unhedged short-premium positions.

Profits vary across the day, though the imprecision of the estimates limits what can be inferred from the variation. Positions opened between 13:00 and 14:00 have higher average profits than those opened earlier, and iron condors opened between 13:00 and 13:30 have the highest.

The timing of these trades is set in advance. Trades concentrate at fixed clock times whatever the conditions at that moment, and the positions pay over the day rather than in the minutes after they are opened.

6. Identifying Scheduled Trades

Trades that arrive in the bursts can be identified without using the clock. A model told only what a trade looks like and what the market was doing beforehand ranks a burst trade above a quiet-window trade about two times in three. The bursts are not the surrounding trading arriving faster. They are different trades.

6.1. Trades we use in our classification

We use every vertical-spread trade in the sample, the most common complex strategy during the spikes. Trades in the first 60 seconds of each half-hour window form the burst sample: 778,694 trades, 7.0% of the total, in windows that make up 3.3% of the regular trading hours. Trades in minutes 25 through 29 form the reference sample: 1,192,268 trades, 10.7% of all trades. The five-minute marks documented in Section 4.2 fall inside the reference window, so any scheduled trading it contains makes the two samples more alike and works against finding differences. The remaining trades are set aside. In total we use 11,130,060 vertical-spread trades over 516 trading days from January 2024 through January 2026, dropping the first 25 and last 5 minutes of each day.

6.2. Differences between the samples

Table 6 compares the two samples. Trades arrive 2.6 times as fast in the bursts, 125.8 per minute against 48.1. The spreads are wider, 56.5 points between strikes against 36.3 (t-statistic = 30.1) during the reference window, and they collect less premium for that width: net credit falls from 14.5% to 8.8% (t-statistic = 52.6). The positions pay less and risk slightly more. Average maximum gain falls by $308.9, from $1,076.4 to $767.5, and average maximum loss rises to

$8,459.3. Greek exposures are small in both samples, which is what defined-risk positions produce, though burst trades carry somewhat more vega and more negative theta. Effective spreads are $0.0054 wider in the bursts, $0.0470 against $0.0416 ($t$-statistic = 4.8), and quoted spreads are $0.0072 tighter, $0.1073 against $0.1145 ($t$-statistic = 3.0). Both are significant at the 1% level, and Section 7 shows the same pattern strategy by strategy.

The market before the trades differs as well. Realized volatility over the preceding half hour is 2.5 percentage points lower before burst trades, 9.2% against 11.7% (t-statistic = 12.6), while implied volatility is the same in both windows, 12.01% against 11.99%. The scheduled trades arrive after calm half hours, and not because option prices signal calm: implied volatility does not move. It is the other trading that changes. Discretionary trading expands when realized volatility is high, which leaves the scheduled trades a smaller share of a busier tape. Table 3 shows the same pattern by volatility quintile.

6.3. Classification of individual trades

Averages across windows do not show that individual trades can be identified. We therefore train a model to sort trades into burst and reference windows (LightGBM; Ke et al., 2017), using 24 characteristics of the trade and of the market before it: spread width, net credit relative to width, maximum gain and loss, how far the short strike sits from the index, the quoted spread, the Greeks, implied volatility, the index return over the preceding half hour, and realized volatility over 5 to 30 minutes. We give the model no information about the time of day, so it cannot use the schedule we are asking it to detect. We train on the first 80% of the sample and hold out the last 20%, tune within the training period using five expanding windows, and test once on the held-out fifth.

Table 7 reports the results. The are under the ROC curve is 0.685, the precision-recall area is 53.1% against a base rate of 36.1%, and the Brier score is 0.218. At the default cutoff the model finds 61% of the burst trades and is right 52% of the time when it flags one. Performance is stable throughout: across the five expanding windows from 2024 to 2025 the ROC area averages 0.702, varies by 0.022, and never falls below 0.668. The model is picking up something about these trades that persists rather than a passing market condition.

Table 8 and Figure 6 show which characteristics the model uses, measured by gain and by SHAP values (Lundberg and Lee, 2017). Spread width matters most, accounting for 22.6% of the total, and the wider the spread the more likely the trade comes from a burst over most of the range

(Figure 6, Panel B). Discretionary traders choose widths across a wide range; platform templates offer a short list of them, and that is what separates the two. Realized volatility over the preceding half hour comes second, for the reason Table 6 shows: high volatility indicates that a trade is not from a burst. Vega and theta account for most of what remains, and they pick out the exposures that defined-risk templates create.

Burst trades therefore come from a different source, not only at a different time. The next section examines what this predictable trading does to the cost of trading.

7. Liquidity and Execution Costs

Liquidity suppliers can anticipate this flow. The schedule is evident in the stream of quotes and trades, and the trades that arrive on it can be identified from their characteristics. This section explores what that predictability does to the cost of trading. Quoted spreads in the single-leg LOB narrow at the spike times, effective spreads of the vertical spreads and iron condors widen, and realized spreads move in opposite directions for calls and puts. Thus, the answer depends on which measure one looks at and on which side of the trade one takes.

7.1. Quoted, effective, and realized spreads

Table 9 reports quoted, effective, and realized spreads for the three main strategies at the spike minutes and over the five minutes before.

There are no resting quotes in the complex order book, and thus no quoted spreads for complex strategies. Instead, for each strategy we computed a synthetic quoted spread from quotes in the single-LOB; specifically, we compute the synthetic quoted spread as the sum of the half-spreads in the single-leg LOB for each of the legs of the complex trades. For example, if a vertical spread involves selling a put with strike K_2 and buying a put with strike K_1, we compute the quoted bid-ask half-spread of the vertical spread as the sum of the half-spreads of the two constituent puts. Thus, the synthetic quoted spread is the cost of the alternative strategy of executing the several legs by crossing the spreads in the single-leg LOB.

The synthetic quoted spreads we compute narrow at the spikes. For short call verticals the quoted spread falls by \$0.009, from \$0.104 to \$0.095. For short put verticals it falls by \$0.005, from \$0.107 to \$0.102, and for short iron condors by \$0.004, from \$0.171 to \$0.167. The narrowing occurs across strategies and is largest in the early afternoon, when the spikes are largest. This is what liquidity suppliers do when they know uninformed orders are coming, and here they know it

twice over: the schedule is evident in the stream of quotes and trades, and the trades that arrive on it can be identified from their characteristics (Section 6).

This is narrowing of quoted spreads consistent with theory, which predicts that predictable order flow attracts competing quotes (Glosten and Milgrom, 1985; Foucault, Kadan, and Kandel, 2005). While the synthetic quoted spreads are from the single-leg LOB, not the complex LOB, complex orders interact with the single-leg LOB when they are marketable against the single-leg LOB, which is more likely to happen when the single-leg LOBs have tighter spreads (Fu et al. 2026). This result is consistent with the hypothesis that spreads tighten in the single-leg LOBs as traders complete to interact with the complex orders.

Effective spreads widen over the same minutes. For short call verticals the effective spread rises by \$0.004, from \$0.039 to \$0.043; for short put verticals from \$0.039 to \$0.044; and for short iron condors from \$0.049 to \$0.055. This is not limited to the three strategies: across all vertical spreads the effective spread is \$0.0054 higher in the burst windows than in the reference windows (Table 6). The traders who send these orders pay more for immediacy at exactly the times when the quotes look best. Many similar orders arriving within seconds move prices, especially for multi-leg positions whose legs must trade together, and orders that arrive together are more likely to trade against stale quotes or to walk the book. Budish, Cramton, and Shim (2015) show that continuous limit order books produce costs of this kind with no asymmetric information. Quoted and effective spreads have been found to move apart in equity markets when order flow clusters (Biais, Hillion, and Spatt, 1995; Brogaard, Hendershott, and Riordan, 2014).

The averages hide differences across the day. Spike by spike, quoted spreads widen at 10:00 for all three strategies, by \$0.0187 for short call verticals (t-statistic = 6.1), \$0.0283 for short put verticals (t-statistic = 4.6), and \$0.0308 for short iron condors (t-statistic = 5.6), and narrow at most later spikes, most of all at 13:30, where they fall \$0.0226 for short call verticals (Internet Appendix Table IA2). Effective spreads widen at every spike, and widen most at 10:00. Liquidity suppliers quote into the scheduled orders only later in the day; at the first spike of the day both spreads widen. The 14:00 window is also unusual, with quoted spreads widening for short iron condors. FOMC statements are released at 14:00 on meeting days and may account for this; we treat that window separately.

Realized spreads show who bears the cost. One minute after the trade, the spike and baseline windows are indistinguishable for all three strategies. Ten minutes after, they move in

opposite directions for calls and puts. For short call verticals the realized spread rises by \$0.031, from −\$0.014 to \$0.017, and for short put verticals it falls by \$0.034, from \$0.016 to −\$0.018; both changes are significant at the 1% level. Short iron condors combine the two exposures and show no change, \$0.013 against \$0.016. What liquidity suppliers earn therefore depends on which side of the strategy they take. Short-premium orders leave them holding positions that must be hedged, the hedge differs for calls and puts, particularly near the money, and realized spreads compensate them for holding it (Stoll, 1978; Ho and Stoll, 1981).

7.2. Interpretation

Quoted spreads narrow at the spike times, which suggests that liquidity compete in the single-leg LOB to interact with the complex trades, but effective spreads widen and realized spreads depend on the strategy. A market can show narrow quotes and still be costly for traders who need to trade immediately (Budish, Cramton, and Shim, 2015). The costs are small relative to the positions themselves: an effective spread \$0.004 to \$0.006 wider is a fraction of the credit collected on a spread held to expiration (Section 5.4).

The cost is not adverse selection. We find no evidence that these traders know where prices are going. Orders sent on a schedule cannot respond to news that arrives after the schedule was set, and the trades carry the marks of a template rather than of a view: standard widths, mechanical credit-to-width ratios, and arrival in calm markets (Table 6). The cost arises because a schedule sends many orders into the book at the same instant, which is a feature of the market's design rather than of what one side knows and the other does not.

The results also have a lesson for measurement. Studying retail options trading without account-level data has required exchange flags, such as the codes that identify wholesaler executions (Bryzgalova, Pavlova, and Sikorskaya, 2023; Ernst and Spatt, 2026). Section 6 shows that this flow identifies itself: it arrives on a schedule and it looks the same every time, so it can be found in public data. This raises questions about execution quality and internalization in markets where no flags exist (O'Donovan, Yu, and Zhang, 2023).

What the model identifies is scheduled automated trading, which is not the same as retail trading, and an institutional algorithm could produce similar arrival patterns. We read these trades as retail because they follow the defaults and constraints of retail platforms rather than the practices of institutional execution. Their timing is also fixed before the conditions they trade into are realized, so the flow is plausibly unrelated to short-horizon information, which is the variation

needed to estimate price impact and the supply of liquidity (Hasbrouck, 1991). This trading is automated and precisely timed, yet it takes liquidity rather than supplying it and does not compete on speed. It is algorithmic but not high-frequency.

8. Conclusion

Independently operating algorithms can generate synchronized order flow without learning, communication, or strategic interaction. We document this in SPX 0DTE options, where trading arrives at fixed clock times, follows a small number of strategy templates, comes in small size, and is predictable in both its timing and its form. The bursts arrive within seconds of the mark and are over within a minute. They have grown each year and have spread to more clock times, and they are concentrated in complex, ultra-short-maturity strategies; simple trades and longer-dated options show little of the same pattern. This is not the discretionary, heterogeneous trading that the literature on retail investors describes.

What trades at these times also differs. Trading shifts toward small, complex, short-premium positions with standardized strikes, widths, and sizes. Strike placement, spread width, net credit, and contract size move together in the way a delta target and a maximum-loss cap imply, and the trades can be identified from these characteristics alone, without using the clock. Many independent traders are running similar templates on similar schedules, supplied by the same platforms.

The synchronization changes execution. Quoted spreads narrow at the spike times, which is what competing liquidity suppliers do when they can anticipate uninformed flow. Effective spreads widen over the same intervals, so the traders who generate the flow pay more for immediacy. Realized spreads move in opposite directions for call-side and put-side verticals, reflecting the hedging demands each imposes. These costs are not adverse selection. Orders sent on a schedule cannot respond to information arriving at the moment they execute; what makes them expensive is that a schedule sends many of them into the book at the same instant.

Algorithmic behavior in financial markets need not be studied only in simulation. Retail automation is built on published rules, so its schedules, its synchronization, and its consequences for market quality can be measured in public transaction data. This is a field counterpart to experimental studies of algorithmic trading (Colliard, Foucault, and Lovo, 2026; Dou, Goldstein, and Ji, 2025a, 2025b).

This is a structural change in options markets, and it is unlikely to remain confined to them. Fintech platforms do not simply reduce the cost of trading; they standardize decision-making. Once many investors draw their rules from the same templates and the same schedules, retail behavior in aggregate becomes synchronized, predictable, and large enough to move market quality, though it originates in many independent accounts. Studying markets in which retail trades through rules therefore requires attention not only to who the traders are but also to the platforms that determine how their trades are generated and executed**.**

References

Admati, Anat R., and Paul Pfleiderer. 1988. A Theory of Intraday Patterns: Volume and Price Variability. *Review of Financial Studies* 1(1): 3–40.

Admati, Anat R., and Paul Pfleiderer. 1991. Sunshine Trading and Financial Market Equilibrium. *Review of Financial Studies* 4 (3): 443–481.

Assad, Stephanie, Robert Clark, Daniel Ershov, and Lei Xu. 2024. Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market. *Journal of Political Economy* 132 (3): 723–771.

Barber, Brad M., and Terrance Odean. 2000. Trading Is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors. *Journal of Finance* 55 (2): 773–806.

Barber, Brad M., and Terrance Odean. 2008. All That Glitters: The Effect of Attention on the Buying Behavior of Individual and Institutional Investors. *Review of Financial Studies* 21 (2): 785–818.

Barber, Brad M., Yi-Tsung Lee, Yu-Jane Liu, and Terrance Odean. 2009. Just How Much Do Individual Investors Lose by Trading? *Review of Financial Studies* 22 (2): 609–632.

Barber, Brad M., and Terrance Odean. 2013. The Behavior of Individual Investors. In *Handbook of the Economics of Finance*, edited by George M. Constantinides, Milton Harris, and René M. Stulz, 2:1533–1570. Elsevier.

Barber, Brad M., Yi-Tsung Lee, and Terrance Odean. 2014. The Cross-Section of Speculator Skill: Evidence from Day Trading. *Journal of Financial Markets* 8, 1-24.

Barber, Brad M., Xing Huang, Terrance Odean, and Christopher Schwarz. 2022. Attention-Induced Trading and Returns: Evidence from Robinhood Users. *Journal of Finance* 77 (6): 3141–3190.

Bali, Turan G., Nusret Cakici, and Robert E. Whitelaw. 2011. Maxing Out: Stocks as Lotteries and the Cross-Section of Expected Returns. *Journal of Financial Economics* 99 (2): 427–446.

Biais, Bruno, Pierre Hillion, and Chester Spatt. 1995. An Empirical Analysis of the Limit Order Book and the Order Flow in the Paris Bourse. *Journal of Finance* 50 (5): 1655–1689.
Black, Fischer. 1986. Noise. *Journal of Finance* 41 (3): 529–543.

Blau BM, Bowles TB, Whitby RJ. 2016. Gambling Preferences, Options Markets, and Volatility. *Journal of Financial and Quantitative Analysis* 51(2): 515-540.

Boyer, Brian, and Keith Vorkink. 2014. Stock Options as Lotteries. *Journal of Finance* 69 (4): 1485–1527.

Boehmer, Ekkehart, Jones, Charles M., Zhang, Xiaoyan and Zhang, Xinran. 2021. Tracking Retail Investor Activity. *Journal of Finance*, 76: 2249–2305.

Bogousslavsky, Vincent, and Dmitriy Muravyev. 2024. The Anatomy of Retail Trading in Options. Working paper.

Brogaard, Jonathan, Terrence Hendershott, and Ryan Riordan. 2014. High-Frequency Trading and Price Discovery. *Review of Financial Studies* 27 (8): 2267–2306.

Brown, Zach Y., and Alexander MacKay. 2023. Competition in Pricing Algorithms. *American Economic Journal: Microeconomics* 15 (2): 109–156.

Bryzgalova, S., Pavlova, A., and Sikorskaya, T. 2023. Retail trading in options and the rise of the big three wholesalers. *Journal of Finance*, *78*(6), 3465-3514.

Budish, Eric, Peter Cramton, and John Shim. 2015. The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response. *Quarterly Journal of Economics* 130 (4): 1547–1621.

Calvano, Emilio, Giacomo Calzolari, Vincenzo Denicolò, and Sergio Pastorello. 2020. Artificial Intelligence, Algorithmic Pricing, and Collusion. *American Economic Review* 110 (10): 3267–3297.

Chan, Kalok, William Christie, and Paul Schultz. 1995. "Market Structure and the Intraday Pattern of Bid-Ask Spreads for NASDAQ Stocks." *The Journal of Business* 68 (1): 35–60.

Colliard, Jean-Edouard, Thierry Foucault, and Stefano Lovo. 2026. Algorithmic Pricing and Liquidity in Securities Markets. *Review of Financial Studies*, forthcoming.

Cookson, J. Anthony, and Marina Niessner. 2020. Why Don't We Agree? Evidence from a Social Network of Investors. *Journal of Finance* 75 (1): 173–228.

Coval, Joshua D., and Tyler Shumway. 2001. Expected Option Returns. *Journal of Finance* 56 (3): 983–1009.

Da, Zhi, Joseph Engelberg, and Pengjie Gao. 2011. In Search of Attention. *Journal of Finance* 66 (5): 1461–1499.

de Silva, Tim, Kevin Smith, and Eric C. So. 2025. Losing is optional: retail option trading and expected announcement volatility. *Review of Finance* 30 (2), 489–535.

Dou, Winston Wei, Itay Goldstein, and Yan Ji. 2025a. AI-Powered Trading, Algorithmic Collusion, and Price Efficiency. NBER Working Paper No. 34054.

Dou, Winston Wei, Itay Goldstein, and Yan Ji. 2025b. "Financial Market Fragility in the Era of AI Planning." Working paper, University of Pennsylvania.

Eaton, Gregory W., Green, T. Clifton, Roseman, Brian S., Wu, Yanbin. 2026. Retail option traders and the implied volatility surface. *Journal of Financial Economics* 117.

Ernst, Thomas, and Chester Spatt. 2026. Payment for Order Flow and Option Internalization. Review of Financial Studies, forthcoming.

Foster, Fred D., and S. Viswanathan. 1990. A Theory of the Intraday Variations in Volume, Variance, and Trading Costs in Securities Markets. *Review of Financial Studies* 3 (4): 593–624.

Foucault, Thierry, Oleg Kadan, and Eugene Kandel. 2005. Limit Order Book as a Market for Liquidity. *Review of Financial Studies* 18 (4): 1171–1217.

Fu, Lei, Su Li, David K. Musto, and Neil D. Pearson. 2026. Hope at a Reasonable Price: Customer Use of Limit Orders in the 0DTE Market. Working paper.

Galati, Luca, Terrence Hendershott, Saad Ali Khan, and Ryan Riordan, 2026, Not dead yet: Options trading floors, working paper.

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 (1): 71–100.

Hasbrouck, Joel. 1991. Measuring the Information Content of Stock Trades. Journal of Finance 46 (1): 179-207.

Hendershott, Terrence, Charles M. Jones, and Albert J. Menkveld. 2011. Does Algorithmic Trading Improve Liquidity? *Journal of Finance* 66 (1): 1–33.

Ho, Thomas S. Y., and Hans R. Stoll. 1981. Optimal Dealer Pricing under Transactions and Return Uncertainty. *Journal of Financial Economics* 9 (1): 47–73.

Hirshleifer, David. 2001. Investor Psychology and Asset Pricing. *Journal of Finance* 56 (4): 1533–1597.

Hu, Jianfeng, Antonia Kirilova, Gilbert Seongkyu Park, and Doojin Ryu. 2024. Who Profits from Trading Options? *Management Science* 70 (7): 4742–4761.

Kaniel, Ron, Gideon Saar, and Sheridan Titman. 2008. Individual Investor Trading and Stock Returns. *Journal of Finance* 63 (1): 273–310.

Ke, Guolin, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. 2017. LightGBM: A Highly Efficient Gradient Boosting Decision Tree. Advances in Neural Information Processing Systems 30: 3146–3154.

Kelley, Eric K., and Paul C. Tetlock. 2013. How Wise Are Crowds? Insights from Retail Orders and Stock Returns. *Journal of Finance* 68 (3): 1229–1265.

Kirilenko, Andrei, Albert S. Kyle, Mehrdad Samadi, and Tugkan Tuzun. 2017. The Flash Crash: High-Frequency Trading in an Electronic Market. *Journal of Finance* 72 (3): 967–998.

Kumar, Alok. 2009. Who Gambles in the Stock Market? *Journal of Finance* 64 (4): 1889–1933.

Kyle, Albert S. 1985. Continuous Auctions and Insider Trading. *Econometrica* 53 (6): 1315–1335.

Lakonishok, Josef, Inmoo Lee, Neil D. Pearson, and Allen M. Poteshman. 2007. Option Market Activity. *Review of Financial Studies* 20 (3): 813–857.

Lee, Charles M. C., and Mark J. Ready. 1991. Inferring Trade Direction from Intraday Data. *Journal of Finance* 46 (2): 733–746.

Li, Sophia, David K. Musto, and Neil D. Pearson, 2023, Costs of executing complex options trades, Wharton School, Working paper.

Lundberg, Scott M., and Su-In Lee. 2017. A Unified Approach to Interpreting Model Predictions. Advances in Neural Information Processing Systems 30: 4765-4774.

Muravyev, Dmitriy. Order flow and expected option returns. 2016. *Journal of Finance* 71 (2): 673-708.

Naranjo, Andy, Mahendrarajah Nimalendran, and Yanbin Wu. 2024. Betting on elusive returns: Retail trading in complex options. Working paper, University of Florida.

Newey, Whitney K., and Kenneth D. West. 1987. A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. Econometrica 55 (3): 703-708.

Ni, Shiu, Jun Pan, and Allen Poteshman. 2008. Volatility Information Trading in the Option Market. *Journal of Finance* 63 (3): 1059–1091.

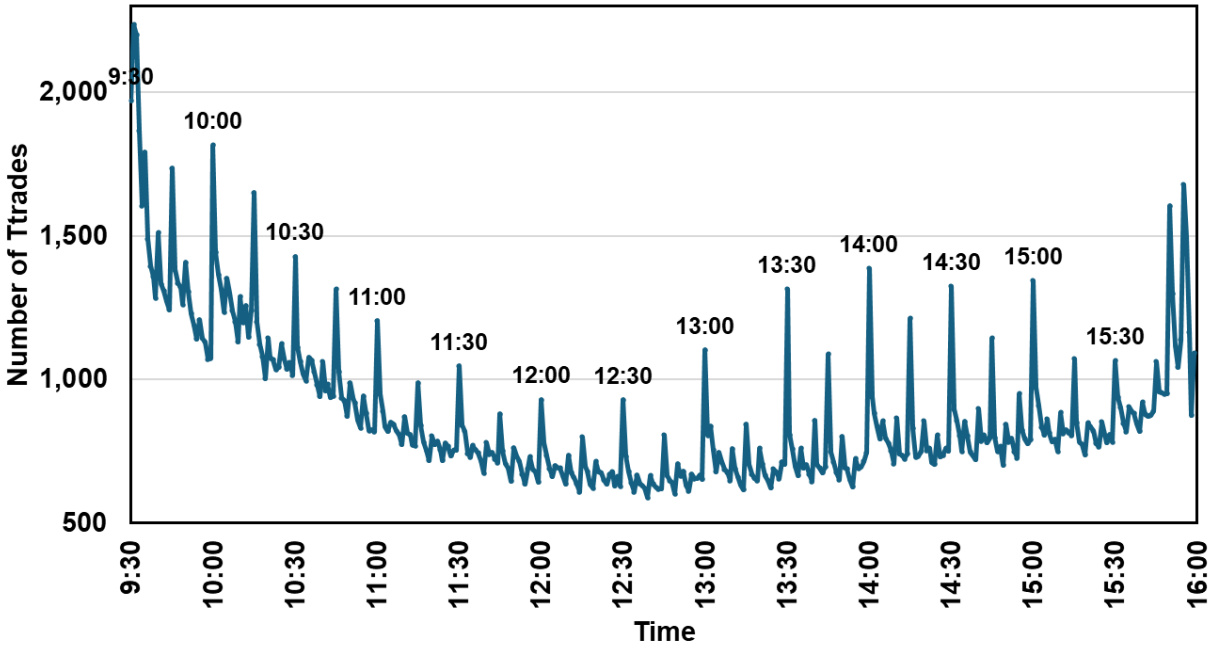
O'Donovan, James, Gloria Yang Yu, and Jinyuan Zhang. 2023. Retail Option Trading and Market Quality: Evidence from High-Frequency Data. Working paper.

Poteshman, Allen M., and Vitaly Serbin. 2003. Clearly Irrational Financial Market Behavior: Evidence from the Early Exercise of Exchange Traded Stock Options. *Journal of Finance* 58 (1): 37–70.

Seru, Amit, Tyler Shumway, and Noah Stoffman. 2010. Learning by Trading. *Review of Financial Studies* 23 (2): 705–739.

Stoll, Hans R. 1978. The Supply of Dealer Services in Securities Markets. *Journal of Finance* 33 (4): 1133–1151.

Panel A: Number of Trades



Panel B: Volume

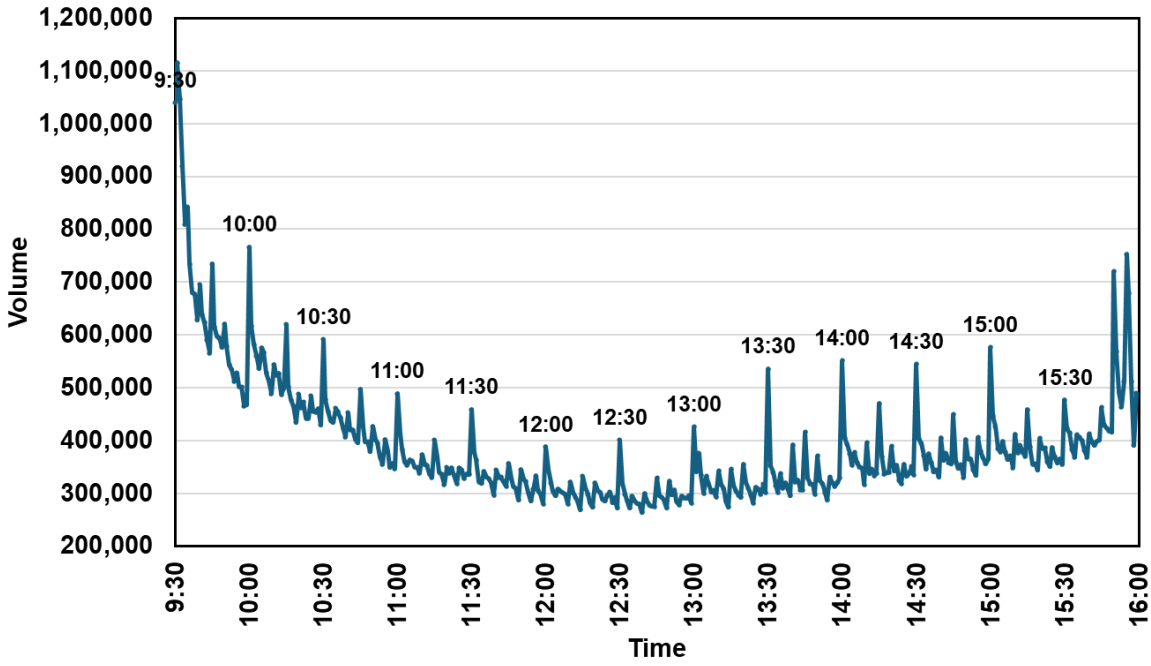


Figure 1
Intraday Patterns in SPX 0DTE Trading Volume and Trade Counts
This figure plots the average intraday trading activity for SPX 0DTE options. Panel A displays the average number of trades per minute, while Panel B displays the average trading volume (contracts) per minute. The x-axis represents the time of day from 9:30 to 16:00 EST. Labels indicate specific times (e.g., 10:00, 10:30, 13:00) where sharp spikes in activity are observed.

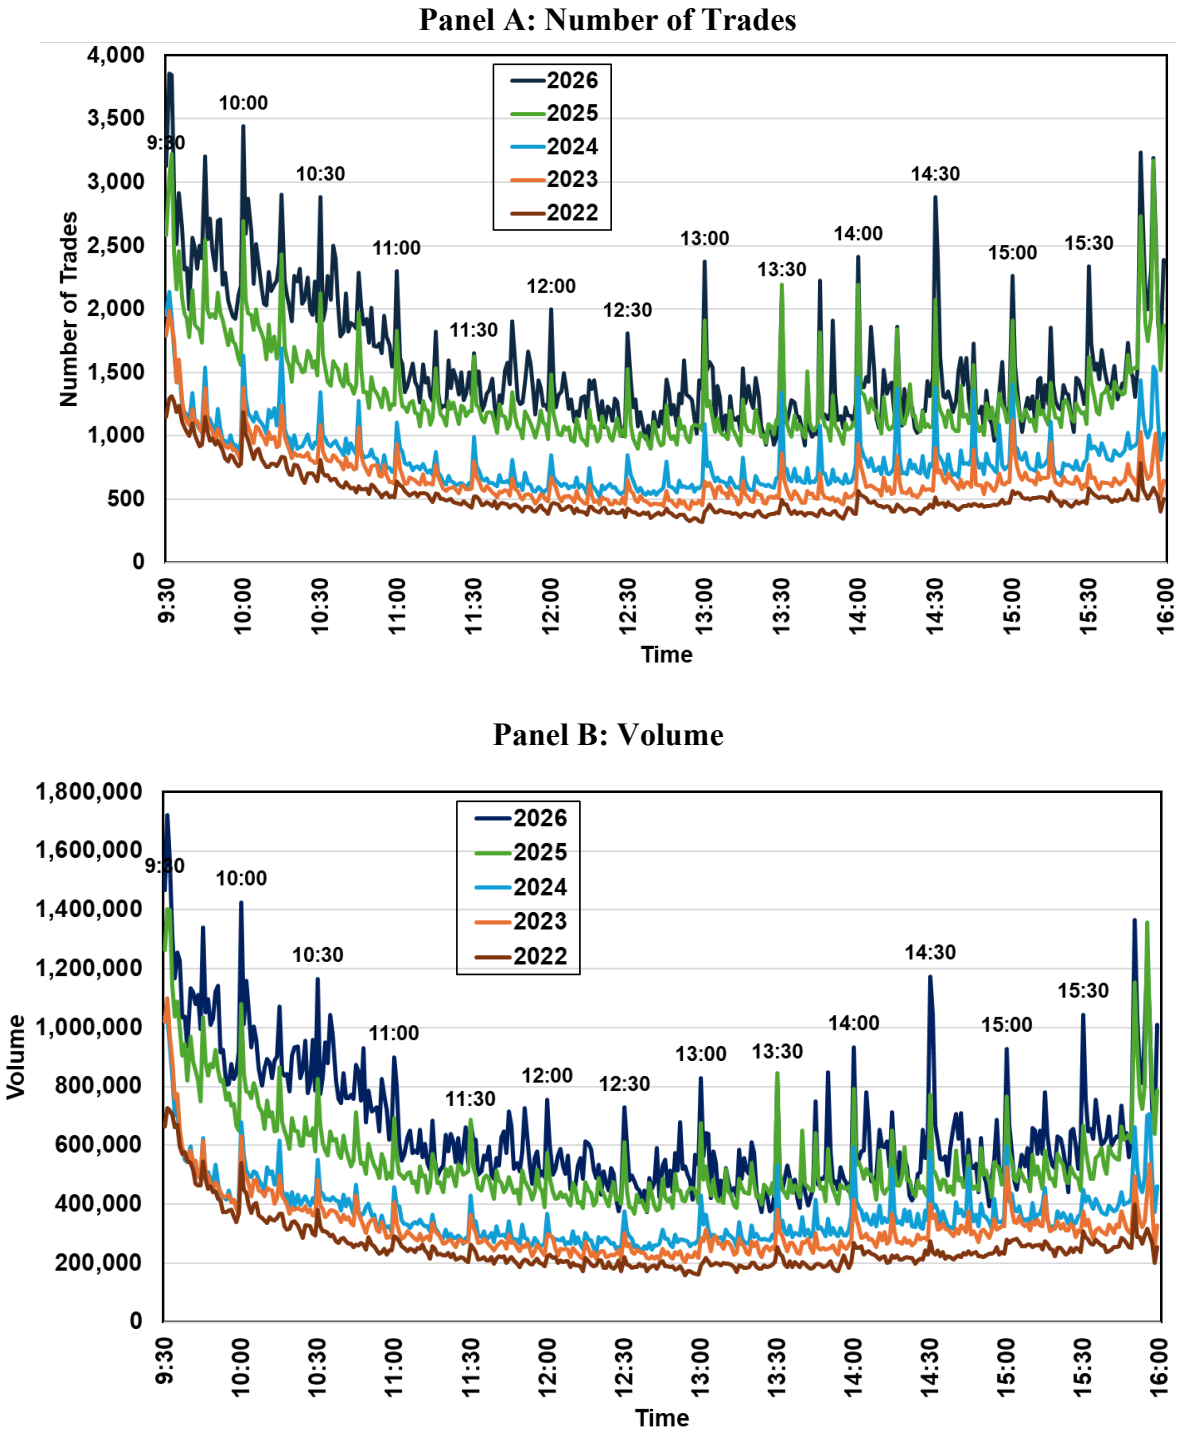


Figure 2
The Evolution of Intraday Trading Activity by Year (2022–2026)
This figure illustrates the evolution of intraday trading patterns over the sample period. Panel A plots the average number of trades per minute, and Panel B plots the average volume per minute, with separate lines for each year from 2022 to 2026.

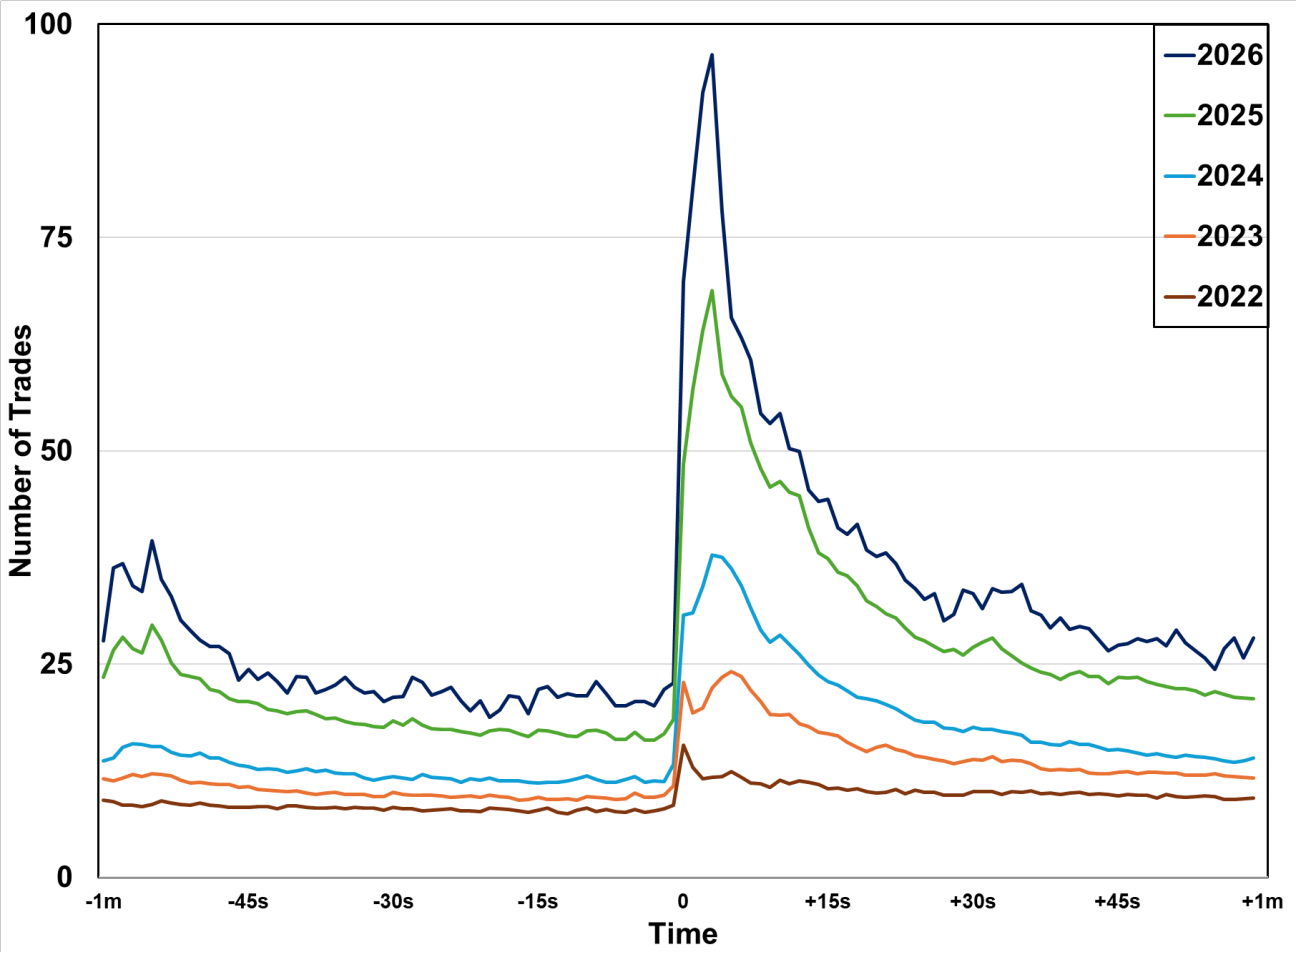
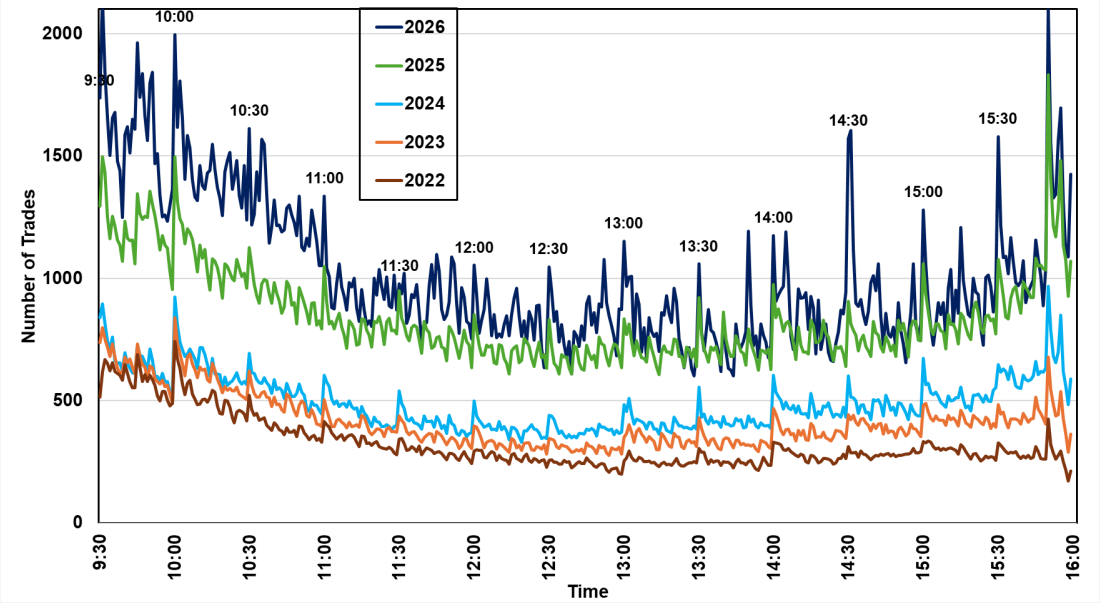


Figure 3
Second-Level Microstructure Dynamics of Intraday Trading Spikes
This figure details the microstructure of trading activity around the deterministic half-hour spikes aggregating option trades to second-level data. It plots the average number of trades per second for a window of 60 seconds centered on the spike time ($t = 0$). Separate lines represent different years (2022–2026).

Panel A: Simple Trades



Panel B: Complex Trades

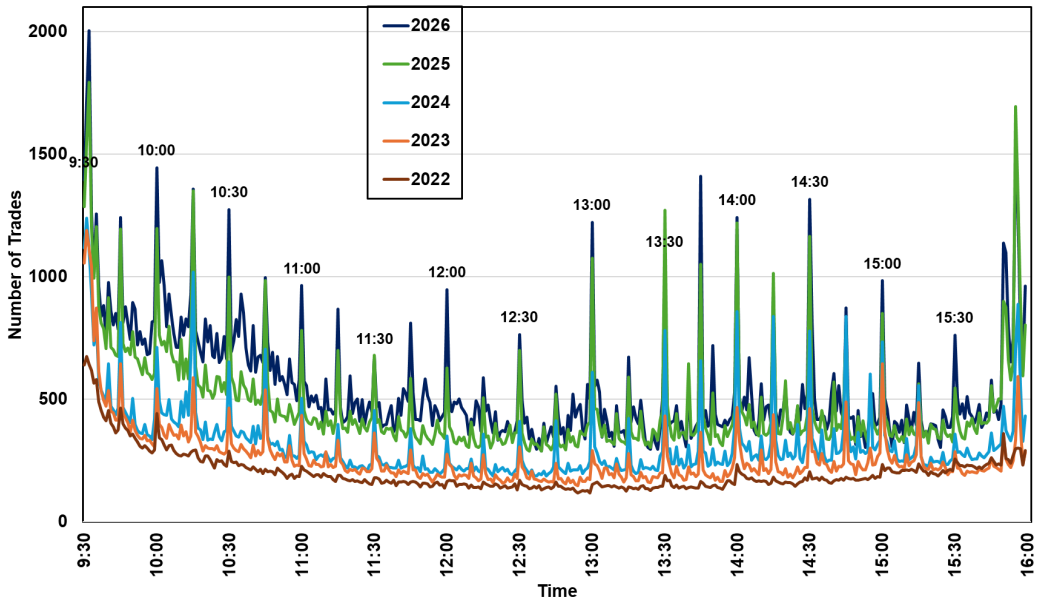


Figure 4
Intraday Trading Intensity: Simple versus Complex Option Strategies
This figure decomposes the average intraday number of trades into two categories: Simple Trades (Panel A) and Complex Trades (Panel B). Simple trades consist of single-leg transactions, while Complex trades involve multi-leg strategies. The figure compares the intraday patterns of these two groups across years (2022–2026).

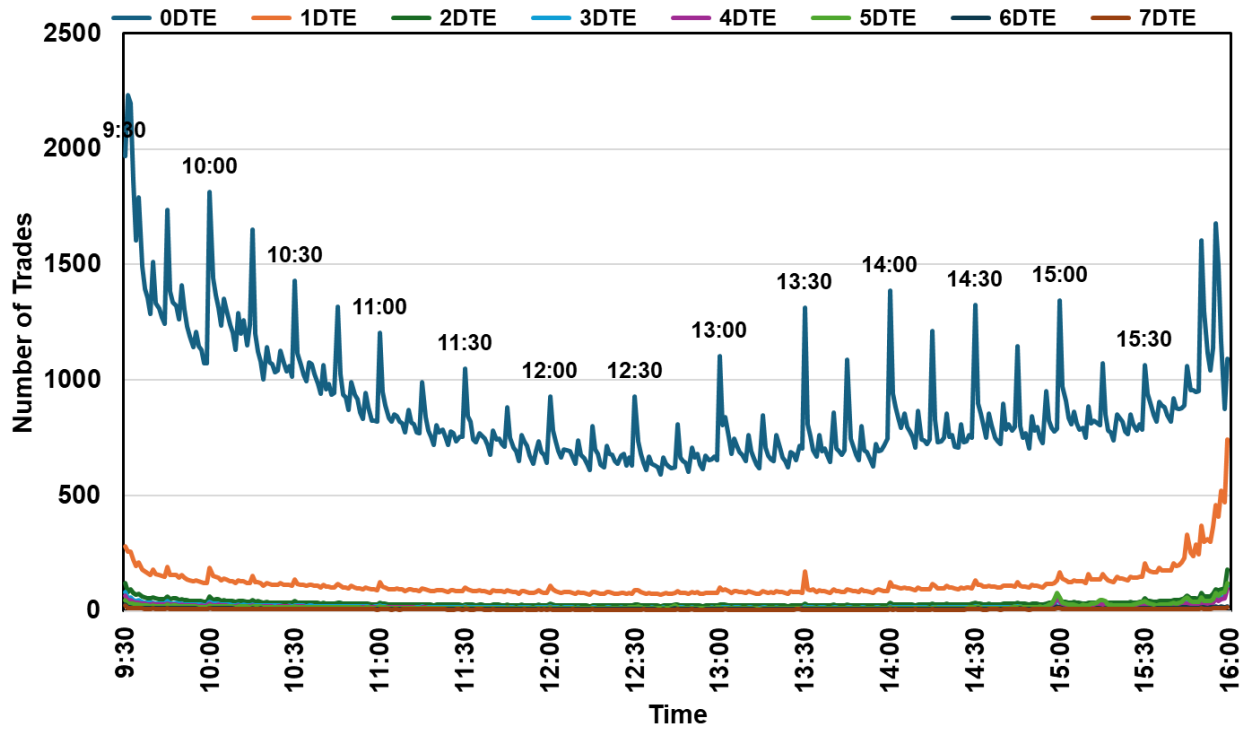
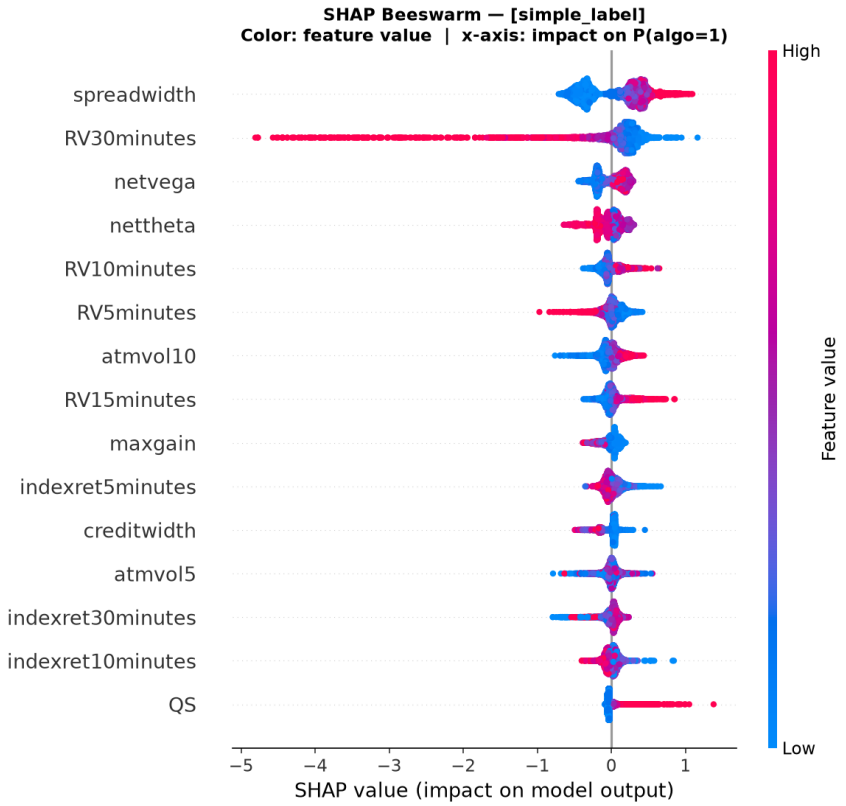


Figure 5
Intraday Trading Intensity across Option Maturities
This figure compares the average intraday number of trades across different option maturities, ranging from 0 days to expiration (0DTE) to 7 days to expiration (7DTE). The blue line represents 0DTE contracts, while other colors represent longer-dated options.

Panel A: SHAP Summary



Panel B: SHAP Dependence for the Three Leading Features

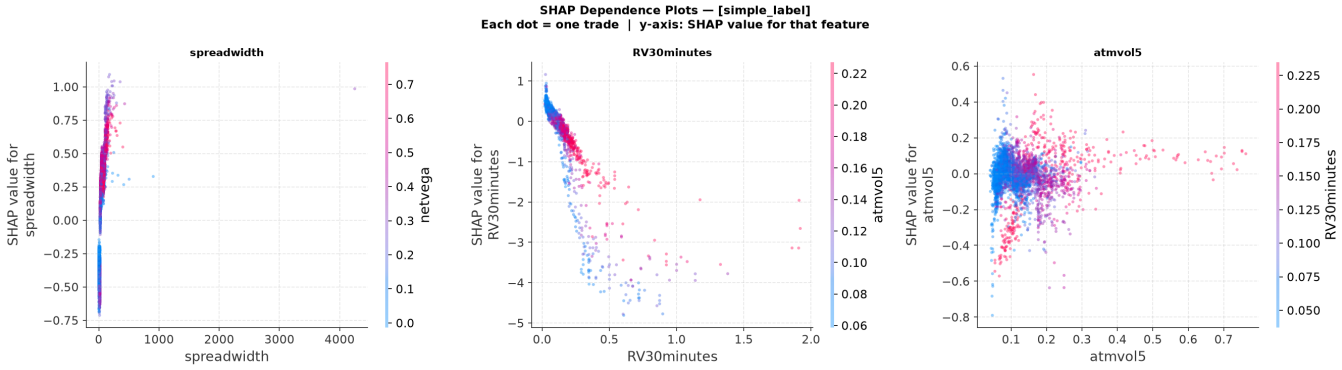


Figure 6

What Identifies Burst-Window Trades: SHAP Evidence

Panel A plots per-trade SHAP values (Lundberg and Lee, 2017) for the fifteen most important inputs of the classifier in Table 7. Each point is one trade; horizontal position is the feature's contribution to the predicted log-odds of burst-window membership, and color indicates the feature's value (red high, blue low). Panel B plots SHAP dependence for the three features with the highest gain. SHAP values are computed on a 5,000-trade sample. The panels display raw variable names; Table 8 reports the corresponding descriptive labels. The sample period runs from January 2024 through January 2026..

Table 1
Trade characteristics at the spike minutes and in the preceding five minutes
This table reports average characteristics of SPX 0DTE option trades in the spike minute of each half-hour window from 10:00 to 15:30 (Event, [0, +1m]) and in the five minutes before it (Pre-event, [−5m, 0)). Values are percentages of trades within each category. Type distinguishes calls from puts. Simple and Complex distinguish single-leg from multi-leg trades. Trade size is reported in contracts and in dollars. Moneyness is $(S - K)/K$ for calls and its negative for puts, where K is the strike and S is the index level. Diff. is the difference between the two windows. Asterisks indicate that the difference is significant at the 1% level, using Newey-West standard errors computed at the daily frequency with the optimal number of lags, on average eight days. The sample runs from January 2024 through January 2026.

		Event-time averages		
Characteristic	Category	Pre-event [−5m, 0)	Event [0, +1m]	Diff.
Type	Call	49.8	48.5	-1.3*
	Put	50.2	51.5	1.3*
Simple vs. complex	Complex	37.1	49.4	12.2*
	Simple	62.9	50.6	-12.2*
Trade size (contracts)	1	52.8	53.3	0.5*
	2 to 5	30.7	31.7	1.0*
	6 to 10	9.9	8.8	-1.1*
	11 to 20	4.0	3.7	-0.2*
	Above 20	2.7	2.5	-0.2*
Trade size (dollars)	Below 25	6.6	11.9	5.3*
	25-50	4.3	4.9	0.6*
	50-100	6.9	7.0	0.1
	100-150	6.3	6.6	0.3*
	150-250	10.4	11.2	0.7*
	250-500	15.9	15.3	-0.6*
	500-1,000	16.4	14.9	-1.5*
	1,000-2,500	17.0	14.8	-2.2*
	Above 2,500	16.3	13.4	-2.8*
Moneyness	Below -0.01	11.7	17.2	5.5*
	-0.01 to -0.005	17.8	18.7	0.9*
	-0.005 to 0.0025	67.2	61.5	-5.7*
	Above 0.0025	3.4	2.7	-0.7*

Table 2
Strategies and trade direction at the spike minutes
This table reports the share of trading accounted for by each strategy in the spike minute of each half-hour window ([0, +1m]) and in the five minutes before it ([-5m, 0)). Simple strategies are single-leg long and short calls and puts. Complex strategies are vertical spreads, iron condors, and other multi-leg positions. Values are percentages of all trades in the window. Diff. is the difference in percentage points between the two windows. Asterisks indicate that the difference is significant at the 1% level, using Newey-West standard errors computed at the daily frequency with eight lags. The sample runs from January 2024 through January 2026.

		Event-time averages		
		Pre-event	Event	
Type	Strategy	[−5m, 0]	[0, +1m]	Diff.
Simple	All	80.3	71.6	-8.7*
	Long Call	16.6	14.6	-2.0*
	Short Call	15.8	14.3	-1.6*
	Long Put	16.2	14.8	-1.4*
	Short Put	15.4	14.4	-1.0*
	Mid Call	8.2	6.7	-1.5*
	Mid Put	8.2	6.8	-1.3*
Complex	ALL	19.7	28.4	8.7*
	Long Call Vertical Spread	1.7	1.6	-0.1
	Short Call Vertical Spread	3.1	5.5	2.4*
	Long Put Vertical Spread	1.8	1.8	0.0
	Short Put Vertical Spread	3.6	6.7	3.1*
	Other 2 Leg Complex	3.1	3.4	0.3*
	3 Leg Complex	1.4	1.6	0.2*
	Short Iron Condor	0.6	1.2	0.6*
	Other 4 Leg Complex	2.4	3.7	1.3*
	Rest	2.0	2.9	0.9*

Table 3
Spike-time strategies by volatility quintile
This table reports the daily average share of short call vertical spreads, short put vertical spreads, and short iron condors traded in the spike minutes, with the spike intervals sorted into quintiles by realized volatility over the preceding 30 minutes. Low is the lowest volatility quintile and High the highest. HML is the difference between the highest and lowest quintiles. Newey-West t-statistics with eight lags are in parentheses. The sample runs from January 2024 through January 2026.

Strategy	Low	2	3	4	High	HML
Short call vertical spread	5.75%	5.82%	5.71%	5.51%	4.91%	-0.84%
	(34.56)	(31.24)	(30.93)	(32.91)	(29.49)	(-5.65)
Short put vertical spread	8.10%	7.67%	6.94%	6.00%	4.69%	-3.41%
	(35.12)	(36.90)	(30.17)	(34.51)	(34.43)	(-16.47)
Short iron condor	1.36%	1.28%	1.16%	1.10%	1.10%	-0.26%
	(28.84)	(21.17)	(20.37)	(18.52)	(19.07)	(-4.58)

Table 4

The structure of the dominant complex strategies at the spike minutes

This table reports characteristics of short call vertical spreads, short put vertical spreads, and short iron condors traded in the spike minute of each half-hour window ([0, +1m]) and in the five minutes before it ([-5m, 0)). Distance to spot is the absolute difference between the short strike and the index level, and spread width is the difference between the strikes, both in index points. Short-leg moneyness is the ratio of the index level to the short strike. Credit as % of width is the net premium divided by the spread width. Maximum loss and maximum gain are the dollar payoff bounds of the position. Diff. is the difference between the two windows. Asterisks indicate that the difference is significant at the 1% level, using Newey-West standard errors computed at the daily frequency with eight lags. The sample runs from January 2024 through January 2026.

		Event-time averages		
Strategy	Characteristic	Pre-event [-5m, 0]	Event [0, +1m]	Difference
Short call vertical spread	Number of trades	23.1	58.1	35.0*
	Trade Size (contracts)	4.7	3.4	-1.3*
	Net Price ($)	3.0	2.7	-0.3*
	Trade Size ($) Mean	1,171.7	862.2	-309.5*
	Trade Size ($) Median	345.7	287.8	-58.0*
	Notional ($ million)	2.8	2.0	-0.8*
	Distance to spot	23.2	22.6	-0.6*
	Spread width	28.8	42.4	13.5*
	Short leg moneyness	1.0	1.0	0.0*
	Delta of short leg	0.3	0.2	0.0*
	Credit as % of width	17.2	11.3	-5.9*
	Max loss	9,767.4	10,588.4	821.1*
	Max gain	1,171.7	862.2	-309.5*
Short put vertical spread	Number of trades	26.0	67.7	41.7*
	Trade Size (contracts)	4.7	3.4	-1.3*
	Net Price ($)	2.5	2.5	0.0
	Trade Size ($) Mean	980.7	764.7	-216.1*
	Trade Size ($) Median	314.6	294.8	-19.8*
	Notional	2.7	2.0	-0.7*
	Distance to spot	28.7	27.2	-1.6*
	Spread width	41.9	62.3	20.4*
	Short leg moneyness	1.0	1.0	0.0
	Delta of short leg	-0.3	-0.2	0.0*
	Credit as % of width	13.3	9.2	-4.2*
	Max loss	14,516.8	15,044.5	527.7*
	Max gain	980.7	764.7	-216.1*
Short iron condor	Number of trades	4.3	12.9	8.5*
	Trade Size (contracts)	4.2	3.0	-1.2*
	Net Price ($)	3.2	3.5	0.3*
	Trade Size ($) Mean	961.9	879.6	-82.3*
	Trade Size ($) Median	696.7	543.7	-153.1*
	Notional	2.5	1.7	-0.7*
	Distance to spot	20.6	21.5	0.9*
	Spread width	21.2	35.9	14.7*
	Short leg moneyness	1.0	1.0	0.0
	Delta of short leg	0.2	0.2	0.0
	Credit as % of width	20.9	15.5	-5.4*
	Max loss	6144.1	7388.8	1244.7*
	Max gain	961.9	879.6	-82.3*

Table 5
Profits by holding horizon
This table reports the daily average dollar profit on short call vertical spreads, short put vertical spreads, and short iron condors opened in the spike minute of each half-hour window ([0, +1m]) and in the five minutes before it ([−5m, 0)). Positions are marked to market five minutes after execution, thirty minutes after execution, and at expiration. Newey-West *t*-statistics with eight lags are in parentheses. The sample runs from January 2024 through January 2026.

		Event-time averages		
Strategy	Holding horizon	Pre-event [−5m, 0]	Event [0, +1m]	Diff.
Short call vertical spread				
	5 min	-990.7	-6,954.2	-5,963.5
		(-0.77)	(-2.54)	(-3.18)
	30 min	2,946.4	-859.2	-3,805.1
		(1.08)	(-0.17)	(-0.87)
	Expiration	3,042.3	2,010.8	-1,031.5
		(0.22)	(0.06)	(-0.05)
Short put vertical spread				
	5 min	-7,356.9	-2,540.9	4,816.0
		(-3.36)	(-0.73)	(2.09)
	30 min	-5,381.6	190.5	5,572.1
		(-1.22)	(0.03)	(1.17)
	Expiration	1,589.0	35,535.5	33,946.5
		(0.08)	(0.82)	(1.33)
Short iron condor				
	5 min	-2,515.4	-1,762.3	753.1
		(-3.04)	(-4.85)	(1.05)
	30 min	-1,423.6	-1,742.1	-318.5
		(-0.69)	(-1.30)	(-0.18)
	Expiration	1,740.6	3,089.8	1,349.2
		(0.70)	(0.44)	(0.27)

Table 6
Vertical spreads in burst windows and reference windows
This table compares vertical-spread trades executed in burst windows, the first 60 seconds of each half-hour window, with trades executed in reference windows, minutes 25 through 29 of the same window. Trades per minute of window time divides each day's trade count by the number of clock minutes the windows occupy, 12 and 48, respectively. Spread width and net premium are in index points, and credit as % of width is the net premium divided by the spread width. Short-leg moneyness is the ratio of the index level to the short strike. Quoted spread is half the sum of the two legs' bid-ask spreads, and effective spread is the difference between the executed net price and the midpoint-implied net price, both in dollars per spread. Realized volatility is annualized over the preceding thirty one-minute index returns. The columns report means of daily means, their difference, and the t-statistic of the daily difference series, computed with Newey-West standard errors with eight lags. The sample runs from January 2024 through January 2026.

	Reference windows	Burst windows	Diff.	t-statistic
Trades per minute of window time	48.1	125.8	77.6*	36.63
Spread width (points)	36.3	56.5	20.2*	30.05
Net premium (points)	2.69	2.52	-0.17*	-7.16
Credit as % of width	14.5	8.8	-5.7*	-52.56
Maximum gain ($)	1,076.4	767.5	-308.9*	-18.28
Maximum loss ($)	7,932.8	8,459.3	526.5*	4.45
Short leg moneyness	1.0008	1.0006	-0.0002*	-6.42
Net delta	-0.001	-0.011	-0.010*	-6.56
Net gamma	0.0084	0.0118	0.0034*	19.87
Net theta	-1.539	-1.978	-0.439*	-23.07
Net vega	0.264	0.380	0.116*	39.79
Quoted spread ($)	0.1145	0.1073	-0.0072*	-2.97
Effective spread ($)	0.0416	0.0470	0.0054*	4.82
ATM IV (%)	11.99	12.01	0.03	0.71
30-minute RV (%)	11.69	9.21	-2.48*	-12.57

Table 7
Classifying burst-window trades from their characteristics
This table reports the out-of-sample performance of a gradient-boosted tree classifier (LightGBM; 250 trees, learning rate 0.05, balanced class weights) trained to predict whether a vertical-spread trade falls in a burst window, using 24 characteristics of the trade and of the market before it and no information about the time of day. Burst and reference windows are defined as in Table 6. ROC-AUC is the probability that a randomly chosen burst-window trade is assigned a higher predicted probability than a randomly chosen reference-window trade. The hold-out rows evaluate one model on the final 20% of the sample under a chronological split; the walk-forward rows report the mean and standard deviation across five expanding-window folds inside the training period. Precision and recall refer to the burst class at the default cutoff. The sample runs from January 2024 through January 2026.

Hold-out evaluation (final 20%)	
ROC-AUC	0.6848
PR-AUC	0.5309
Brier score	0.2183
Precision, burst class	0.52
Recall, burst class	0.61
Walk-forward cross-validation (mean ± sd, five folds)	
ROC-AUC	0.7022 ± 0.0215
PR-AUC	0.5930 ± 0.0311
Brier score	0.2161 ± 0.0076
Sample	
Test observations	394,193
Burst share of test (%)	36.1

Table 8
What the classifier uses
This table reports the ten most important inputs to the classifier of Table 7. Gain is the percentage of total LightGBM split gain attributable to each characteristic. Mean |SHAP| is the mean absolute SHAP value (Lundberg and Lee, 2017), computed on a sample of 5,000 trades, which measures the average size of a characteristic's contribution to the predicted log-odds that a trade falls in a burst window. Each pair of columns is sorted by its own measure, and raw variable names appear in parentheses. Variable definitions follow Table 6, and the trailing-window variables are measured over the indicated number of preceding minutes. The sample runs from January 2024 through January 2026.

Rank	Feature (by gain)	Gain (%)	Feature (by |SHAP|)	Mean |SHAP|
1	Spread width (spreadwidth)	22.63	Spread width (spreadwidth)	0.371
2	Realized volatility (30 min) (RV 30minutes)	9.87	Realized volatility (30 min) (RV30minutes)	0.347
3	ATM implied volatility (5 min) (atmvol5)	5.63	Net vega (netvega)	0.161
4	Net vega (netvega)	5.61	Net theta (nettheta)	0.122
5	Realized volatility (10 min) (RV10minutes)	5.09	Realized volatility (10 min) (RV10minutes)	0.092
6	Realized volatility (15 min) (RV15minutes)	4.62	Realized volatility (5 min) (RV5minutes)	0.088
7	ATM implied volatility (10 min) (atmvol10)	4.30	ATM implied volatility (10 min) (atmvol10)	0.088
8	Index return (5 min) (indexret5minutes)	4.28	Realized volatility (15 min) (RV15minutes)	0.081
9	Net theta (nettheta)	4.19	Maximum gain (maxgain)	0.080
10	ATM implied volatility (30 min) (atmvol30)	4.05	Index return (5 min) (indexret5minutes)	0.076

Table 9
Quoted, effective, and realized spreads at the spike minutes
This table reports execution costs for short call vertical spreads, short put vertical spreads, and short iron condors in the spike minute of each half-hour window ([0, +1m]) and in the five minutes before it ([-5m, 0)). Quoted spread is half the difference between the national best offer and the national best bid at the time of the trade. Effective spread is the difference between the trade price and the prevailing NBBO midpoint, signed by trade direction, and realized spread 1' and 10' are the same differences against the midpoint one and ten minutes after the trade. All measures are in dollars. Diff. is the difference between the two windows. Asterisks indicate that the difference is significant at the 1% level, using Newey-West standard errors computed at the daily frequency with eight lags. The sample runs from January 2024 through January 2026.

		Event-time averages		
		Pre-event	Event	
Strategy		[−5m, 0]	[0, +1m]	Diff.
Short call vertical spread				
	Quoted Spread	0.104	0.095	-0.009*
	Effective Spread	0.039	0.043	0.004*
	Realized Spread 1'	0.015	0.028	0.014
	Realized Spread 10'	-0.014	0.017	0.031*
Short put vertical spread				
	Quoted Spread	0.107	0.102	-0.005*
	Effective Spread	0.039	0.044	0.004*
	Realized Spread 1'	0.017	0.009	-0.009
	Realized Spread 10'	0.016	-0.018	-0.034*
Short iron condor				
	Quoted Spread	0.171	0.167	-0.004*
	Effective Spread	0.049	0.055	0.006*
	Realized Spread 1'	0.031	0.035	0.003
	Realized Spread 10'	0.013	0.016	0.003

Internet Appendix to

The Rise of Algorithmic Retail Option Traders

Diego Amaya[1], Pedro A. Garcia-Ares[2], Neil D. Pearson[3,4], and Aurelio Vasquez[5]

August 12, 2026

[1] Lazaridis School of Business & Economics Business, Wilfrid Laurier University, 75 University Avenue West, Waterloo, Ontario, Canada, N2L 3C5. Tel.: +1 (519) 886-9351 ext. 4428. E-mail: damaya@wlu.ca.
[2] Instituto Tecnológico Autónomo de México (ITAM), Río Hondo No. 1, Col. Progreso Tizapán, México, D.F., México, 01080. Tel.: +52 (55) 5628 4000 ext. 6518. E-mail: pedro.garcia@itam.mx.
[3] Department of Finance, University of Illinois at Urbana-Champaign, 1206 South Sixth Street, Champaign, Illinois, 61820. Tel.: +1 (217) 244-0490. E-mail: pearson2@illinois.edu.
[4] Canadian Derivatives Institute, 3000, Chemin de la Côte-Sainte-Catherine, Montréal, Québec H3T 2A7, Canada.
[5] Instituto Tecnológico Autónomo de México (ITAM), Río Hondo No. 1, Col. Progreso Tizapán, México, D.F., México, 01080. Tel.: +52 (55) 5628 4000 ext. 6518. aurelio.vasquez@itam.mx.

Table IA1
Composition Changes at Each Spike Window

This table reports, for each half-hour spike window, the change in the listed share at the spike minute relative to the preceding five-minute window, in percentage points. The column headed "All" pools all twelve windows and corresponds to the differences reported in Tables 1 and 2. *t*-statistics computed from daily observations appear in parentheses. The sample comprises SPX 0DTE options traded on Cboe C1 from January 2024 through January 2026.

	All	**10:00**	**10:30**	**11:00**	**11:30**	**12:00**	**12:30**	**13:00**	**13:30**	**14:00**	**14:30**	**15:00**	**15:30**
Complex trade share	12.2	6.1	11.5	9.1	9.6	7.9	11.9	19.4	21.3	18.6	18.1	10.0	3.3
	(52.5)	(16.8)	(35.4)	(28.1)	(27.1)	(18.8)	(31.8)	(26.3)	(48.7)	(41.3)	(35.2)	(24.5)	(13.2)
Trades with premium below $25	5.3	2.2	4.4	3.7	4.0	3.3	5.6	7.5	9.6	8.1	8.3	5.0	2.3
	(37.7)	(17.5)	(26.3)	(17.4)	(18.0)	(21.5)	(22.1)	(24.0)	(32.9)	(36.1)	(34.3)	(23.0)	(17.3)
Short Put Vertical Spread share	3.1	0.7	2.0	2.0	2.3	1.5	3.7	4.8	7.8	5.3	5.3	1.5	0.0
	(29.1)	(6.5)	(13.2)	(16.0)	(12.6)	(11.2)	(20.8)	(15.0)	(22.7)	(22.8)	(21.7)	(9.8)	(0.2)
Short Call Vertical Spread share	2.4	0.6	3.2	2.0	2.0	1.4	1.6	4.1	3.6	4.3	3.0	2.9	0.2
	(30.8)	(7.0)	(25.0)	(16.6)	(15.6)	(14.9)	(11.2)	(18.4)	(21.2)	(20.5)	(17.9)	(16.1)	(3.3)
Short Iron Condor share	0.6	0.4	0.4	0.5	0.4	0.6	0.5	1.2	0.7	0.9	0.9	0.5	0.1
	(16.4)	(7.3)	(9.7)	(15.2)	(7.7)	(12.7)	(10.9)	(15.9)	(11.9)	(14.7)	(11.3)	(13.0)	(6.9)

Table IA2
Execution Quality at Each Spike Window

This table reports, for each half-hour spike window, the change in the quoted spread, the effective spread, and the ten-minute realized spread at the spike minute relative to the preceding five-minute window, in dollars, for the three dominant short-premium strategies. The All column pools all twelve windows and corresponds to Table 8. t-statistics computed from daily observations appear in parentheses. The sample comprises SPX 0DTE options traded on Cboe from January 2024 through January 2026.

	All	10:00	10:30	11:00	11:30	12:00	12:30	13:00	13:30	14:00	14:30	15:00	15:30
Short Call Vertical Spread													
Quoted spread	-0.0091	0.0187	-0.0096	-0.0102	-0.0132	-0.0119	-0.0155	-0.0175	-0.0226	-0.0001	-0.0131	-0.0084	-0.0055
	(-16.3)	(6.1)	(-8.7)	(-7.0)	(-10.7)	(-9.6)	(-9.7)	(-14.1)	(-7.5)	(0.0)	(-6.2)	(-4.0)	(-3.8)
Effective spread	0.0040	0.0092	0.0058	0.0030	0.0031	0.0031	0.0023	0.0038	0.0025	0.0073	0.0030	0.0036	0.0016
	(18.0)	(10.5)	(10.0)	(5.9)	(8.5)	(6.2)	(6.1)	(7.1)	(3.0)	(7.5)	(6.9)	(8.7)	(3.5)
Realized spread, 10 minutes	0.0306	-0.0100	0.0527	-0.0109	0.0173	0.0201	-0.0033	0.0480	0.0721	0.0383	0.0601	0.0589	0.0231
	(3.3)	(-0.3)	(1.9)	(-0.5)	(0.9)	(0.9)	(-0.2)	(2.0)	(2.8)	(1.2)	(2.2)	(2.1)	(0.8)
Short Put Vertical Spread													
Quoted spread	-0.0049	0.0283	-0.0105	-0.0036	-0.0046	-0.0135	-0.0131	-0.0103	-0.0207	0.0033	-0.0104	-0.0019	-0.0023
	(-5.8)	(4.6)	(-5.9)	(-1.6)	(-2.0)	(-6.7)	(-6.7)	(-4.2)	(-7.4)	(0.7)	(-5.5)	(-0.5)	(-1.5)
Effective spread	0.0042	0.0084	0.0022	0.0046	0.0044	0.0020	0.0047	0.0038	0.0030	0.0066	0.0039	0.0044	0.0021
	(16.1)	(9.4)	(6.6)	(6.4)	(8.1)	(5.2)	(10.1)	(5.6)	(6.8)	(6.7)	(6.3)	(6.7)	(4.1)
Realized spread, 10 minutes	-0.0336	-0.0054	-0.0506	-0.0412	-0.0425	-0.0373	0.0115	-0.0479	-0.0113	0.0033	-0.0444	-0.1221	-0.0145
	(-3.6)	(-0.2)	(-2.1)	(-1.5)	(-1.7)	(-1.5)	(0.5)	(-1.8)	(-0.4)	(0.1)	(-1.3)	(-3.9)	(-0.5)
Short Iron Condor													
Quoted spread	-0.0042	0.0308	-0.0112	-0.0117	-0.0177	-0.0048	-0.0129	-0.0079	-0.0137	0.0201	-0.0017	-0.0102	-0.0076
	(-4.2)	(5.6)	(-3.2)	(-4.8)	(-5.3)	(-1.1)	(-5.7)	(-3.5)	(-3.5)	(4.3)	(-0.6)	(-5.6)	(-2.0)
Effective spread	0.0064	0.0136	0.0040	0.0074	0.0052	0.0082	0.0050	0.0056	0.0027	0.0111	0.0040	0.0080	0.0031
	(19.9)	(12.7)	(5.3)	(7.5)	(7.3)	(10.3)	(6.4)	(6.8)	(3.4)	(8.9)	(4.6)	(9.5)	(3.0)
Realized spread, 10 minutes	0.0030	-0.0138	0.0013	-0.0171	0.0178	0.0021	-0.0079	-0.0209	0.0443	0.0241	0.0314	-0.0348	0.0166
	(0.5)	(-0.7)	(0.1)	(-1.2)	(1.4)	(0.2)	(-0.6)	(-1.4)	(2.0)	(1.2)	(1.3)	(-1.5)	(0.5)