The Evolution of Large U.S. Option Order Flow, 2021–2026

Nathaniel Peng (optionwhales.io)

April 11, 2026

Abstract

This paper studies the evolution of large U.S. option order flow over the five-year window from April 8, 2021 to April 7, 2026. The analysis draws on processed intraday option-flow data from OptionWhales (optionwhales.io), covering 1,253 trading sessions, roughly 5,509 unique underlying tickers, and 188,577,198 large orders with a minimum size threshold of 50 contracts. Three descriptive findings stand out. First, large-order activity became materially larger and broader: average daily large orders rose from 118,440 in the April–December 2021 sample slice to 206,205 in the January–April 2026 slice, while average daily premium increased from $6.3 billion to $19.2 billion and average daily active tickers increased from 1,619 to 2,039. Second, cluster-based evidence points to a gradual increase in inferred strategy complexity: the single-leg share of order clusters fell from 81.93% to 73.35%, while the average size of multi-leg clusters rose from 2.51 to 2.96 legs. Third, contract preferences shifted sharply toward shorter maturities: the 0DTE share of large-order activity rose from 12.6% to 36.0%, with longer-dated buckets losing share and at-the-money contracts remaining dominant. The evidence supports a descriptive interpretation of a larger, broader, shorter-dated, and more complex large-order options market, while highlighting important limitations associated with missing open/close flags and cluster-based strategy inference.

1 Introduction

The U.S. options market changed meaningfully between 2021 and 2026. Daily expirations became common in major index products, intraday macro trading through index-linked options intensified, and large option orders became a more visible part of market structure. These changes matter because options are not only directional instruments. They are also vehicles for hedging, volatility trading, yield generation, and tactical risk transfer. When the composition of option trading changes, the market's aggregate preferences over maturity, strike location, execution method, and strategy design change with it.

This paper focuses on a specific slice of that market: large option orders of 50 contracts or more. That threshold does not capture the full population of option activity, but it provides a useful window into behavior that is more likely to reflect institutional or otherwise sophisticated trading. Large-order flow is especially informative for studying how market structure evolves because it is where execution tactics, maturity selection, and multi-leg design are easiest to observe in aggregate.

We utilize data processed through OptionWhales (optionwhales.io), a platform designed to detect statistically abnormal option flow by aggregating OPRA trade data and reconstructing trade direc-

tion using NBBO-based methods. The platform provides normalized measures of directional option flow, which form the basis of our signal construction.

The paper addresses three empirical questions. First, how did the level and composition of large option activity change over the sample? Second, did cluster-based evidence point to greater use of multi-leg structures over time? Third, how did preferences shift across maturity, moneyness, trade size, and timing? The goal is descriptive rather than structural. The manuscript does not estimate a formal demand model. Instead, it documents how large-order flow reallocated across contract features and trading patterns over a five-year period.

Three findings organize the analysis. Large-order activity expanded in both scale and breadth. Inferred strategy complexity increased gradually, with single-leg clusters losing share and multi-leg clusters capturing a larger share of premium. Most strikingly, short-dated trading expanded sharply, especially in 0DTE contracts, while the market remained heavily concentrated near at-the-money strikes. These patterns are consistent with a market that became more tactical, more maturity-sensitive, and more reliant on defined-risk structures.

2 Data and Sample Construction

The empirical input is an intraday option-flow dataset from OptionWhales (optionwhales.io) containing processed large-order records and session-level summary statistics for U.S.-listed equities and exchange-traded funds. The central unit of observation is an individual large order with at least 50 contracts. Each record contains the underlying ticker, OCC-format option symbol, option type, strike, execution time, total contracts, premium, notional exposure, trade direction, execution type, and a set of option Greeks when available. The project also relies on precomputed daily aggregates and cluster summaries that condense order-level information into ticker-day and session-level trend statistics.

Throughout the paper, the directional classification and normalized flow measures are taken from the OptionWhales processing pipeline, which aggregates OPRA trade data and reconstructs trade direction relative to NBBO conditions.

The usable sample runs from April 8, 2021 through April 7, 2026 and contains 1,253 trading sessions. Coverage expands over the sample, with monthly average daily large orders increasing from roughly 92,000 in mid-2021 to roughly 170,000–210,000 by 2025 and early 2026. Across the full window, the processed dataset contains 188,577,198 large-order records and approximately 5,509 unique underlying tickers. The average number of tickers with at least one large order on a given day rises over time, reinforcing the view that the market became broader rather than simply more concentrated in a handful of names.

Tickers are classified into three broad segments for descriptive analysis. *Index proxies* include a short list of major index-linked products such as SPY, QQQ, IWM, DIA, and leveraged variants. *ETFs* are identified from instrument metadata and naming rules. Remaining common stocks are treated as *equities* and can be further grouped into broad sector and market-cap buckets. This classification is useful for presentation, but it is not perfect historically because the available metadata reflects a later snapshot of the ticker universe.

Two sample-construction details are important for interpretation. First, bid and ask quote fields

are sparse in 2021 and improve substantially over time, so aggressiveness statistics are treated as supplementary rather than headline results. Second, the annual summary tables presented below do not correspond to six full calendar years. The 2021 row covers April–December 2021, while the 2026 row covers January–April 2026. Those partial-year labels are made explicit throughout the paper to avoid overstating comparability.

3 Measurement and Classification Methodology

3.1 Order-level metrics

Time to expiration (DTE) is recovered from the OCC option ticker rather than from a month/day expiration field that does not include the year. Contracts are then assigned to six maturity buckets: 0DTE, 1DTE, 2–7DTE, 8–30DTE, 31–90DTE, and 91+DTE. This bucketing is central to the paper because the sample period coincides with the market-wide rise of daily expirations in major index-linked contracts.

Moneyness is measured from the strike-to-spot relationship and assigned to ordered categories spanning deep low, low, at-the-money, high, and deep high strikes. To reduce obvious data errors, contracts with extreme strike-to-spot ratios satisfying $|K/S - 1| > 5$ are excluded from moneyness summaries. Spot prices less than or equal to zero are also excluded from calculations that require a spot reference.

Premium is measured in dollar terms at the trade level, and summary tables report average daily premium in billions of dollars. The dataset also includes notional exposure and option Greeks. The trade-direction labels used in the descriptive summaries follow the OptionWhales NBBO-based reconstruction rather than direct exchange-provided open/close intent fields. For Greek-based summaries, zero values are excluded from median calculations to reduce contamination from missing or placeholder entries. When quote data are available, an aggressiveness flag is constructed from the relationship between execution price and contemporaneous bid/ask quotes; because quote coverage improves materially only after mid-2022, this metric is not used as a headline result.

3.2 Daily aggregation

The underlying research pipeline first converts raw session-level files into daily aggregate statistics. For each trading day and ticker, the processed aggregates summarize order counts, contract counts, premium, notional, directional mix, option type, execution type, DTE composition, moneyness composition, size percentiles, Greek medians, and intraday timing buckets. Those daily summaries are then aggregated to monthly or quarterly frequency for the figures and to annual slices for the summary tables. This procedure emphasizes persistent changes in composition rather than noise from individual sessions.

3.3 Cluster-based strategy inference

The paper does not observe explicit spread identifiers or open/close flags. Strategy evidence is therefore inferred from order clusters formed by grouping trades that share the same underlying ticker and the same timestamp rounded to the nearest second. Cluster size is interpreted as a first-pass proxy for strategy complexity. Single-leg clusters are the benchmark case. Size-two clusters are classified into common structures such as vertical spreads, straddles, strangles, calendar spreads, diagonals, and split fills using option type, strike, expiration, and direction patterns. Size-three clusters can be tagged as butterflies when strike spacing matches the expected arithmetic pattern. Size-four clusters can be tagged as iron condors or iron butterflies when the joint call-put structure matches those templates. Clusters with more legs or ambiguous patterns are retained as complex but not fully classified.

This method is best understood as a lower-bound proxy for multi-leg activity. Genuine strategies split across multiple seconds, routed through different brokers, or executed in staggered fashion will appear as separate trades and therefore bias complexity downward. The reverse problem is also possible: unrelated same-second trades in the same underlying may occasionally be grouped together. For those reasons, the paper uses restrained language such as *inferred strategy complexity* and *cluster-based evidence* rather than claiming direct observation of trader intent.

4 Results I: Trade Structure

Figure 1 shows that large-order activity increased substantially over the sample. The monthly average of daily large orders rises from roughly 92,000 near the start of the sample to a range of roughly 170,000–210,000 by 2025 and early 2026. Table 1 confirms the same pattern in annual slices. Average daily orders increase from 118,440 in April–December 2021 to 206,205 in January–April 2026. Average daily premium rises even faster, from $6.3 billion to $19.2 billion. This growth does not appear to be driven solely by a narrower set of names. Average daily active tickers increase from 1,619 to 2,039 over the same period.

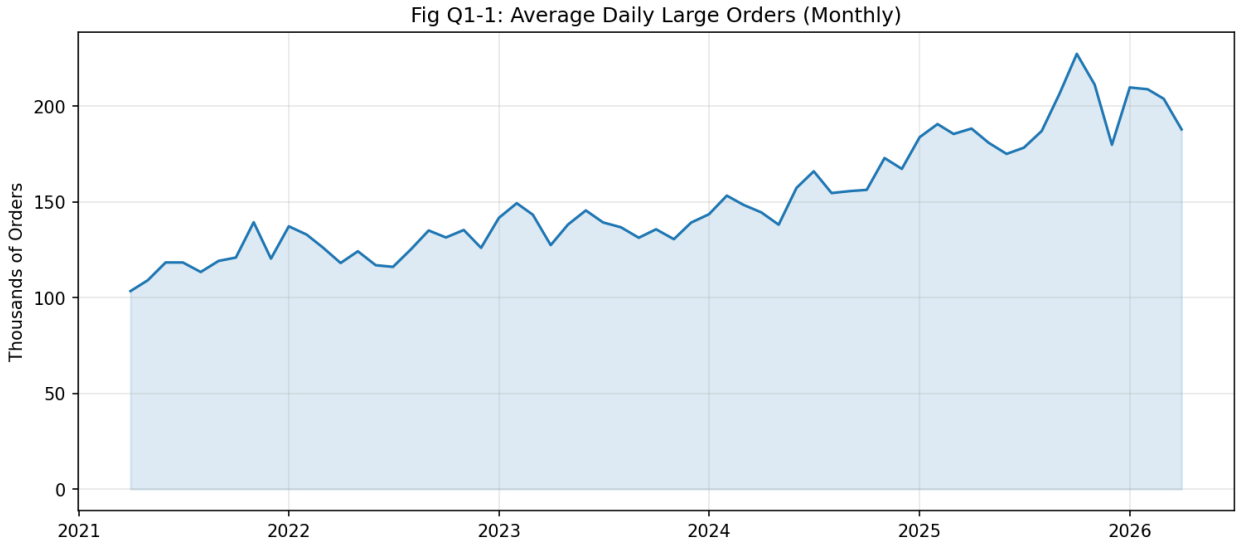

Figure 1: Average daily large orders by month. Large-order activity rises materially over the sample and remains elevated in the final two years.

Figure 2 highlights a second structural shift: the mix of execution methods changes over time. Sweep orders lose share while burst orders become the dominant execution type. Table 1 shows the same pattern in annual slices. Sweep share declines from 33.7% in the 2021 slice to 26.7% in the 2026 slice, whereas burst share rises from 38.5% to 46.2%. The result suggests that the growth in large-order activity is accompanied by changed execution tactics rather than simply more activity through an unchanged mechanism.

Figure 3 shows a third shift in composition: the balance of large orders moves toward the sell side. The sell share of records rises from 44.8% in the 2021 slice to 56.6% in the 2026 slice, while the buy share falls from 44.0% to 37.8%. This evidence is consistent with greater use of overwriting, volatility-selling, or hedge monetization strategies, but the data do not identify open versus closing trades. Accordingly, the result should be interpreted as a change in trade direction classification rather than a direct measurement of net speculative sentiment.

The trade-structure evidence therefore points to a market that became larger, broader, and operationally different. Growth in activity is accompanied by a more burst-heavy execution mix and a higher sell share, not just by a proportional scaling of the 2021 market structure. Supplemental evidence on active-ticker breadth and segment mix is reported in the appendix.

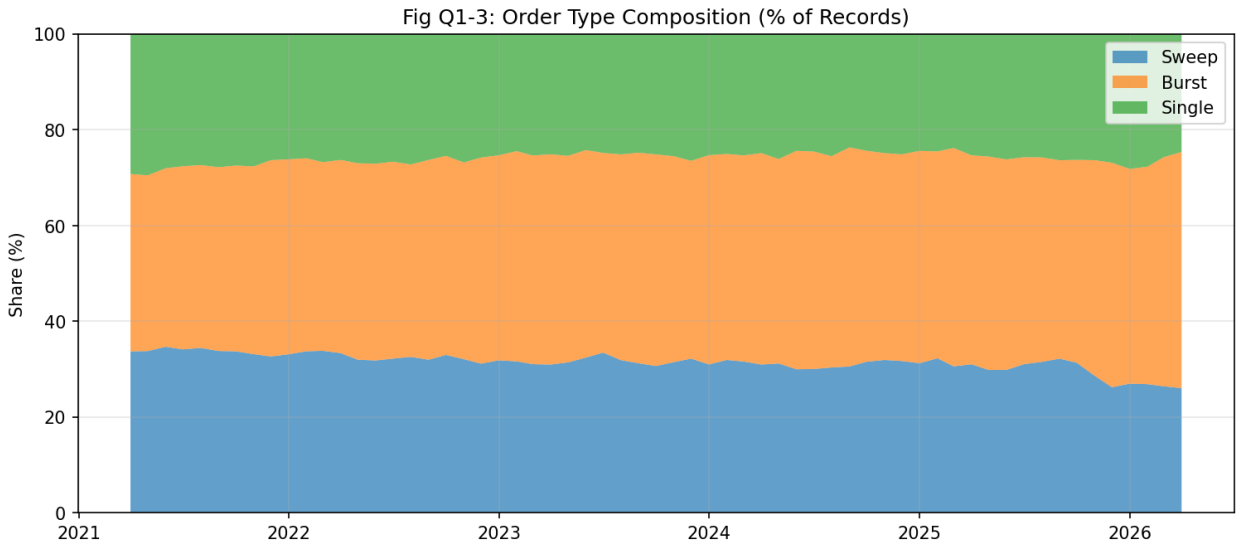

Figure 2: Order-type composition of large-order records by month. Burst orders steadily overtake sweeps, indicating a changed execution mix.

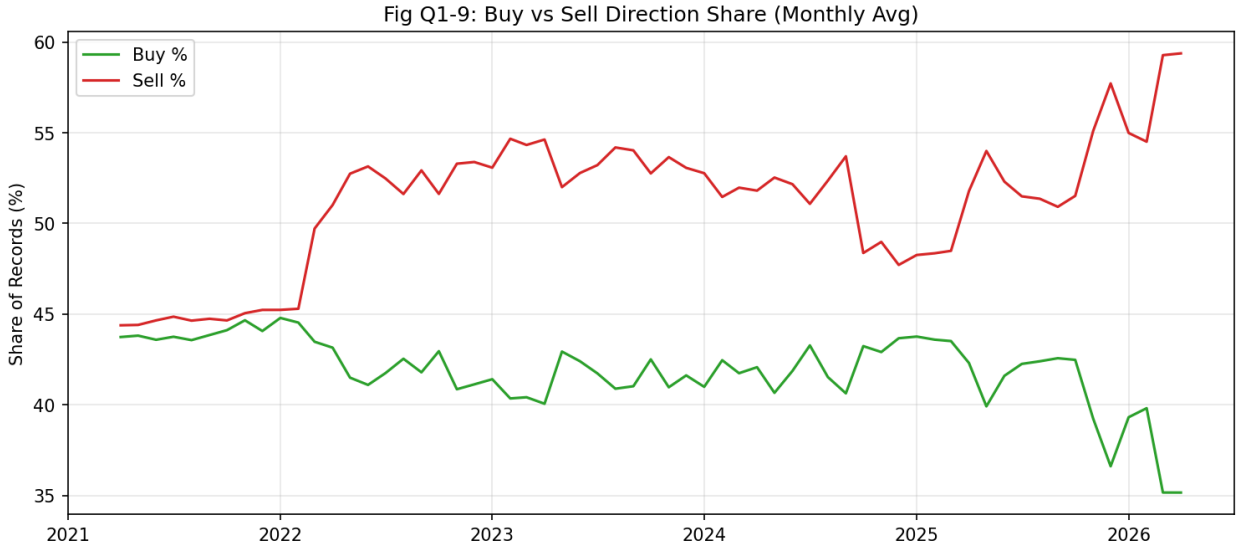

Figure 3: Buy- versus sell-initiated share of large-order records by month. The sell share trends upward over the sample, especially in the final year slice.

Table 1: Annual trade-structure summary. The first and last rows are partial-year slices.

Period	Avg. Daily Orders	Avg. Daily Premium ($B)	Avg. Daily Tickers	Buy Share	Sell Share	Sweep Share	Burst Share	ETF+Index Premium Share
2021 (Apr–Dec)	118,440	6.3	1,619.0	44.0	44.8	33.7	38.5	27.1
2022	127,072	9.5	1,568.3	42.6	51.0	32.6	41.0	32.8
2023	138,251	8.6	1,567.2	41.6	53.4	31.7	43.2	32.9
2024	154,753	8.1	1,663.7	42.2	51.1	31.1	44.0	29.0
2025	191,385	12.2	1,882.5	41.8	51.8	30.5	43.9	28.5
2026 (Jan–Apr)	206,205	19.2	2,039.0	37.8	56.6	26.7	46.2	26.2

Notes: premium is reported in billions of dollars. The ETF+Index Premium Share column combines index proxies and ETFs, while the residual share is attributable to equities.

5 Results II: Inferred Strategy Evolution

Figure 4 shows a gradual decline in the dominance of single-leg clusters. In the 2021 slice, 81.93% of observed clusters are single-leg, leaving 18.07% in multi-leg groupings. By the 2026 slice, the single-leg share falls to 73.35%, implying that the multi-leg share rises to 26.65%. The changes are not abrupt, but they are persistent: two-leg, three-leg, and four-plus-leg clusters all gain share over time.

Figure 5 shows that the premium footprint of multi-leg trading is even larger than the cluster-count evidence suggests. Multi-leg clusters account for roughly 60% of premium at the start of the sample and more commonly fall in the 70%–80% range later in the window, albeit with visible volatility and a temporary dip in early 2025. The contrast between cluster shares and premium shares indicates that complex clusters tend to carry more premium per event than simple single-leg trades.

Table 2 provides more detail on the composition of inferred strategy activity. Average cluster size rises from 2.51 legs in the 2021 slice to 2.96 in the 2026 slice. The supplementary size-two breakdown in Appendix Figure 10 shows that vertical spreads lose relative share, while strangles become more common, rising from 9.3% to 14.8% of two-leg classifications. These changes are consistent with a market that increasingly uses multi-leg structures to express views on volatility, timing, and payoff shape rather than relying exclusively on simple directional positions.

The strategy evidence should nonetheless be read as suggestive rather than definitive. Same-second clustering is informative at scale, but it does not recover true execution intent perfectly. The most defensible claim is that the processed order flow becomes more consistent with complex, multi-leg trading over time.

Table 2: Annual cluster-based strategy summary. The first and last rows are partial-year slices.

Period	Clusters (mn)	Single	Multi-Leg	Two-Leg	Three-Leg	Four+ Leg	Avg. Size
2021 (Apr–Dec)	17.43	81.93	18.07	13.25	2.65	2.17	2.51
2022	23.29	78.60	21.40	14.17	3.65	3.57	2.74
2023	23.91	76.45	23.55	14.83	4.20	4.52	2.86
2024	27.23	76.01	23.99	15.21	4.35	4.43	2.82
2025	33.18	75.57	24.43	15.39	4.40	4.65	2.83
2026 (Jan–Apr)	8.85	73.35	26.65	16.08	4.78	5.79	2.96

Notes: shares are percentages. Multi-leg share is computed as 100 minus the single-leg share. Strategy types are inferred from same-ticker, same-second clusters and should be interpreted as cluster-based proxies rather than exact observations of true trader intent.

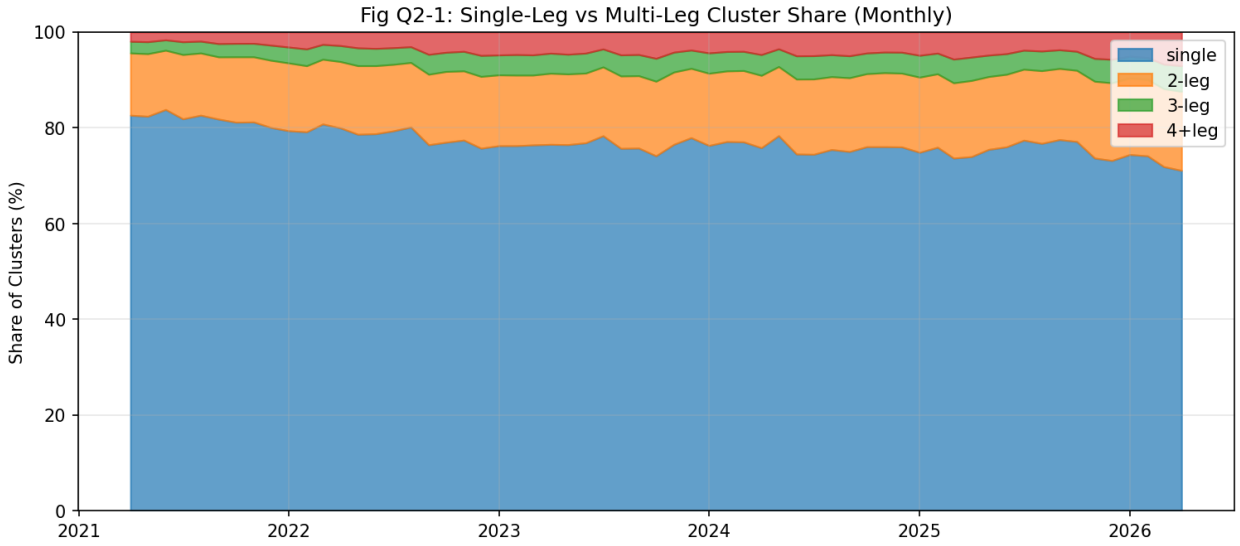

Figure 4: Single-leg versus multi-leg cluster share by month. Single-leg orders remain dominant, but multi-leg structures gain share steadily over time.

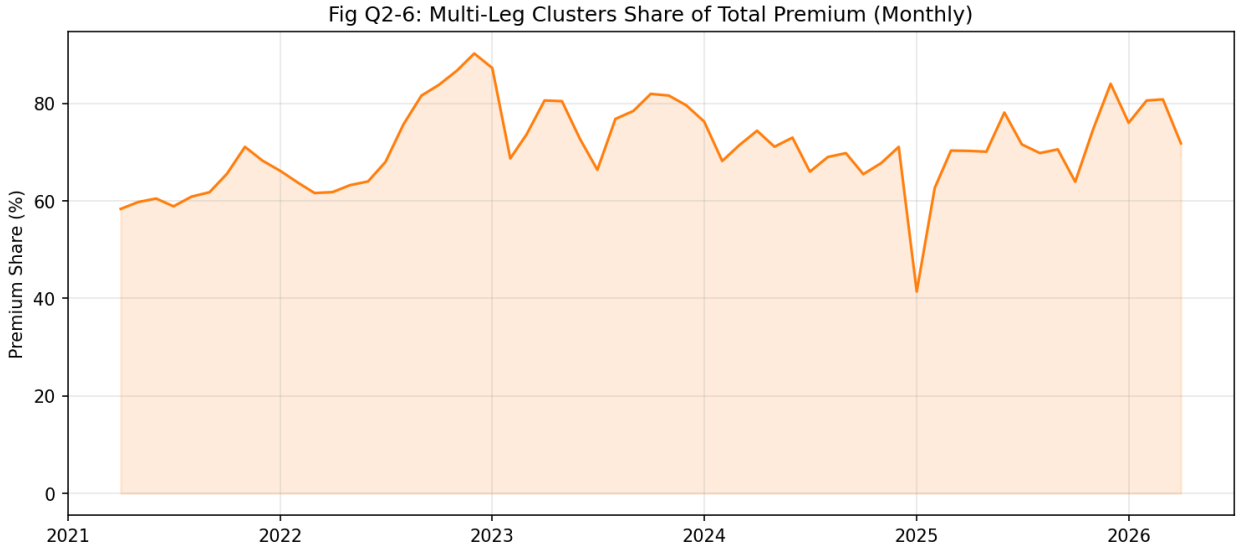

Figure 5: Share of total premium attributable to multi-leg clusters by month. Complex clusters account for a disproportionate share of premium relative to their share of events.

6 Results III: Preference Shifts Across Contract Features

Figure 6 presents the strongest headline result in the paper. Short-dated trading, and especially 0DTE trading, gains substantial share over the sample. In the April–December 2021 slice, 0DTE contracts account for 12.6% of large-order activity. By the January–April 2026 slice, that share reaches 36.0%. The rise in 0DTE is accompanied by broad declines in longer maturity buckets. The 2–7DTE share falls from 27.9% to 18.9%; the 8–30DTE share falls from 24.6% to 17.0%; and the 31–90DTE share falls from 14.4% to 9.7%. The implication is not merely that the market becomes shorter-dated in isolated episodes, but that shorter duration becomes a durable organizing feature of large-order flow.

Figure 7 shows how maturity and moneyness interact. At-the-money contracts dominate the joint distribution throughout the sample, but the center of gravity within that at-the-money region shifts toward shorter maturities over time. In 2022, the strongest demand cells are at-the-money 2–7DTE and 8–30DTE contracts. By 2026, the darkest cell is at-the-money 0DTE, with at-the-money 2–7DTE and 8–30DTE still important but less dominant. This is a compact way to summarize the contract-choice change documented in the maturity tables: the market remains heavily focused on near-the-money exposure, but increasingly expresses that exposure through very short duration.

Figure 8 adds an intraday dimension. The 10–11 a.m. Eastern hour is the most active interval across most quarters, while 15–16 ET is the second main concentration point. This pattern is stable enough to suggest that large orders tend to cluster during early-session price discovery and again into the closing hour, rather than being evenly distributed across the day. Although this timing pattern is secondary to the maturity result, it reinforces the interpretation of the modern large-order options market as tactical and intraday-sensitive.

Table 3 adds three secondary results. First, call share remains above 50% throughout the sample, but it is not monotonic enough to support a simple one-directional narrative; the call-versus-put mix is better described as volatile than steadily trending. Second, trade size drifts upward only gradually: median trade size increases from 85 to 95 contracts and the 99th percentile rises from 193 to 209 contracts. Third, median implied volatility rises from 0.416 to 0.533 across the first and last sample slices. Taken together, these results point to a market that became shorter-dated much more dramatically than it became larger in per-trade size.

Table 3: Annual contract-preference summary. The first and last rows are partial-year slices.

Period	Call	0DTE	1DTE	2–7DTE	8–30DTE	31–90DTE	91+DTE	Median Size	Median IV
2021 (Apr–Dec)	63.8	12.6	10.5	27.9	24.6	14.4	10.1	85	0.416
2022	55.1	17.3	12.1	27.2	21.7	12.3	9.3	90	0.490
2023	55.8	24.2	11.3	24.6	19.8	11.4	8.8	93	0.404
2024	59.2	25.4	10.8	24.7	19.4	10.9	8.7	95	0.407
2025	60.2	28.8	9.5	22.9	18.6	10.7	9.4	95	0.477
2026 (Jan–Apr)	57.2	36.0	9.5	18.9	17.0	9.7	9.0	95	0.533

Notes: shares are percentages. Median IV is annualized implied volatility. The omitted secondary metrics reported in the underlying summary files include the 99th-percentile trade size, median delta, and ETF+Index premium share.

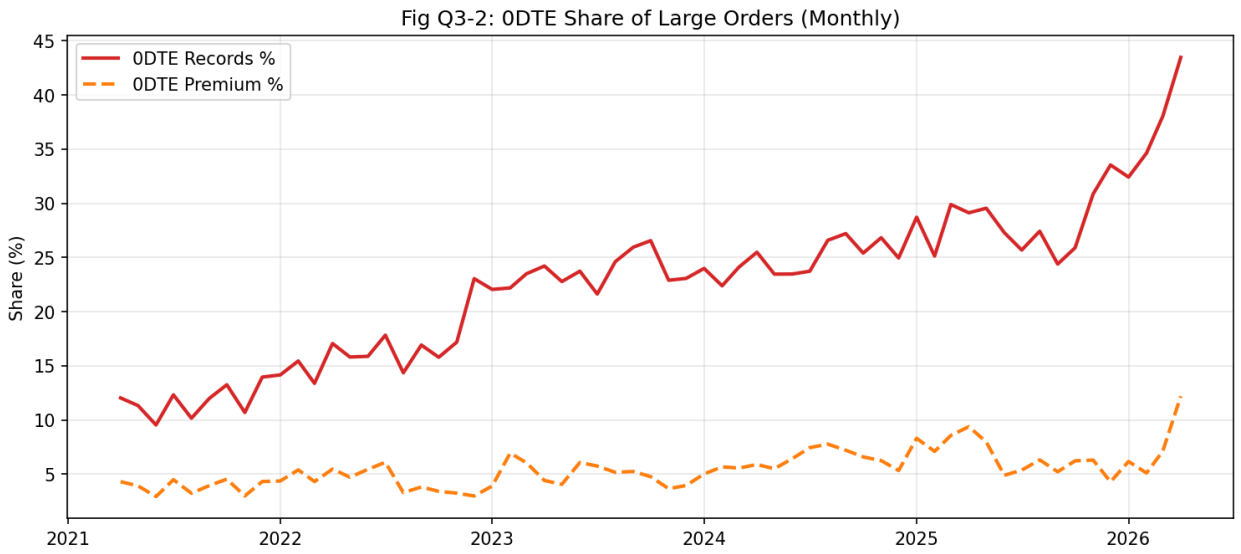

Figure 6: 0DTE share of large-order flow by month. Same-day expiration contracts move from a modest share early in the sample to a central feature of the market by 2026.

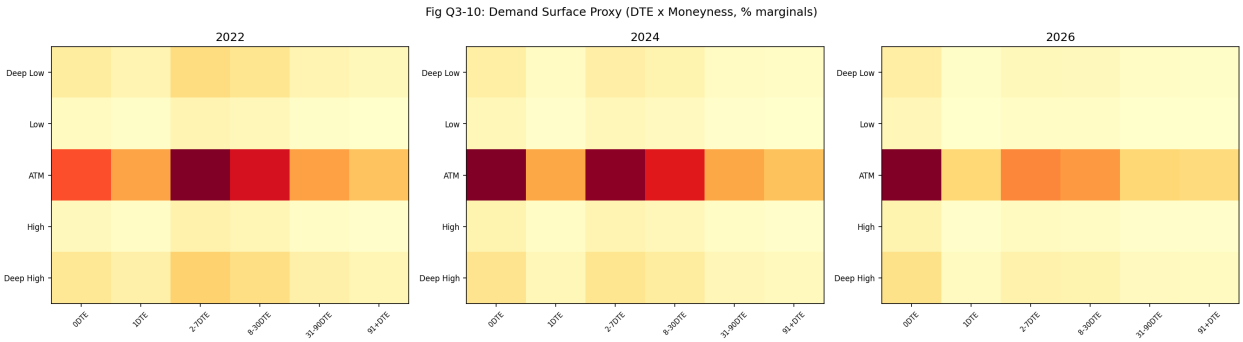

Figure 7: Joint distribution of DTE and moneyness for selected years. At-the-money contracts remain dominant, but the mass shifts toward shorter maturities by 2026.

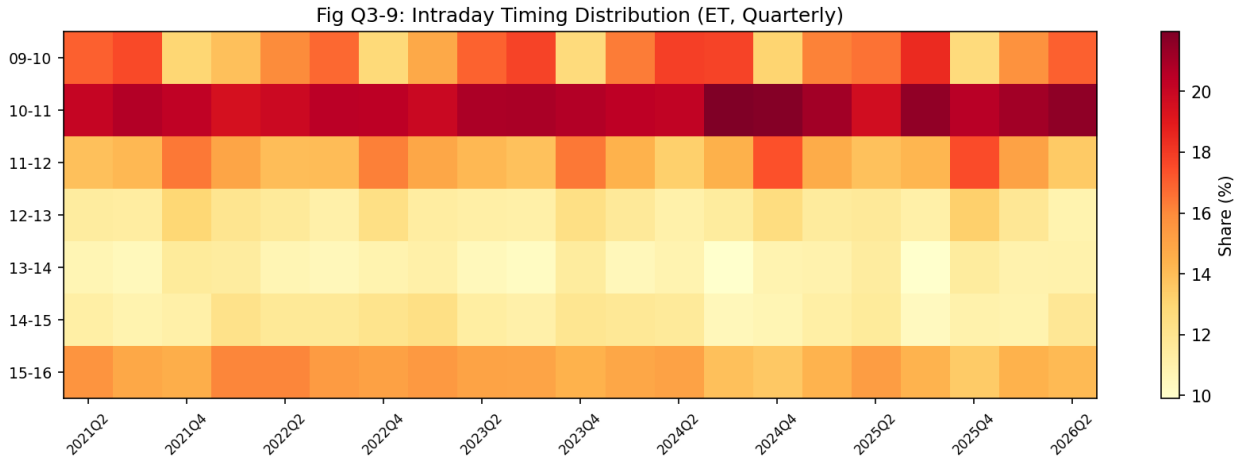

Figure 8: Quarterly intraday timing distribution of large-order activity. The 10–11 ET bucket is consistently the most active interval, with a secondary concentration near the close.

7 Limitations

Several limitations shape the interpretation of the results. First, the paper does not observe open/close flags, inventory context, or ultimate beneficial ownership. Even though OptionWhales reconstructs trade direction using NBBO-based methods, those direction classifications should not be read as direct evidence of net bullish or bearish intent. Second, strategy evidence is generated from same-second clustering, which likely understates the true use of multi-leg structures when legs are executed in staggered fashion and may occasionally overstate it when unrelated trades coincide in the same second.

Third, the 50-contract threshold means that the paper studies large-order flow rather than the entire options market. The results are informative about a behaviorally important slice of trading, but they should not be generalized mechanically to smaller retail-sized orders. Fourth, some data fields improve over time. Quote coverage is sparse in 2021 and substantially better later in the sample; Greek coverage also increases. Cross-period comparisons that rely heavily on those fields should therefore be treated with caution.

Finally, the annual summary tables intentionally relabel the first and last rows as partial-year slices, but the sample still does not span six complete calendar years. The results are best read as evidence from a rolling five-year window rather than from six directly comparable yearly panels.

8 Conclusion

This paper documents a clear evolution in large U.S. option order flow between April 2021 and April 2026. The most robust conclusions are descriptive. Large-order activity becomes larger and broader. Cluster-based evidence becomes more consistent with multi-leg trading. And the market shifts sharply toward short-dated contracts, especially 0DTE, while remaining centered on at-the-money exposure.

These findings matter because they describe a market that is becoming more tactical in both maturity choice and execution style. A larger share of economically important option flow is concentrated in instruments with very short time horizons and in clusters that are more consistent with defined payoff structures than with isolated single-leg bets. At the same time, the paper's results should be interpreted within the limits of the available fields. The dataset is rich enough to describe how large-order flow changed, but not rich enough to fully identify why every trade occurred.

Future work could extend the present descriptive evidence in several directions: linking flow patterns to realized volatility and macro events, studying differences between index-linked and single-name 0DTE adoption, and building richer models of strategy inference that relax the same-second clustering assumption. Those extensions would deepen interpretation, but they would not change the central descriptive result of this paper: by 2026, the large-order options market is materially larger, more complex, and much shorter-dated than it was in 2021.

A Supplementary Figures

Figure 9 shows that growth in large-order activity is accompanied by broader participation across underlying tickers. Figure 10 summarizes the changing mix of size-two clusters. Figures 11 and 12 provide complementary views of the shift toward shorter maturities and the persistent concentration of trading near at-the-money strikes.

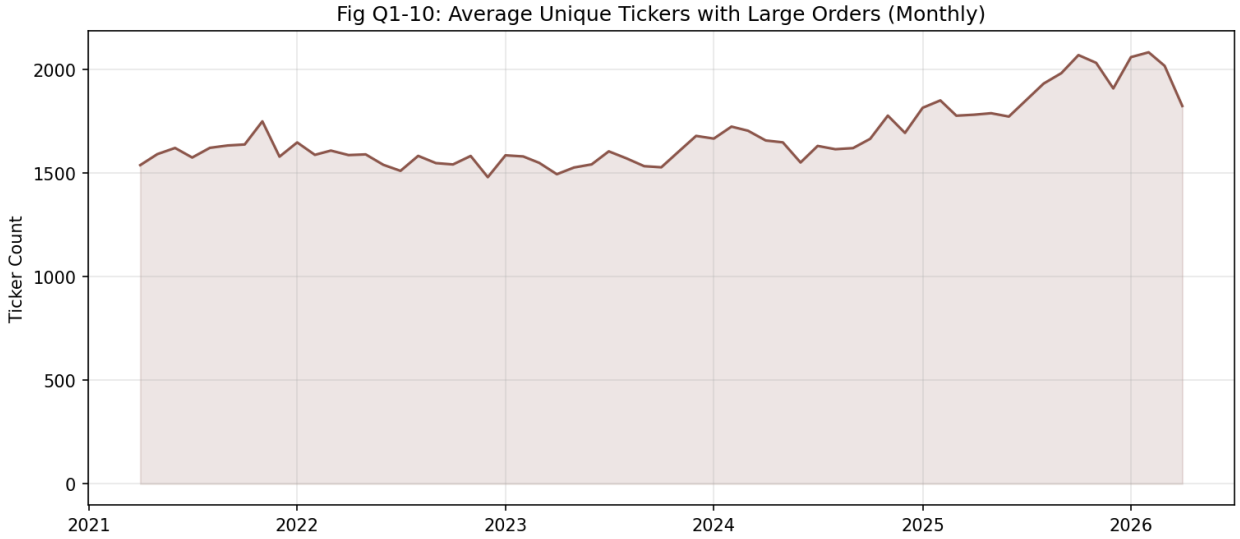

Figure 9: Average number of unique tickers with large orders by month. Breadth rises meaningfully alongside total activity.

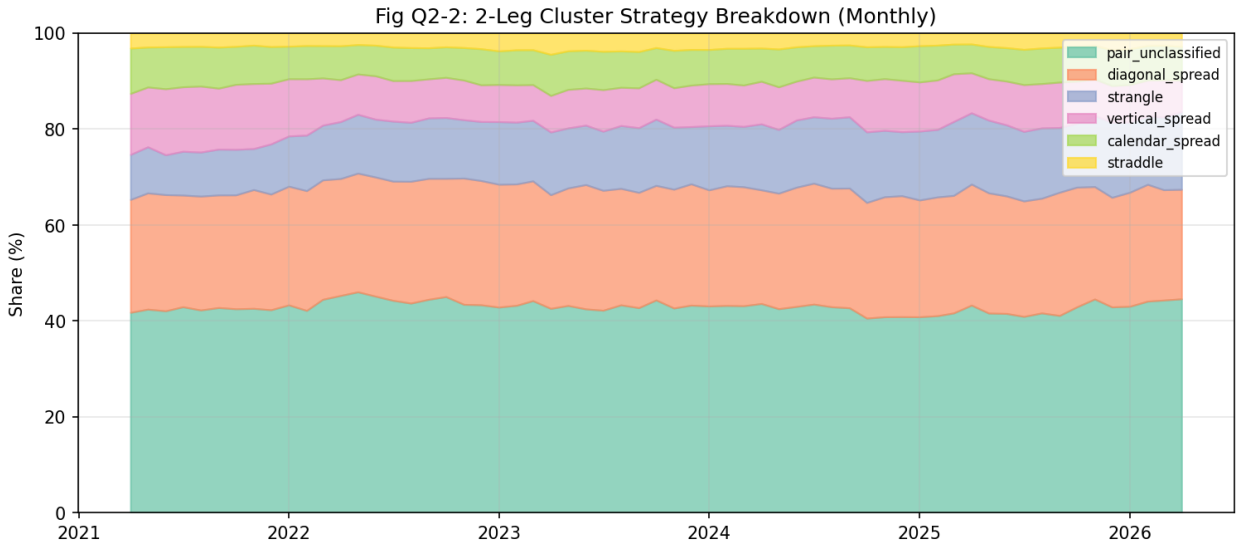

Figure 10: Composition of size-two clusters by month. Strangles gain share over time, while vertical spreads lose relative importance.

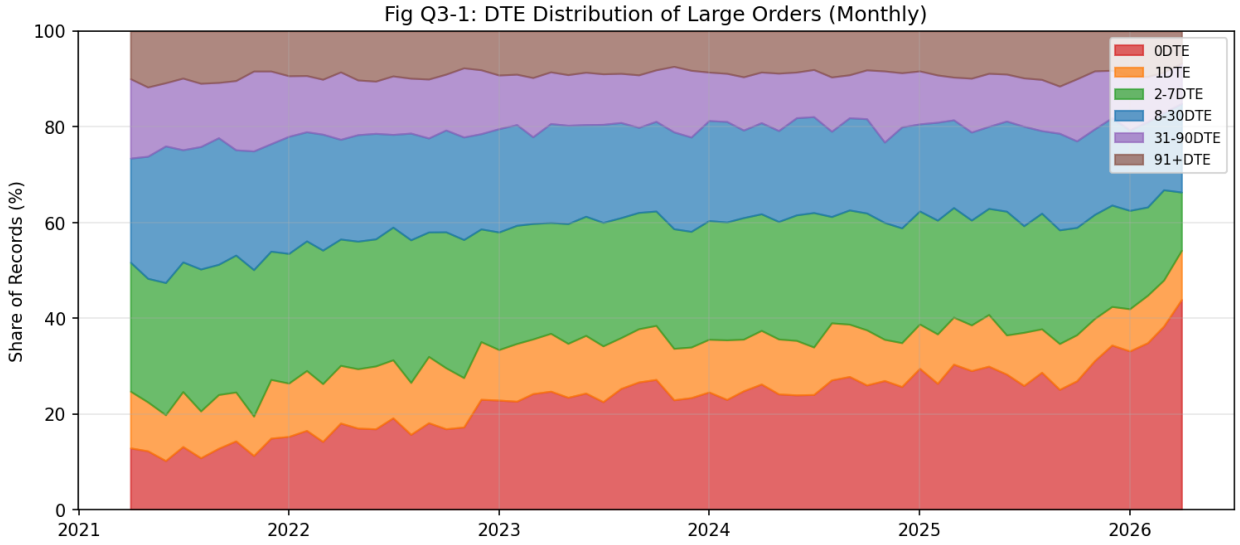

Figure 11: DTE distribution of large-order records by month. Short-dated buckets gain share over time, with the most visible growth in 0DTE contracts.

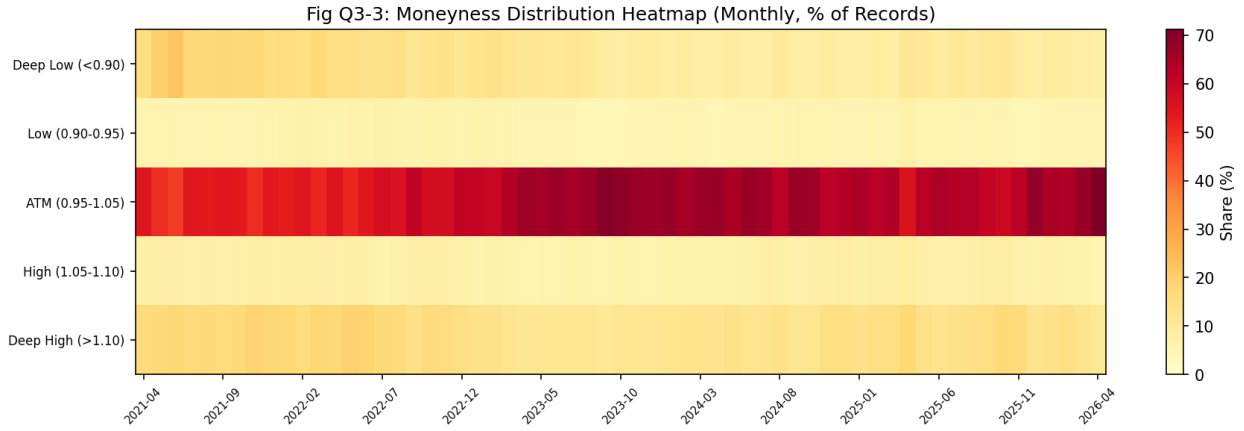

Figure 12: Moneyness distribution by year and bucket. The market remains strongly centered on at-the-money contracts throughout the sample.

B Key Data Fields

Table 4: Selected fields used in the analysis.

Field	Description
underlying_ticker	Underlying stock or ETF symbol.
option_ticker	OCC-format option symbol containing expiration information.
option_type	Call or put.
strike	Strike price of the option contract.
total_contracts	Number of contracts in the order.
premium	Dollar premium, measured as contracts times option price times 100.
total_notional	Notional exposure field provided in the processed large-order data.
order_type	Execution-type classification such as sweep, burst, or single.
direction	Trade direction label such as buy, sell, or neutral.
delta, gamma, vega, theta	Option Greek fields when populated.
implied_volatility	Annualized implied volatility estimate.
avg_bid, avg_ask	Best bid and ask reference values when quote data are available.

C Cluster Classification Rules

Table 5: Rule-based strategy patterns used in cluster classification.

Cluster Size	Strategy	Conditions
2	Vertical spread	Same option type, same expiration, different strikes, with one buy and one sell.
2	Straddle	One call and one put with the same strike and same expiration.
2	Strangle	One call and one put with different strikes and the same expiration.
2	Calendar spread	Same option type and strike, but different expirations.
2	Diagonal spread	Same option type with different strikes and different expirations.
2	Split fill	Same option type, same strike, same expiration.
3	Butterfly	Three same-type options with the same expiration and arithmetic strike spacing.
4	Iron condor	Two calls and two puts with the same expiration and puts below calls.
4	Iron butterfly	Iron-condor pattern with equal middle strikes.
5+	Complex cluster	Multi-leg pattern that does not map cleanly to the simpler templates above.