Systematic Short-Dated Options Strategies on US Large Caps

Iron Condors, Calendar Spreads and Double Calendars on the 40 Largest S&P 500 and Nasdaq 100 Constituents

A five-year backtest, an economic fair-value volatility model, and an honest account of what the evidence will and will not support.

Report date: 19 August 2026. All market data through the 18 August 2026 close.

Scope and standing caveat

Every price path in this study is real: daily open, high, low and close bars for the 40 names and the two indices, the Cboe volatility index complex, Treasury yields and macroeconomic series, and quarterly filing dates taken from the companies' own submissions to the Securities and Exchange Commission.

Option prices are not observed. No historical option quote database was available for this work. Every option in the backtest is priced on a volatility surface reconstructed from those observable inputs, fitted on 2016 to 2021 and then frozen, so that the entire five-year test window is out of sample. The construction is set out in full in chapter 10 and validated by rebuilding the VIX itself from S&P 500 price history alone. Readers should treat every performance figure in this report as conditional on that model, and chapter 15 quantifies exactly how much the conclusions move when its central assumption is varied across the range the literature supports.

Executive Summary

Three short-dated option structures — the iron condor, the calendar and the double calendar — were tested on a 40-name US large-cap universe over a five-year window ending at the 18 August 2026 close, under a specification supplied as a constraint rather than derived from a search. That mandated specification is an implied volatility rank above 78, a front leg 7 to 12 days from expiry, and short-leg deltas of 25 to 35 for the condor and 30 to 40 for the calendars. Exits follow the managed rule: a 50% profit target, a 200% stop and a close at two days to expiry, whichever comes first. The finding is that all three programmes work, none works well, and the margin by which they work is smaller than the uncertainty attached to the two assumptions that generate it.

0.1 What the study found

0.1.1 Under the mandated specification the iron condor returned +4.12% a year with a Sharpe of 0.84 and an 8.2% maximum drawdown; the calendar +3.57% at 0.61 and the double calendar +2.48% at 0.52

Strategy, managed rule	Trades	Win rate	CAGR	Sharpe	Maximum drawdown	CAGR across the assumed premium range, 1.061 to 1.262
Iron condor	822	67.2%	+4.12%	0.84	-8.2%	+4.12% to +11.85%
Calendar	807	59.7%	+3.57%	0.61	-5.6%	+2.50% to +3.57%
Double calendar	807	57.4%	+2.48%	0.52	-7.4%	+0.79% to +2.48%

The right-hand column is the same backtest re-run at three assumed levels of single-stock option richness, and it is wider for the condor than the entire spread between the three strategies. Two further facts belong beside the headline. Held to front expiry, without any exit rule, all three lost money — the condor at -0.76% a year, the calendar at -2.36% and the double calendar at -2.09% — so what is being measured is management as much as geometry, and chapter 16 attributes most of the improvement to the time stop rather than to the profit target or the stop loss. And the ranking is not interpretable at conventional confidence: against 38 configurations and 570 configuration-rule cells searched, the deflated Sharpe probability is 0.334 for the condor, 0.071 for the calendar and 0.025 for the double calendar, where a value below 0.95 means the result cannot be distinguished from the product of the search itself. Chapter 14 adds a second caution against reading the 0.84 as evidence of edge: the Sharpe ratio is maximised, in the complete absence of skill, by precisely this payoff shape. A block bootstrap of 4,000 paths puts the condor's probability of ending the period down at 24.7%.

0.1.2 The entire result rests on one assumption: at a single-stock implied-to-expected-realised ratio of 1.061 the condor earns +4.12%, at 1.262 it earns +11.85%, and the literature supports 1.03 to 1.07

This is the study's dominant sensitivity and the reason its conclusions are stated cautiously. The same 822 condor entries, on the same names and dates, over the same price paths, with the same strikes and the same managed exits, compound at +4.12% a year under the calibrated assumption and +11.85% under the richest one tested. No price path differs between those two figures; a single scalar governing how expensive single-stock options are assumed to be relative to subsequently realised volatility does all of the work. Chapter 4 places the published evidence at the low end of that range: for liquid large caps the implied-to-realised ratio runs about 1.03 to 1.07, this study calibrates to a median modelled ratio of 1.060 against 1.364 for the index, and the 1.128 and 1.262 rows sit outside the supported range entirely. Part of the literature, testing at stock level, cannot reject a premium of zero at all. The direction of the exposure is also not uniform: raising the assumption drives the double calendar monotonically down, from +2.48% to +1.82% to +0.79%, because a calendar buys as well as sells. Only the net short-premium structure gains, which is why the condor's headline is the one most exposed. A reader who takes a lower view of single-stock richness than this study does should mark the condor's return down accordingly, and there is no result in this report that survives that adjustment intact.

0.1.3 Costs of about $3.25 per contract per leg consume 4.4% to 7.0% of the premium; at twice that slippage the condor earns +0.85% and both calendars lose money

Round-trip execution costs measured in the backtest are $18.88 for the condor against a median premium of $247, $22.46 for the calendar and $43.95 for the double calendar — 7.0%, 4.4% and 4.4% of premium respectively, on average holding periods of 4.8 to 5.3 days. At zero slippage the three structures are barely distinguishable, at +7.42%, +7.75% and +7.44%, so every performance difference reported anywhere in this study is a difference in what each pays to trade rather than in what it earns from the volatility surface. Doubling the base assumption leaves the condor at +0.85%, the calendar at -0.78% and the double calendar at -2.23%; trebling it takes all three to losses between -2.42% and -6.58%. Chapter 17 shows that a multiple of that size is not hypothetical: the published measurements it assembles for the first quarter of 2020 record quoted option spreads more than doubling in percentage terms while posted size roughly halved, with individual names dispersing far beyond the market-wide average — and the widening arrives in exactly the week a short-premium book most needs to close. The base cost is itself anchored to a January 2020 measurement, with no 2026-vintage replacement located.

0.1.4 The VIX at 15.83 sits 1.27 points below an economic fair value of 17.10, a 0.52 standard deviation discount, so premium is cheap rather than rich

An economic fair-value model for the VIX, regressing log VIX on realised volatility, credit spreads, financial conditions, the yield-curve slope, policy uncertainty, jobless claims and the equity drawdown across 2,493 daily observations, explains 79.8% of the variation with a root mean squared error of 3.53 VIX points. On 18 August 2026 it puts fair value at 17.10 against an actual 15.83 — a discount of 1.27 points, a residual of -0.52 standard deviations. The macro-only variant, excluding realised volatility, places fair value slightly higher still at 17.26, so the cheapness is not

an artefact of quiet realised markets alone. A 0.52 standard deviation gap is inside the normal range and is not a trading signal, and the model is contemporaneous rather than predictive. The implication for a programme that sells premium is nonetheless unhelpful: this is not a moment when volatility is being overpaid. It is a separate question from the variance risk premium of 0.1.2, which compares implied variance to subsequent outcomes rather than to a fitted level, but the two do not currently point the same way.

Figure 0.1 places the three equity curves on one axis, and the visual impression is the report's position in miniature: three lines that rise, stay close together, and are separated by less than the assumptions underneath them move.

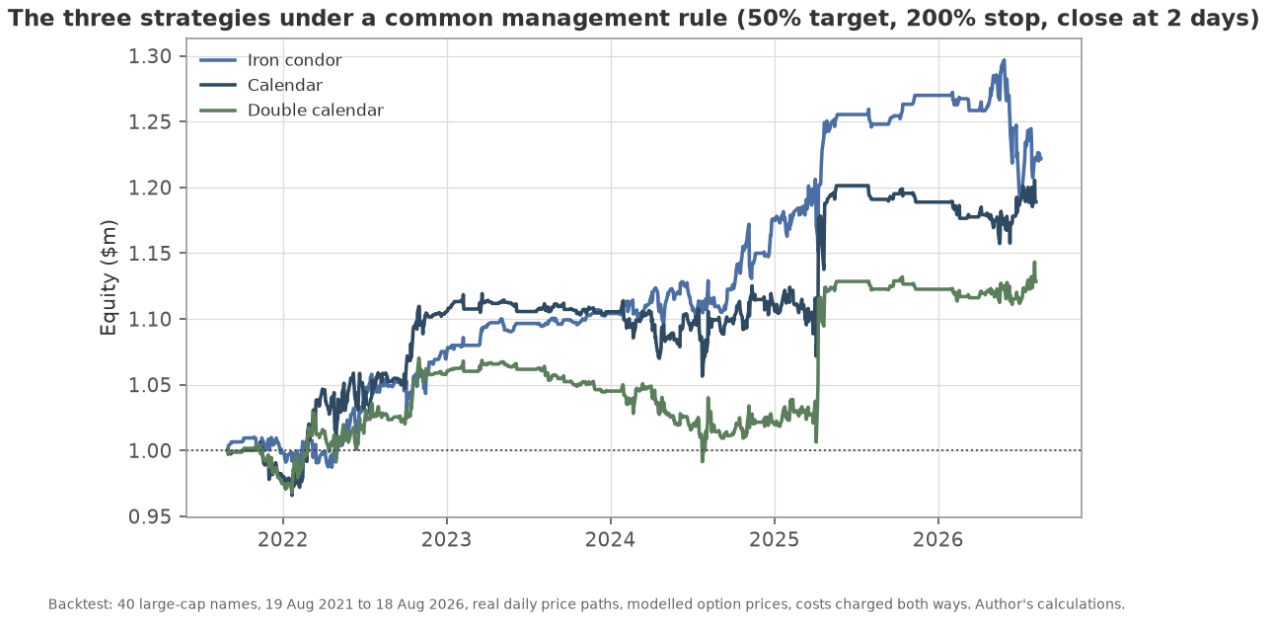


Figure 0.1 — The three strategies compounding side by side under the managed rule: a narrow band of outcomes, long runs of small gains, and isolated deep months.

One qualification conditions everything above. No historical option quotes were available to this study; every implied volatility behind every trade is produced by a frozen volatility model rather than observed in the market. The model's out-of-sample reconstruction of the VIX from S&P 500 prices alone achieves a correlation of 0.818 and a root mean squared error of 3.15 points across 1,254 days, which is a fair result for an index and no guarantee at single-stock level, where the premium is assumed rather than measured. The reconstruction also runs +0.60 points rich at nine days and -1.70 points cheap at three months, a tilt that flatters short-dated selling. Read the returns above as conditional on that surface. What follows begins by establishing the market state the surface was fitted to, and then, chapter by chapter, exposes each assumption holding these numbers up.

---

1. Market Conditions in August 2026

1.1 Where the indices stand

1.1.1 The S&P 500 closed at 7,691.76 on 18 August 2026, 1.37% below its running peak, and the Nasdaq 100 at 29,490.96

The S&P 500 closed at 7,691.76 on 18 August 2026 and the Nasdaq 100 at 29,490.96. These are the settlement levels against which every backtest, calibration and trade example in this report is anchored; no observation after that close enters the study. Measured against its running peak — the highest close reached at any point up to and including the observation date — the S&P 500 sat 1.37% lower. The reference point is recent: Reuters reported a record close of 7,798.99 on 13 August, three sessions earlier.[1]

Table 1.1 — Index state at the 18 August 2026 close, from this study's own price data.

Measure	Value
S&P 500 close	7,691.76
S&P 500 drawdown from running peak	1.37%
Nasdaq 100 close	29,490.96
S&P 500 realised volatility, Yang–Zhang, 21 day	11.6%
Nasdaq 100 realised volatility, Yang–Zhang, 21 day	21.2%

Two features of Table 1.1 shape everything that follows. The first is that a 1.37% shortfall from the peak is a pause, not a correction; the market enters the sample window in an undisturbed uptrend, which matters because short-dated premium strategies are path-dependent and their measured performance is conditioned on the drawdown environment they were tested in. The second is the volatility gap between the two indices. Realised volatility here is the Yang–Zhang estimator over a 21-day window, an open-high-low-close construction that uses the intraday range as well as the close-to-close move and so is less noisy than a simple return standard deviation. At 11.6% for the S&P 500 against 21.2% for the Nasdaq 100, the technology-weighted index was carrying materially more realised risk than the broad index on the same estimator and window. A programme trading single large-cap names is exposed to the second number far more than the first.

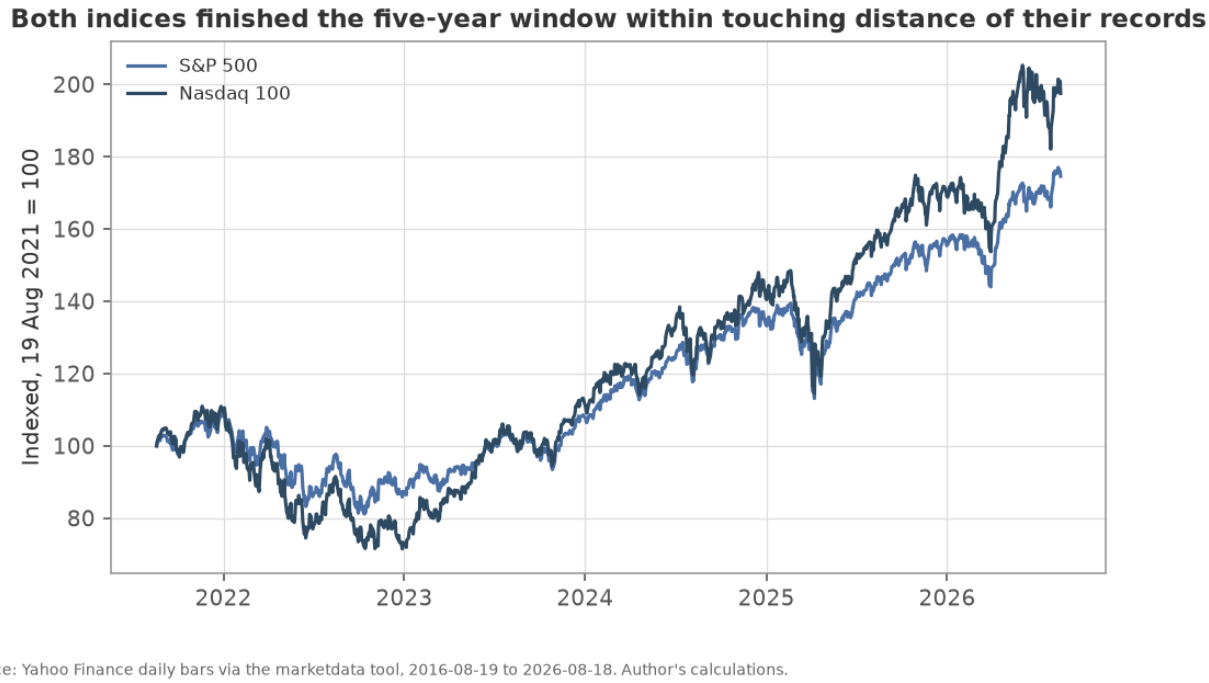


Figure 1.1 — Index levels for the S&P 500 and Nasdaq 100 through the 18 August 2026 close, showing both indices finishing the sample within a short distance of their highs.

1.1.2 The August record was set on earnings rather than re-rating: Q2 2026 blended growth of 50.4%, falling to 32.0% excluding Alphabet and Amazon

Blended year-on-year earnings growth for the S&P 500 in the second quarter of 2026 was 50.4% as of 7 August, against a five-year average of 15.2% and a ten-year average of 11.2%, and against the 23.1% expected at the 30 June quarter end. It was the strongest quarter since the 91.6% posted in the second quarter of 2021.[2] The headline overstates the operating result. Excluding Alphabet and Amazon, growth falls to 32.0%; Alphabet's reported earnings per share included a net gain of $98.0bn, primarily unrealised gains on equity securities, with comparable valuation gains at Amazon.[2] [3] The index advanced more than 6% in twelve trading days into the record.[4]

The distinction between an earnings-driven and a multiple-driven record is not decorative for a premium-selling programme. Earnings growth of this size is delivered discontinuously, on scheduled announcement dates, and it concentrates single-stock risk into a small number of sessions. The study's own event detection identifies 1,577 earnings reaction days across the 40-name universe, on which the mean absolute price move was 5.48% against 1.38% on ordinary days — a ratio of 3.96. A market whose gains arrive through the earnings channel is a market in which the gap risk facing a short option position is repeatedly concentrated rather than diffuse.



Source: Yahoo Finance daily bars via the marketdata tool, 2016-08-19 to 2026-08-18. Author's calculations.

Figure 1.2 — S&P 500 drawdown from its running peak, ending at 1.37% on 18 August 2026; the shallow recent troughs are the profile against which the backtest drawdowns in later chapters should be read.

1.1.3 Forward multiples of 20.0 times for the S&P 500 sit above the five-year average of 19.9 and the ten-year average of 19.0

The forward twelve-month price-to-earnings multiple for the S&P 500 was 20.0 on 7 August, above the five-year average of 19.9 and the ten-year average of 19.0, but below the 20.4 recorded at 30 June.[2] The index reached roughly 23 times forward earnings in October 2025, so the record close was set on a lower multiple than the market carried then: the denominator moved, not the rating.[4] For the Nasdaq 100, aggregator estimates cluster near 22 times forward, with one source putting it at 22.02 on 18 August against a five-year median of 24.78 — the 16th percentile of its five-year range.[5] That figure comes from a data aggregator rather than the index provider and should be treated as indicative rather than authoritative; no official index-provider valuation was obtained for either index.

The reading is genuinely two-sided, and this report does not resolve it. Against the ten-year average of 19.0 the S&P 500 is expensive, and the International Monetary Fund's April 2026 assessment describes stretched valuations and concentration, particularly among artificial-intelligence-related firms, as raising downside risks, warning that a reassessment of expectations regarding AI could trigger an abrupt financial market correction.[6] Against October 2025 the index has de-rated. Both statements can hold simultaneously, and for structures whose average holding period in the backtest here runs from 4.8 days for the iron condor to 5.3 days for the double calendar, neither is a timing signal. What the valuation backdrop does supply is a prior on the shape of the tail: a market priced above its ten-year average with an identified single-factor vulnerability is one in which the left tail of a short-premium distribution deserves more weight than

its historical frequency alone would assign.

1.2 Concentration and breadth

1.2.1 The top ten names are close to 40% of the index against a 27% peak in 1999-2000, but on a P/E of 31 against 21 for the rest

The ten largest S&P 500 constituents account for approximately 40% of index weight, roughly double the 18% of a decade ago.[7] Columbia Threadneedle placed the top ten near 39%, well above the 27% peak reached during the technology bubble of 1999-2000, but on a price-to-earnings multiple of 31 against 21 for the remainder of the index — a premium of about 50%, against more than 100% in 2000 — while contributing 30% of index earnings against under 20% at that peak.[8] That analysis is dated September 2025, so the weight and multiple comparisons may have moved; the 1999-2000 comparison point is historical and stable. The Magnificent 7 group represents about one-third of index weight.[9]

Concentration on this scale is the mechanical reason an index-level volatility reading understates the risk carried by a portfolio of individual names. When a handful of constituents dominate the weight, index variance is largely their variance plus their covariance with one another, and the remaining constituents contribute little to the index while retaining their own idiosyncratic variance in full. A programme that sells premium on forty individual large caps is not selling a diluted version of index premium; it is selling forty separate distributions, most of which the index barely registers.

1.2.2 Breadth is genuinely broad: the equal-weight index rose more than 13% against 8.5% cap-weighted

For 2026 to 29 July, the equal-weight S&P 500 rose more than 13% against 8.5% for the standard capitalisation-weighted index, with mid-caps up around 15% and small-caps around 19%. Since the 2 June record, about two-thirds of components rose and eight of eleven sectors were higher, with the equal-weight index up 4% against a cap-weighted decline of 2.4%, and only one Magnificent 7 name had beaten the index for the year to date.[9] The precise size of the breadth gap is not firmly established — a second, undated source reports materially different magnitudes — but the direction is consistent across sources, and the direction is what matters here.

Breadth of this kind is the opposite of the 2023-24 pattern in which index gains were delivered by a narrow leadership group. It means the index is being held up by rotation among many names rather than by the top ten, and rotation is precisely the condition under which individual constituents move substantially while the index does not. For a single-stock premium programme this is favourable in one specific sense: the option premium available on any given name reflects that name's own volatility, while the offsetting index exposure that would arise from trading many names at once is suppressed by their disagreement.

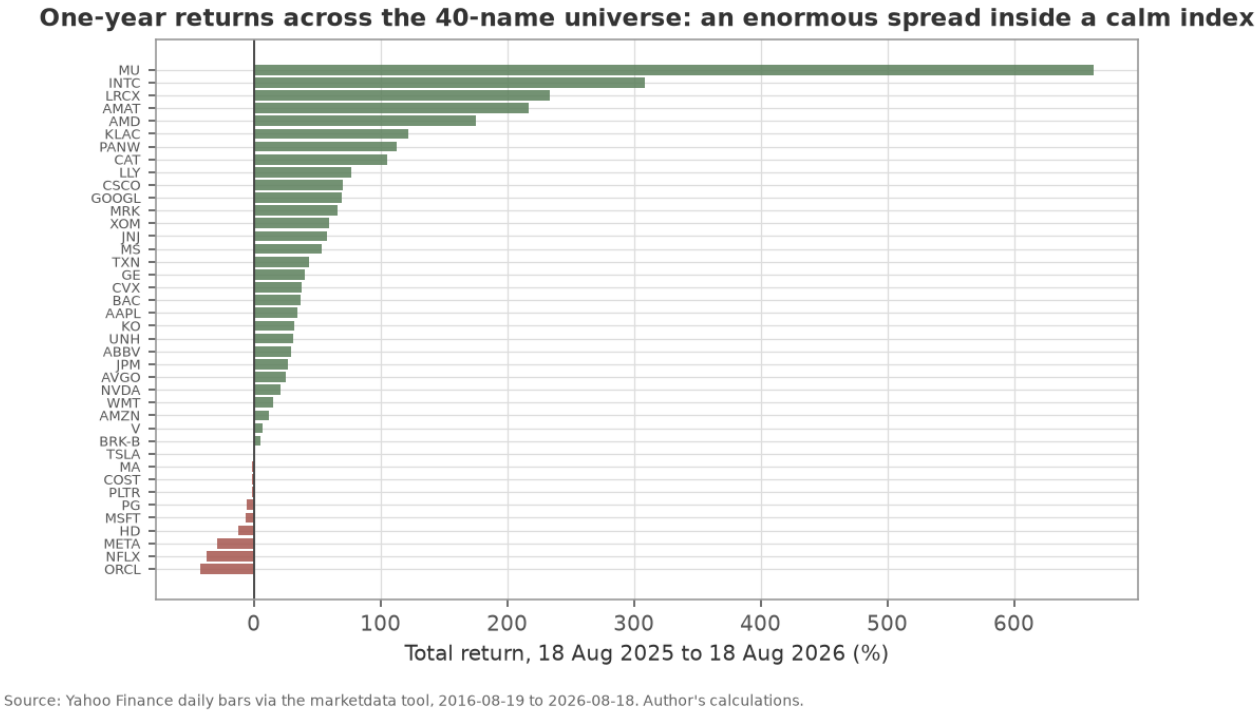
1.2.3 One-year returns across the 40-name universe span an enormous range, and that dispersion is what a single-stock premium programme actually trades

The cross-section of one-year returns across the study's forty names is extraordinarily wide.

Table 1.2 — One-year total returns across the 40-name universe to the 18 August 2026 close, this study's calculation.

Measure	Value
Names in the universe	40
Best one-year return	+662.7% (MU)
Median one-year return	+33.4%
Worst one-year return	-42.0% (ORCL)

The span in Table 1.2 runs from -42.0% to +662.7%, and the two extremes are not obscure names — Micron and Oracle are both large, liquid, heavily optioned index members. The median of +33.4% confirms that the distribution is not a narrow cluster with two outliers pinned to it: the typical name in the universe delivered a return that would itself be a strong year for the index. This is the practical content of the breadth observation in 1.2.2. Where index-level dispersion is high, the volatility embedded in individual name options is materially higher than the index volatility, and the two are not substitutes. A short-dated premium programme is compensated for bearing single-name outcomes, and the width of this cross-section is the direct measure of what those outcomes look like. It also implies that position sizing cannot be calibrated from index statistics; the relevant tail is the constituent tail.



Source: Yahoo Finance daily bars via the marketdata tool, 2016-08-19 to 2026-08-18. Author's calculations.

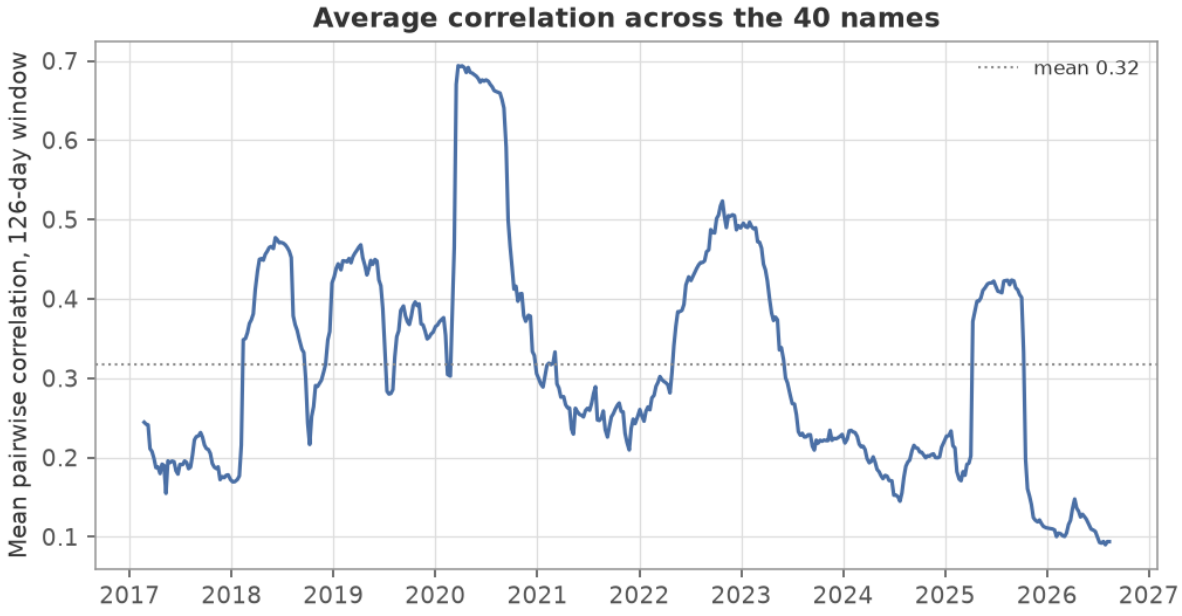
Figure 1.3 — One-year returns for the 40 constituents, with the range running from Oracle at -42.0% to Micron at +662.7% around a median of +33.4%.

1.3 The correlation regime

1.3.1 Mean pairwise correlation across the 40 names is 0.094 against a ten-year mean of 0.317, near the bottom of its range

Mean pairwise correlation across the forty constituents — the average of the correlation coefficients between every distinct pair of names in the universe — stood at 0.094 on 18 August 2026, against a ten-year mean of 0.317 and a ten-year range of 0.090 to 0.694. The current reading therefore sits just above the lowest value observed in a decade and well below the decade average.

A mean pairwise correlation of 0.094 means that constituent returns are close to unrelated to one another on the measurement horizon. Historically this regime is asymmetric: the ten-year range is bounded below at 0.090, barely beneath the current reading, and extends to 0.694 above it, so there is little room to fall and a great deal of room to rise. That asymmetry is the single most important piece of context for a strategy that runs many concurrent single-stock positions, because the diversification such a book relies on is a direct function of this number. The portfolio construction chapter records a maximum of 39 positions open simultaneously against a nominal cap of 12; the risk of that concurrency is benign at a correlation of 0.094 and is not benign at 0.694.



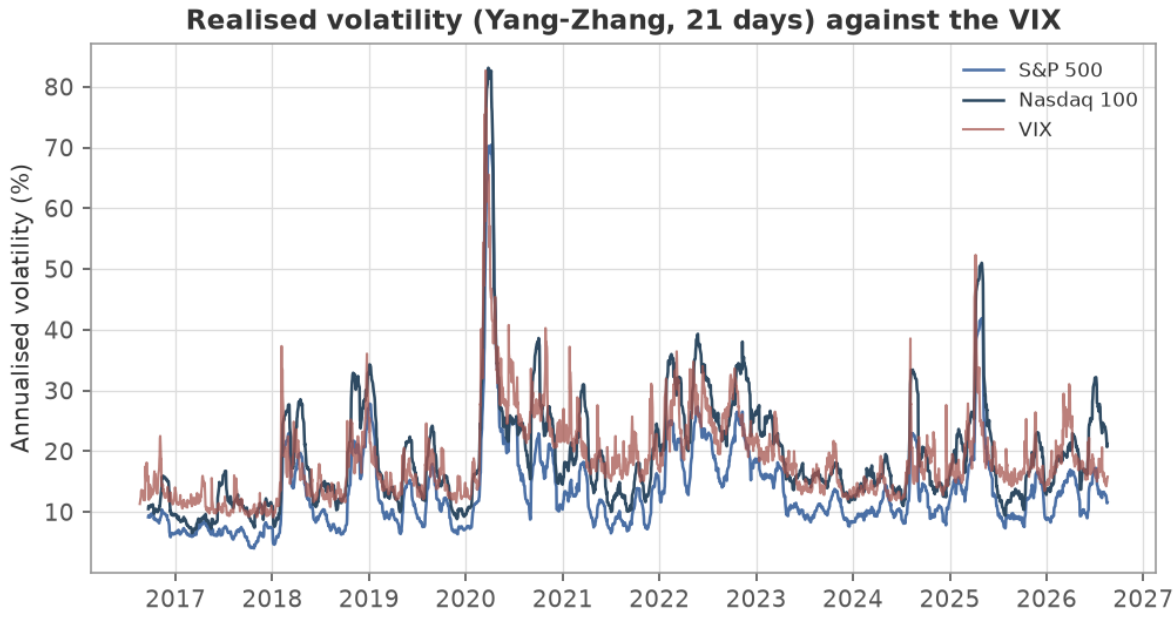
Source: Yahoo Finance daily bars via the marketdata tool, 2016-08-19 to 2026-08-18. Author's calculations.

Figure 1.4 — Mean pairwise correlation across the 40 names, ending at 0.094 against a ten-year mean of 0.317 and a decade low of 0.090.

1.3.2 Low index volatility is therefore a correlation artefact, not evidence of low single-stock risk

S&P 500 realised volatility of 11.6% and a VIX of 15.83 describe an index whose components are moving in different directions, not an index whose components have stopped moving. The arithmetic is straightforward in direction if not in magnitude: index variance is the weighted sum of constituent variances plus all the pairwise covariance terms, and at a mean pairwise correlation of 0.094 those covariance terms contribute close to nothing. Suppress the correlation and the index becomes quiet while every constituent remains as volatile as before. The 21.2% realised volatility of the Nasdaq 100 against the S&P 500's 11.6%, and the constituent return range in Table 1.2, are both consistent with that account.

The practical consequence is that a quiet index is not evidence that single-stock premium is thin. It is evidence that the index has stopped aggregating the risk that is still present in its members — and, symmetrically, that a correlation shock would raise index volatility sharply without any individual name needing to become more volatile than it already is.



Source: Yahoo Finance daily bars via the marketdata tool, 2016-08-19 to 2026-08-18. Author's calculations.

Figure 1.5 — S&P 500 realised volatility against the VIX; the index-level volatility that frames the strategy is measured on a market whose constituents are close to uncorrelated.

The volatility indices themselves — their history, their term structure, and the skew and volatility-of-volatility measures that sit alongside them — are the subject of the next chapter, which establishes whether the 15.83 reading is cheap, fair or expensive relative to the conditions that generated it.

---

2. The Volatility Regime and the Cboe Complex

2.1 Where volatility sits

2.1.1 The VIX at 15.83 is the 40th percentile of the last ten years against a mean of 18.62

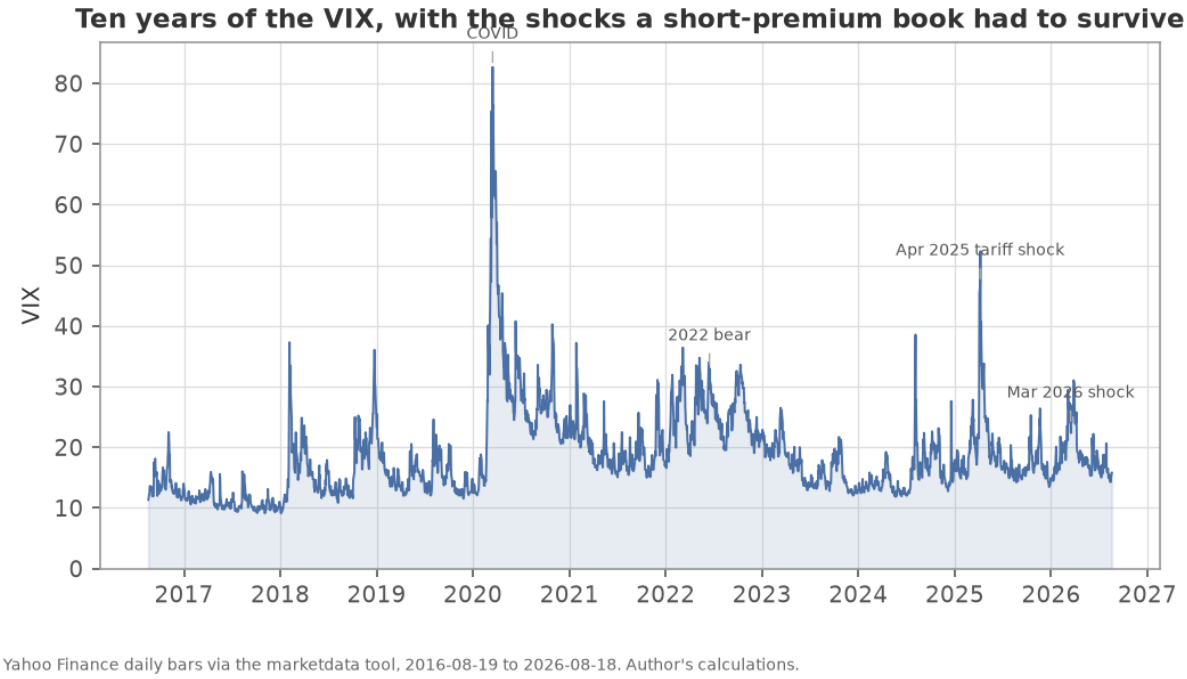
The Cboe Volatility Index (VIX) closed at 15.83 on 18 August 2026, the 40th percentile of the last ten years of daily closes against a ten-year mean of 18.62. The two are consistent because the VIX is right-skewed: infrequent spikes lift the mean above the median, so a level below the mean is not by itself unusual.

Reading, 18 August 2026	Level	Ten-year context on this sample
VIX	15.83	40th percentile; mean 18.62
VIX3M / VIX	1.094	in contango on 92.3% of days

Cboe SKEW Index	143.6	73rd percentile; mean 136.1
VVIX	92.9	not ranked on this sample
S&P 500 realised volatility (Yang-Zhang, 21 day)	11.6%	not ranked on this sample
Nasdaq 100 realised volatility (Yang-Zhang, 21 day)	21.2%	not ranked on this sample

The rows disagree with one another, and that disagreement is the substance of this chapter. Implied volatility sits below its ten-year average while the shape statistic sits above its own; the term structure is upward sloping but not flat; realised volatility on the S&P 500 is far under the implied level, while the Nasdaq 100 carries nearly twice the S&P 500 figure. Three rows carry no percentile: the Cboe VIX of VIX Index (VVIX), which prices options written on the VIX itself, and both realised-volatility estimates were not ranked against history in this study's calculations, and none is asserted for them here. A summary of the form "risk is low" is not supported by the table; each row has to be read on its own basis.

The VIX had dipped to 14.2 on Friday 14 August, its lowest of 2026, and named desks read that as complacency. Jonathan Krinsky of BTIG placed it in "a window that historically sees downside volatility", entered "with the market at all-time highs and VIX at YTD lows", and counted no 80%-downside-volume day in 2026 since October 2025 against an average year's 21[4]. That is a view on positioning, not a measurement.



Source: Yahoo Finance daily bars via the marketdata tool, 2016-08-19 to 2026-08-18. Author's calculations.

Figure 2.1 — Ten years of the VIX: the reading sits in the lower half of a range dominated by a few tall, short-lived spikes.

2.1.2 S&P 500 realised volatility of 11.6% against a VIX of 15.83 leaves roughly four points of gross premium at index level

Realised volatility on the S&P 500, estimated with the Yang-Zhang range estimator over 21 days, was 11.6% at the same close, leaving a gap to the VIX of roughly four volatility points. Three qualifications apply. The two are not measured on the same basis — the VIX is a 30-day risk-neutral expectation extracted from option prices, 11.6% a backward-looking estimate from realised ranges — so the difference indicates a premium rather than a tradeable spread; it is gross of execution; and it is an index quantity, while every position in this study is written on a single name.

That last point is the material one. The Nasdaq 100 realised 21.2% on the same estimator and window, and the 40-name universe traded here is closer in character to that figure. Chapter 1 set out why the index reading is depressed relative to its constituents; what carries forward is that the four-point index gap is not the premium these strategies collect. Chapter 4 measures the index and single-stock premia separately, and the backtest rests on the second.

2.2 The shocks the sample contains

2.2.1 April 2025 took the VIX to 60.13 intraday; March 2026 added six points to near 25 on Middle East and private-credit stress; June 2026 was idiosyncratic, with the chip index down 10.3% in a day

The five-year window contains three episodes large enough to matter to a short-dated premium book, and they are of three kinds. The tariff announcement of 2 April 2025 took the VIX to 60.13 intraday on 7 April, a level unseen since 2020 or 2008[10]; SIFMA's metrics put the April 2025 average at 31.97, peaking at 52.33 on the 8th from a 21.51 trough on the 2nd[11]. The Bank for International Settlements puts the orthogonalised shock at –8.7 standard deviations on 3 April and estimates it as taking the S&P 500 down roughly 13% on impact[12]. A second tariff episode on 10 October 2025 produced the largest single-session decline until June 2026[13].

March 2026 was macro-financial rather than policy-driven. Cboe recorded a 5% fall in the S&P 500 and 5.17% in the Russell 2000, the VIX climbing roughly six points to finish near 25 and breaching 30 several times[14]; mid-month the credit volatility index VIXIG had "almost tripled from its January low" on geopolitical risk and private-credit concerns, with the oil index OVX at its highest on record outside the 2020 pandemic[15]. It unwound as fast as it arrived: the S&P 500 rose 10.4% in April 2026 and the VIX fell more than 33% to near 16.9[16].

June 2026 was neither. On 5 June the S&P 500 fell 2.64% and the Nasdaq Composite 4.18% on strong jobs data and AI weakness[17], with the PHLX semiconductor index down 10.3% in "its deepest one-day loss since March 2020", erasing about $1.3trn on weak custom-AI guidance from one issuer[18]. For a programme selling single-name premium that is the most instructive of the three: a sector dislocation of that size need not register as a regime change at index level, and it

barely marks the series in Figure 2.1. No evidence was found that systematic selling aggravated any of the three: through the worst month of 2026 the Cboe S&P 500 PutWrite Index (PUTD) fell 3.2% and the Cboe S&P 500 One-Month Covered Call Index (CALD) 4.3%, both outperforming the underlying benchmark on premium income[14].

2.2.2 Spike decay has halved from about 25 days to nine, so the window in which a short-premium book bleeds has shortened

The shape of these episodes has changed as well as their frequency. Allianz Global Investors quantifies it: spikes "that used to occur every three to five years now appear every year. And the time for volatility to subside has halved, from around 25 days to only nine days"[19]. The record fits: April 2025 took "only five days for the VIX to peak, and another 14 days to revert to the level from which it started"[10]. For the mandated 7-to-12-day window this cuts both ways. A trade now spans a period comparable to the whole decay window rather than a fraction of it, so a position opened into a spike may see it unwind within its own life; but an annual rather than quinquennial arrival rate means a five-year sample should contain several, and the loss is taken at the top of the spike however fast the decay that follows. Tail statistics by structure are in chapters 11 to 14.

The structural account behind the change is contested and is reported as such. Allianz names five forces, among them growth in zero-days-to-expiry (0DTE) options, electronic liquidity that retreats under stress and volatility targeting, and argues that the dampening effect of overwriting withdraws exactly when it is needed, since rising volatility lets investors "sell a smaller amount of call options to earn similar premiums"[19]. The complex is large: US options-based exchange-traded funds exceed $250bn, with record 2025 net inflows of $82.4bn[20]. Cboe's own measurement of 0DTE dealer gamma points the other way, finding "no increase in frequency of… gap moves"[21], though it dates from 2023 and predates this sample. No documented episode was found in which short-dated selling amplified a 2025 or 2026 move; the claim is unevidenced for the period rather than refuted.

2.3 Term structure and the tails

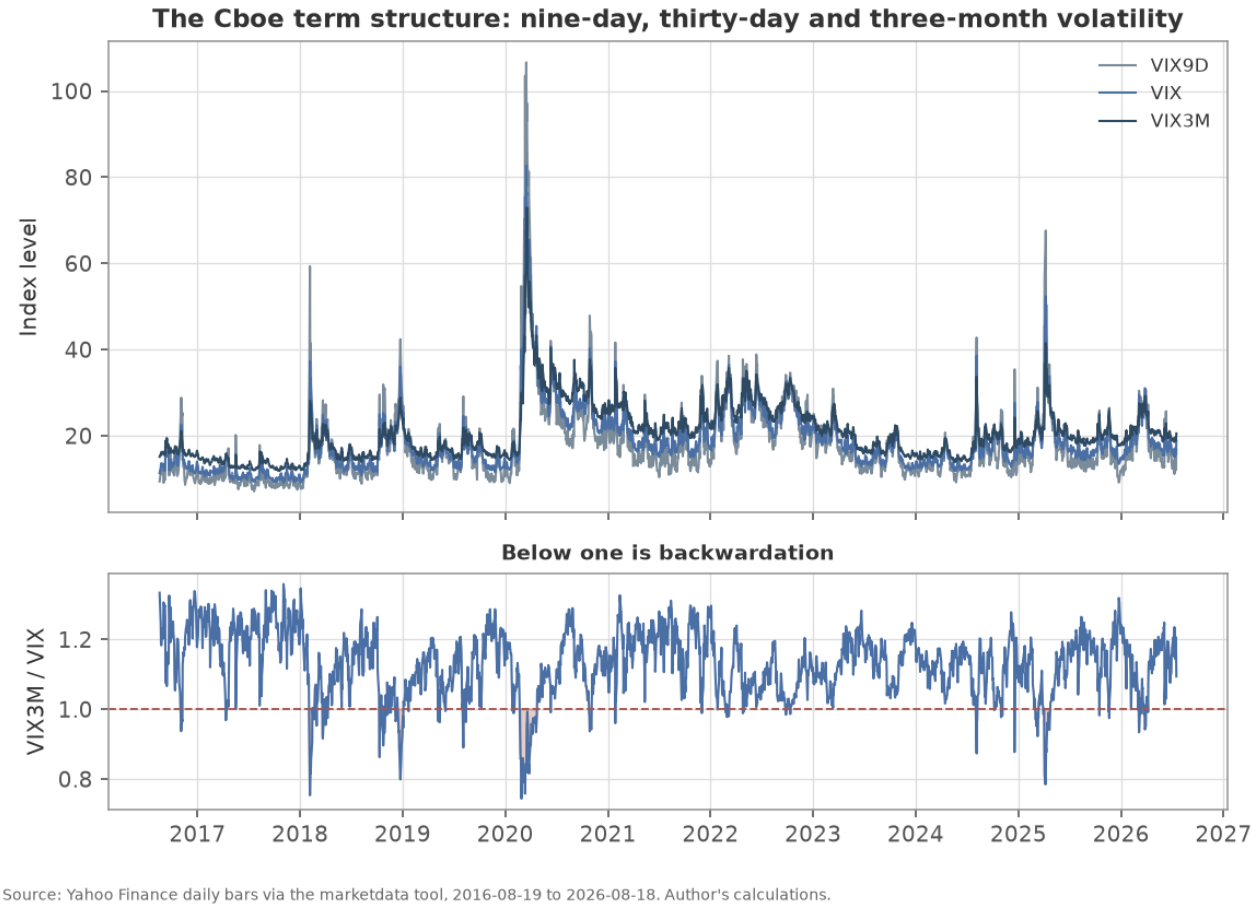
2.3.1 VIX3M over VIX stands at 1.094 and the curve was in contango on 92.3% of the last ten years of days

The ratio of the three-month index, VIX3M, to the 30-day VIX stood at 1.094 on 18 August 2026, and on that pair the curve was in contango on 92.3% of the last ten years of daily closes.

That percentage is definition-dependent and the definition must travel with it. The figure above is measured on the **cash indices**, as VIX3M divided by VIX, contango being any day on which the ratio exceeds one — not the commonly cited statistic. Cboe states that "the VIX futures curve is in contango more than 80% of the time since 2010"[22], and OptionMetrics counts 18.8% of daily observations backwardated since February 2006, implying about 81.2% contango[23] — both measured on the **futures** curve, and a crisis-heavy 2004 to 2009 sample reads lower again at

67.6%[24]. What is safe to assert is directional: the further apart the two tenors, the rarer the inversion, and the near end inverts first and most often. No source at this study's evidence bar counts contango on the cash pairs, so 92.3% is this study's own measurement and is not comparable to the published futures figures.

That bears on the mandate. Cboe's five S&P 500 gauges — VIX9D, VIX, VIX3M, VIX6M and VIX1Y — share one variance-swap methodology, with constituent options "weighted inversely proportional to the square of their strike"[25][26]. A 7-to-12-day short leg sits at the near end of that structure, where inversion is materially more frequent than a 92.3% contango share on the widest available cash pair implies. What the curve looks like name by name when the entry screen fires is measured in chapter 5.



Source: Yahoo Finance daily bars via the marketdata tool, 2016-08-19 to 2026-08-18. Author's calculations.

Figure 2.2 — The curve is upward sloping at 1.094 on the VIX3M/VIX cash pair, its state on 92.3% of the last ten years of days.

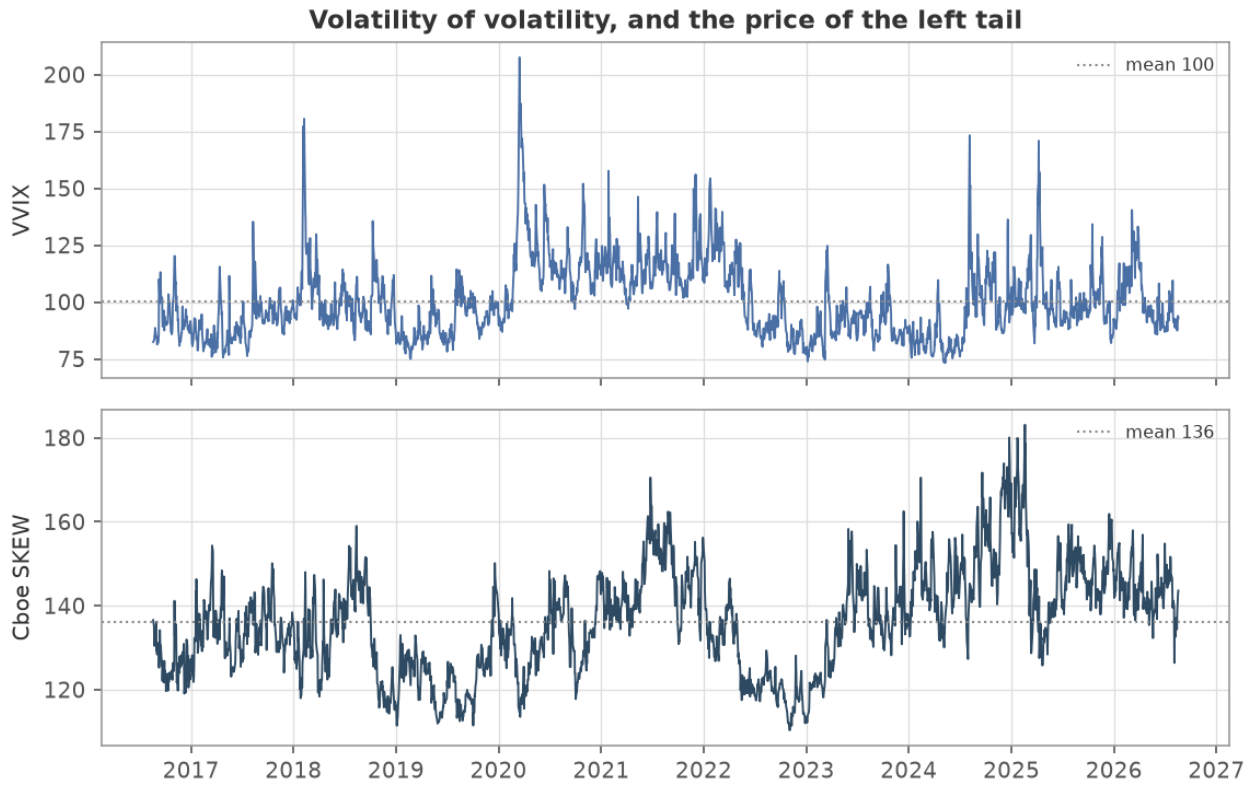
The slope shown in Figure 2.2 carries information, but not the information usually attributed to it. Johnson finds that the shape "conveys information about the price of variance risk rather than expected changes in the VIX, a rejection of the expectations hypothesis", the slope factor negatively predicting returns on 18 variance assets across 36 tests, 33 of them statistically significant, while the level factor does not predict at all[27]. Simon and Campasano reach the same conclusion from the futures basis, which "does not have significant forecast power for the change in the VIX spot index" but does forecast futures returns and "reflects a risk premium that can be

harvested"[28]. At index level, then, steep contango forecasts good returns to selling volatility and inversion poor ones. Whether that sign survives translation to single-name calendar spreads is tested in chapter 12.

2.3.2 Cboe SKEW at 143.6 is the 73rd percentile, and the index measures risk-neutral skewness rather than the cost of a hedge

The Cboe SKEW Index stood at 143.6, the 73rd percentile of the last ten years against a mean of 136.1, with the VVIX at 92.9, and Figure 2.3 sets both against the level. SKEW is the reading most often misused. It is a normalised measure of the asymmetry of the 30-day risk-neutral distribution — a shape statistic — and because it is normalised it says nothing about what protection costs, which is governed by the level of implied volatility. A high SKEW alongside a below-average VIX describes a left tail that is fat relative to its own body, not an expensive hedge: the same shape at 15.83 and at 30 implies very different premiums. No percentile was computed for the VVIX, so none is claimed for 92.9.

What is documented is where August 2026's demand sat, and it was not in downside protection. The VIX "rose a full point in the midst of Tuesday's breakout 1.8% rally"; a record 4,023,796 S&P 500 index (SPX) calls traded on 4 August 2026; and the put-to-call ratio fell to 0.83, the second-lowest reading on record[29][30]. Tanvir Sandhu of Bloomberg Intelligence attributed it directly: "The options market is pricing FOMO… strong demand for upside calls is keeping implied volatility firm even as equities rally"[31]. For a programme short the wings on both sides, that composition matters more than the level: premium concentrated in upside calls raises assignment risk into a melt-up rather than paying for the downside that the SKEW reading describes.



Source: Yahoo Finance daily bars via the marketdata tool, 2016-08-19 to 2026-08-18. Author's calculations.

Figure 2.3 — SKEW at the 73rd percentile against a VIX at the 40th: distribution shape and the price of volatility are separate readings, not moving together.

Describing where the complex sits does not establish whether the level is justified. The next chapter regresses the VIX on realised volatility, credit spreads, financial conditions, the yield-curve slope, policy uncertainty, jobless claims and the equity drawdown; the residual from that fit, not the percentile of the level, is what settles whether index volatility is adequately paid.

---

3. An Economic Fair Value for the VIX

A VIX of 15.83 against a ten-year mean of 18.62 is a description, not a verdict. The level is low, but the conditions that usually accompany low volatility — a shallow equity drawdown, subdued realised volatility, a labour market that has not yet broken — are also present, and a reading that is low for the decade may be entirely ordinary for the week. This chapter separates those two statements by fitting the VIX to observable economic conditions and asking what those conditions, taken together, would justify.

3.1 Specification

3.1.1 Log VIX is regressed on realised volatility, the Baa credit spread, financial conditions, the yield-curve slope, policy uncertainty, jobless claims and the equity drawdown, on 2,493 daily observations

The estimated relationship is a single ordinary least squares regression of the logarithm of the VIX on seven contemporaneous conditioning variables, fitted on 2,493 daily observations. The dependent variable and four of the regressors enter in logarithms, which makes their coefficients elasticities; the remaining three enter in levels, in their native units.

$$\ln \mathrm{VIX}_t = \beta_0 + \beta_1 \ln \mathrm{RV}_t + \beta_2 \ln \mathrm{CS}_t + \beta_3 \mathrm{NFCI}_t + \beta_4 \mathrm{SLOPE}_t + \beta_5 \ln \mathrm{EPU}_t + \beta_6 \ln \mathrm{CLAIMS}_t + \beta_7 \mathrm{DD}_t + \varepsilon_t$$

Here RV is S&P 500 realised volatility, CS the Baa corporate credit spread, NFCI the Chicago Fed financial conditions index — scaled so that higher readings denote tighter conditions — SLOPE the ten-year minus two-year Treasury spread in percentage points, EPU an economic policy uncertainty index entered as a 21-day average, CLAIMS initial jobless claims, and DD the drawdown of the S&P 500 from its running peak, expressed as a negative fraction. Substituting the fitted values gives the working form used throughout this study.

$$\widehat{\ln \mathrm{VIX}_t} = 1.202 + 0.358 \ln \mathrm{RV}_t - 0.113 \ln \mathrm{CS}_t - 0.464 \,\mathrm{NFCI}_t + 0.060 \,\mathrm{SLOPE}_t + 0.087 \ln \mathrm{EPU}_t + 0.135 \ln \mathrm{CLAIMS}_t - 2.958 \,\mathrm{DD}_t$$

Table 3.1 — Fair-value VIX regression: full specification, all coefficients

Term	Entered as	Coefficient
Constant	—	+1.202
Realised volatility (S&P 500)	log	+0.358
Baa credit spread	log	-0.113
Chicago Fed financial conditions index	level	-0.464
10-year minus 2-year Treasury spread	level	+0.060
Economic policy uncertainty, 21-day average	log	+0.087
Initial jobless claims	log	+0.135
S&P 500 drawdown from running peak	level	-2.958

Two coefficients carry signs opposite to naive expectation and should not be glossed over. The Baa credit spread enters negatively (-0.113) and the financial conditions index enters negatively (-0.464), so that on the partial derivative, wider credit spreads and tighter financial conditions map to a *lower* fitted VIX once realised volatility and the drawdown are already in the equation. This is

what a partial coefficient in a collinear system looks like: credit spreads, financial conditions and equity volatility move together, realised volatility and the drawdown absorb the shared component first, and what remains for the credit and conditions terms is a residual that need not carry the unconditional sign. The drawdown coefficient (-2.958) is by far the largest in absolute terms, and with the drawdown measured as a negative fraction it means a deeper decline from the peak raises the fitted level sharply. The equation is therefore best read as a joint mapping, not as seven separable economic channels.

3.1.2 The full model explains 79.8% of the variation in log VIX; the macro-only variant, excluding realised volatility, still explains 70.6%

In-sample R-squared for the full specification is 0.798. A restricted variant that drops realised volatility and retains only the six macro and market-stress terms achieves 0.706. The difference between the two is the marginal explanatory content of contemporaneous realised volatility once the macro block is already present, and it is modest.

That is the more interesting of the two numbers. Realised volatility is mechanically the closest thing to the VIX in the regressor set, and a specification in which it dominated would be near-tautological: implied volatility tracks realised volatility, therefore implied volatility is fair. The macro-only variant removes that circularity and still accounts for slightly over seven-tenths of the variation in log VIX, which is the evidence that the exercise is measuring something economic rather than restating the volatility series against itself. Both figures are in-sample; neither is a held-out result.

3.1.3 Root mean squared error is 3.53 VIX points, which is the honest precision of any such exercise

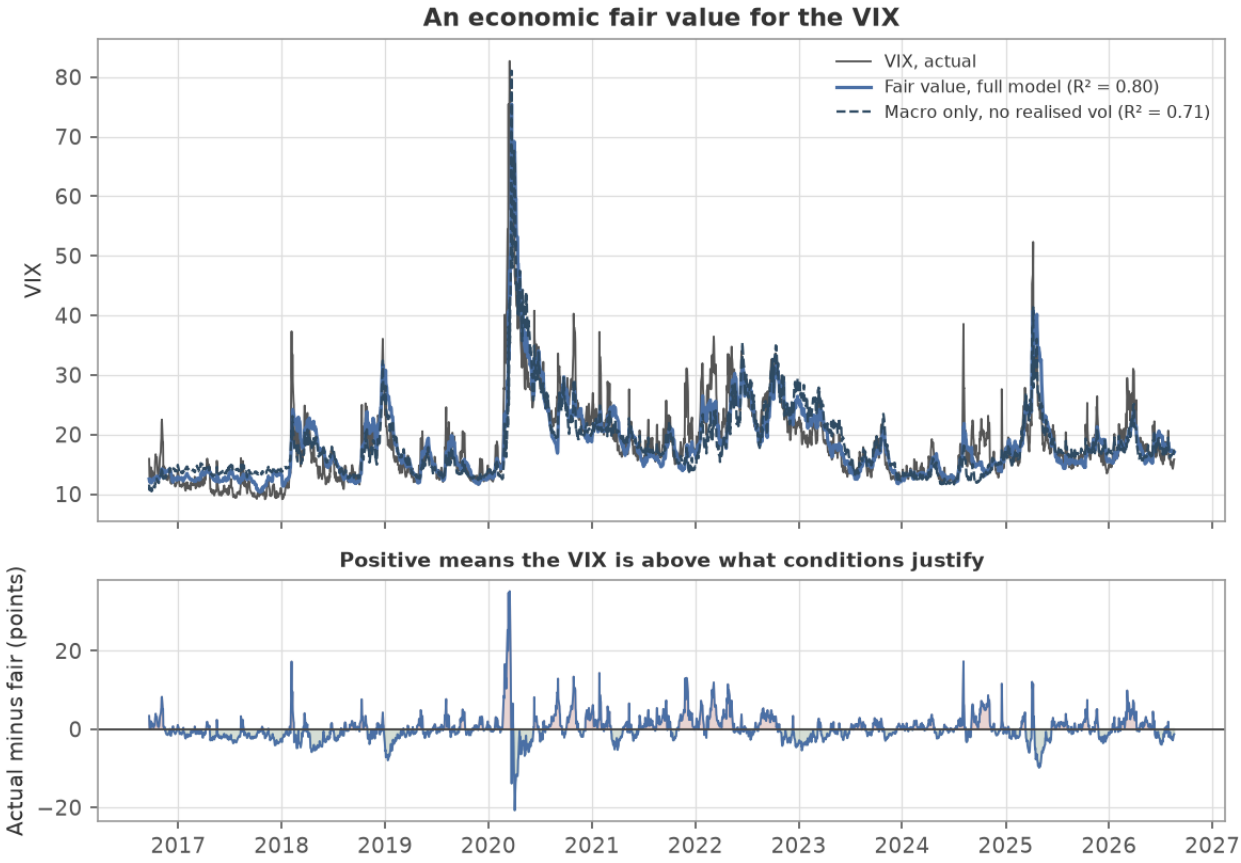
Root mean squared error on the fitted sample is 3.53 VIX points. This is the single most important number in the chapter, and it should be carried alongside every fair-value estimate that follows. A model that explains 79.8% of the variation in log VIX still misses the level by an average of over three points, because the residual variance of a series that ranges from the low teens to above 60 is large in absolute terms even when it is small in proportional terms.

The practical consequence is a floor on resolution. Any gap between actual and fitted VIX that is small relative to 3.53 points is inside the model's own noise and cannot be distinguished from estimation error. Stating a fair value to two decimal places, as this study does for reproducibility, should not be mistaken for stating it to two decimal places of precision. The estimate is a centre of a wide distribution, not a target.

3.2 What the model says today

3.2.1 Fair value on 18 August 2026 is 17.10 against an actual 15.83: the VIX is 1.27 points cheap, a residual of -0.52 standard deviations

Evaluated at the 18 August 2026 close, the full specification returns a fair value of 17.10 against an actual VIX of 15.83. The VIX is -1.27 points from fair, and the corresponding residual is -0.52 standard deviations. Because the regression is estimated in logarithms, the standardised residual is measured in log space and the 1.27 points is the level difference after exponentiating; the two are not linear rescalings of one another and should be quoted as the separate quantities they are.



Sources: Yahoo Finance daily bars, Cboe volatility indices, and FRED macro series, all to 2026-08-18. Author's calculations.

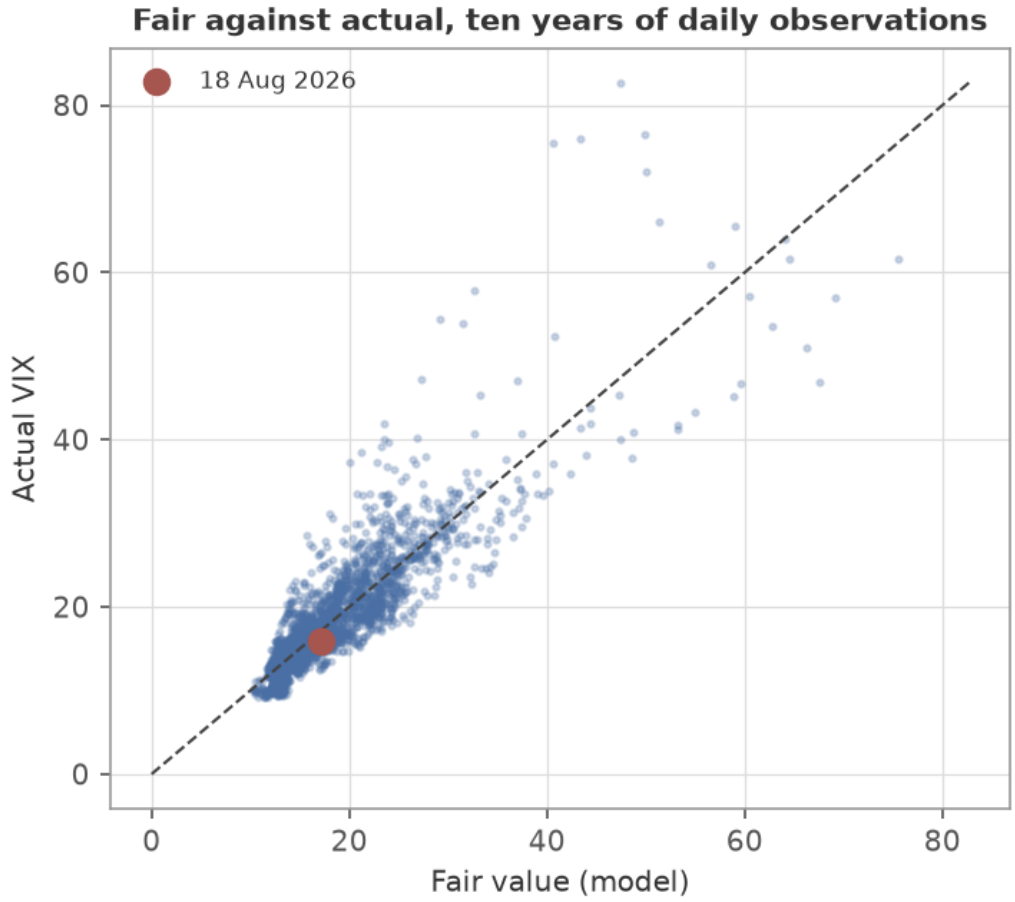
Figure 3.1 — Fitted against actual VIX over the estimation sample: the model captures the large moves and leaves its largest proportional errors in quiet markets.

The fit in Figure 3.1 has the shape typical of this class of model. Sharp volatility episodes are tracked because the drawdown and realised volatility terms move violently at the same time as the VIX; the persistent errors accumulate in calm stretches, where the macro block varies little and the VIX drifts on flow and positioning that the regressors do not observe. August 2026 sits in exactly such a stretch, which is a reason for caution about the current reading rather than confidence in it.

3.2.2 The macro-only variant puts fair value at a higher level still, so the cheapness is not an artefact of low realised volatility alone

The macro-only variant returns a fair value of 17.26, above the full model's 17.10. The direction matters. Realised volatility on the S&P 500 is 11.6% on a 21-day Yang-Zhang basis, which is low, and the obvious objection to calling the VIX cheap at 15.83 is that it is simply following realised volatility down. Removing realised volatility from the specification tests that objection directly: if the

low realised reading were the source of the apparent discount, deleting it would raise the fitted level only trivially or lower it. Instead the macro block alone, without any volatility input, calls for a higher level than the full model does.



Sources: Yahoo Finance daily bars, Cboe volatility indices, and FRED macro series, all to 2026-08-18. Author's calculations.

Figure 3.2 — Actual against fitted VIX across the sample: the current observation lies below the forty-five degree line, but well inside the cloud of ordinary residuals.

Figure 3.2 places today's observation in the context of every other day in the sample. It is below the diagonal, which is the definition of a negative residual, but it sits within the dense body of the scatter rather than at its edge. The discount is a real feature of the current configuration and it survives the removal of the volatility term; it is not, on this evidence, a large one.

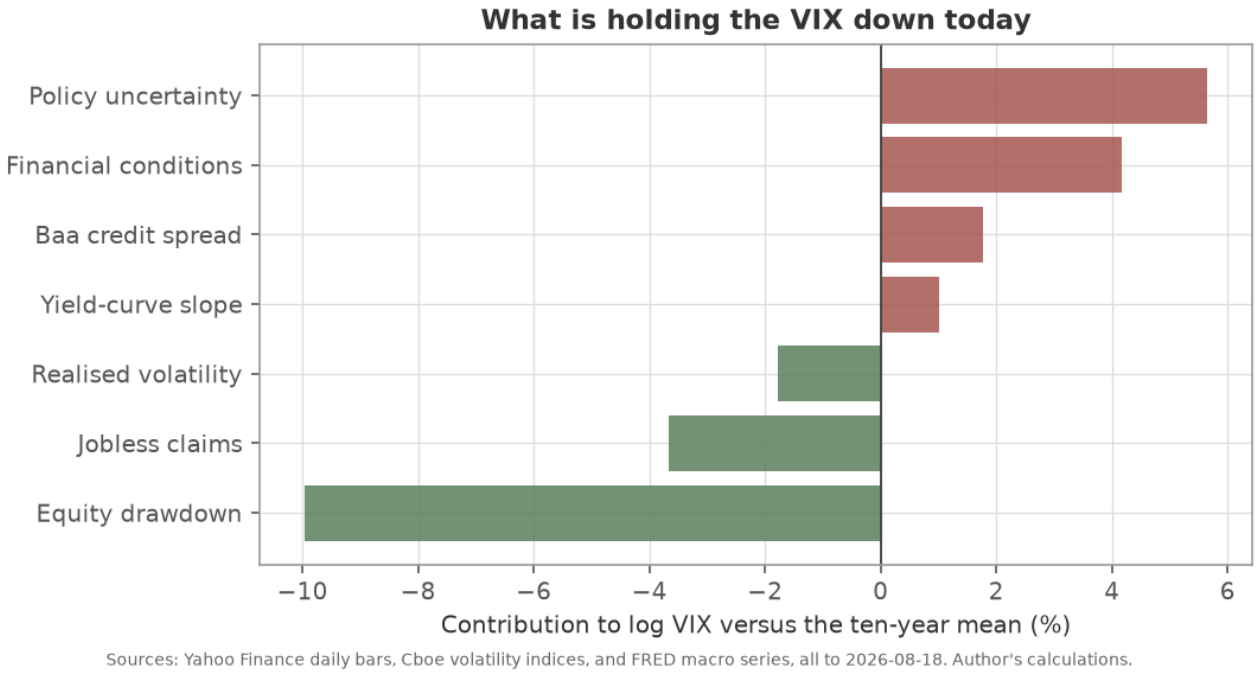
3.2.3 Decomposition: which drivers are holding the fitted level down relative to the ten-year mean

The fitted level can be decomposed into each driver's contribution relative to its sample mean, expressed as a percentage effect on fitted VIX.

Table 3.2 — Contribution of each driver to the fitted log VIX on 18 August 2026, relative to the sample mean

Driver	Contribution to fitted VIX vs sample mean
Equity drawdown	-10.0%
Initial jobless claims	-3.7%
Realised volatility	-1.8%
Yield-curve slope	+1.0%
Baa credit spread	+1.8%
Financial conditions	+4.2%
Policy uncertainty	+5.6%

The dominant term is the equity drawdown, at -10.0%, and it is larger in magnitude than any positive contribution. With the S&P 500 only 1.37% below its running peak, the market is closer to its high than it typically is, and the model reads proximity to the peak as the single strongest argument for a low VIX. Benign initial jobless claims add a further -3.7%. Realised volatility, despite carrying the largest elasticity among the logged terms at +0.358, contributes only -1.8%, because the term is logged and today's realised reading is not extreme relative to its own history. Working the other way, policy uncertainty at +5.6% and financial conditions at +4.2% are arguing for a materially higher VIX than the market is paying for, with the credit spread and the curve slope adding smaller positive amounts.



Sources: Yahoo Finance daily bars, Cboe volatility indices, and FRED macro series, all to 2026-08-18. Author's calculations.

Figure 3.3 — Driver decomposition: the shallow equity drawdown pulls the fitted level down, while policy uncertainty and financial conditions push it up.

The two-sided shape visible in Figure 3.3 is the substance of the result: this is not a market in which every indicator points the same way, but one in which a benign equity and labour picture is outweighing an unsettled policy picture. The policy uncertainty contribution has an identifiable source. The Federal Open Market Committee held its target range at 3.50–3.75% on 29 July 2026 on a 9–3 vote, with all three dissenters preferring a quarter-point *increase*.[32] Chair Kevin Warsh has argued for giving markets fewer signals about the next move, a deliberate reduction in forward guidance.[33] Implied odds of a September hike stood at roughly 82% on 23 July and 35% by 13 August, a 47-point round trip in implied hike probability inside three weeks.[34] The elevated policy uncertainty term in Table 3.2 is that instability, entering the model as a reason for the VIX to be higher than it is.

3.3 How to use it, and how not to

3.3.1 A 0.52 standard deviation discount is inside the normal range and is not a trading signal on its own

A residual of -0.52 standard deviations is a routine reading. Under any roughly symmetric distribution of residuals, a deviation of that size is reached or exceeded on a large fraction of all days, and the 1.27-point level gap is smaller than the model's own root mean squared error of 3.53 points. On both measures the current discount fails to clear the model's noise floor.

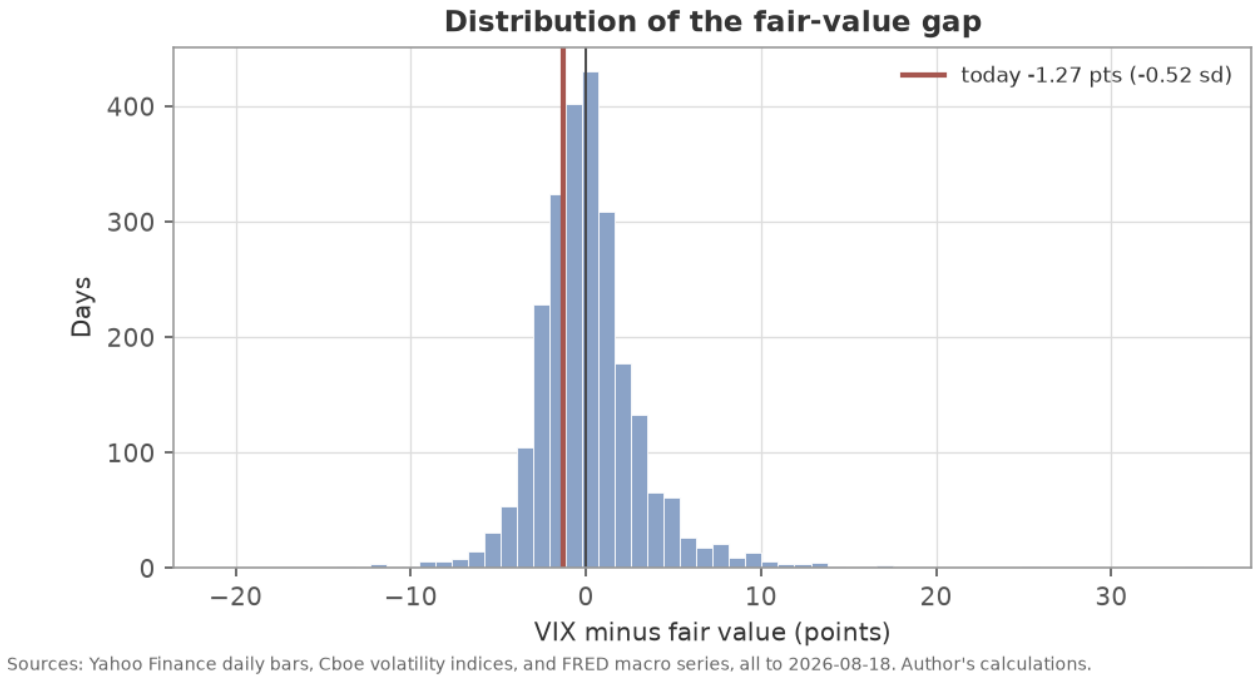


Figure 3.4 — Distribution of the historical fair-value gap: today's -1.27 point discount sits inside the central mass, not in either tail.

Figure 3.4 makes the point visually. Treating a half-standard-deviation residual as an entry trigger would generate signals on a substantial share of all trading days, which is another way of saying it would not be a signal at all. The correct use of the number is as context for position sizing and structure selection: a modest discount argues against paying up for long volatility exposure and marginally against aggressive premium selling, and it does not argue for either with any force. Named desks reached a directionally similar view from different evidence in mid-August 2026, with BTIG's Jonathan Krinsky observing that the market was entering a seasonally weak window at all-time highs with the VIX at its low for the year.[35] That is a judgement about complacency, not a measurement, and it is offered here as corroboration rather than as input.

3.3.2 The model is contemporaneous, not predictive: it says what conditions justify, not what happens next

Every regressor in Table 3.1 is measured on the same day as the dependent variable. The model therefore answers one question — what level of implied volatility is consistent with today's observed economic and market conditions — and it is silent on the question of what volatility will do tomorrow. A negative residual does not mean the VIX will rise; it means the conditions on the same day were, on this fit, associated with a higher level.

Three limitations should be stated plainly. First, the R-squared of 0.798, the R-squared of 0.706 and the root mean squared error of 3.53 points are all **in-sample** statistics from a fit estimated over the whole 2,493-observation window; there is no held-out period, no walk-forward estimation and no out-of-sample validation of the specification anywhere in this chapter. Second, and more importantly, **the fair-value gap has not been tested as a trading signal** — this study contains no backtest in which positions are taken on the basis of the residual, and consequently no evidence at all about whether the gap predicts subsequent VIX changes, subsequent option returns, or anything else. Any such use would be an extrapolation beyond what has been measured here. Third, the coefficients are fitted over a decade in which the behaviour of volatility itself changed: spikes that once occurred every three to five years now appear annually, and the time for volatility to subside has roughly halved, from around 25 days to nine.[19] A relationship estimated across both halves of that shift is an average of two regimes rather than a description of the current one.

Known event risk sharpens the caveat. The remaining 2026 policy calendar has decisions on 16 September, 28 October and 9 December, with the September and December meetings carrying a Summary of Economic Projections and a press conference.[36] A programme trading 7 to 12 day structures will straddle these dates repeatedly, and a reduced-guidance communication regime concentrates more of the reaction into the hours after the announcement. The fair-value model has no term for a scheduled event; the calendar is an input the reader must supply. The same holds for seasonality: in every mid-term election year since 1990, the equal-weight S&P 500 fell at least 7% from its 18 August average peak into mid-October.[35]

The distinction that carries into the next chapter is one of benchmark. This chapter compared implied volatility to *contemporaneous macroeconomic conditions* and found it 1.27 points below

what those conditions justify. Chapter 4 compares implied volatility to *subsequently realised volatility* — the variance risk premium — which is a forward comparison against an outcome that had not yet occurred when the option was priced, and which is the quantity a premium seller is actually paid. The two can disagree without contradiction: volatility can be cheap against the macro backdrop and still expensive against what the market goes on to deliver, and it is the second gap, not the first, that determines whether the strategies in this study earn anything.

---

4. The Variance Risk Premium and How These Strategies Exploit It

Two distinct claims about the price of volatility run through this report, and conflating them misreads everything that follows. The preceding chapter measured the VIX at 15.83 against an economic fair value of 17.10 — a discount of 1.27 points, a residual of -0.52 standard deviations against a model root mean squared error of 3.53 VIX points. That is a statement about the *level* of implied volatility relative to what macro and financial conditions have historically justified. The variance risk premium is a different quantity: the gap between option-implied variance and the variance that subsequently occurs. One compares today's implied level to a fitted conditional mean of itself, the other compares implied variance to the realised outcome; the two are independent. Today they point in opposite directions — implied volatility sits below its macro fair value while still exceeding realised volatility, which stands at 11.6% on a 21-day Yang-Zhang basis. Chapter 3 asked whether premium is cheap; only this chapter's question — whether premium is systematically overpriced relative to outcomes — determines whether a short-premium programme has an edge.

4.1 The index premium is large and well documented

4.1.1 Roughly four volatility points on the S&P 500 across three independent sources, and about 18 in monthly variance units

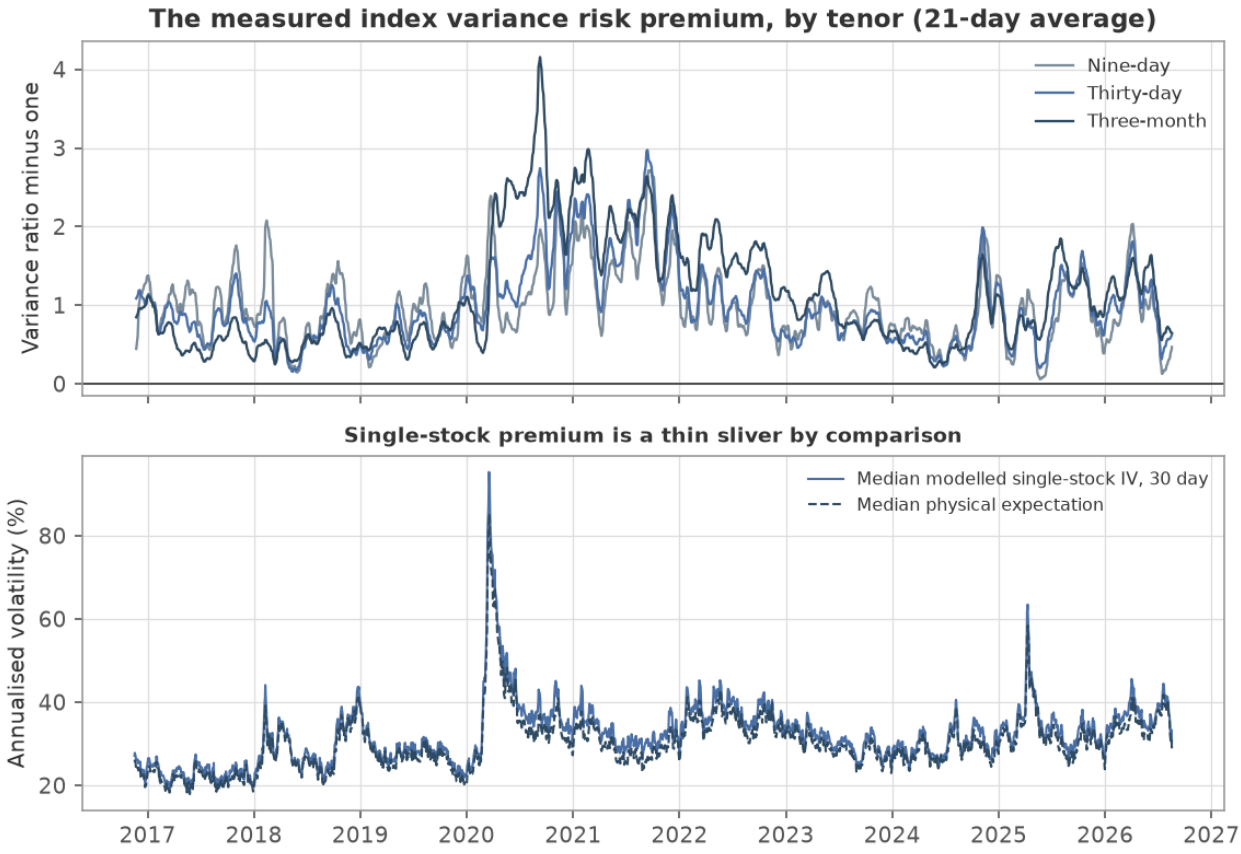
At index level the premium is well evidenced. Cboe, over 1990 to 2018, reports an average VIX of 19.3% against average S&P 500 realised volatility of 15.1%, a difference of 4.2 percentage points[37]. A CFA Institute study spanning 35 years finds average 30-day-forward realised volatility of 15.50% against an average VIX of 19.59%, a spread of 4.09%[38]. Cboe's strategy work puts the long-run average at 4.4 volatility points, about 28% above subsequent realised since 1990 and 5.0 points or 31% since the March 2009 low[39]. In variance units, Bollerslev, Tauchen and Zhou report mean implied variance of 33.23 against mean realised variance of 14.93 over January 1990 to December 2007, a premium of 18.30 in percentages squared on a monthly basis[40]. Two cautions attach. These are not competing estimates of one parameter: the same means convert to roughly 6.6 annualised volatility points rather than 4.2, on five-minute realised variance and a sample ending in 2007 — a conversion made in the research note, not by the authors. And the premium is an average, not a per-period edge: the VIX explains only 20% to 25% of subsequent 30-day volatility by adjusted R-squared[41].

4.1.2 Measured here at a variance ratio of 0.877 at 30 days, 0.918 at nine days and 0.740 at three months

Measured against the Cboe indices over the fitting window, the premium declines monotonically with tenor. The values are variance ratios less one — the excess of implied over expected realised variance — so +0.877 means implied variance at thirty days ran 87.7% above expected realised variance.

Tenor	Cboe index	Variance ratio less one, median 2016–2021	Reconstruction correlation, 2021–2026	RMSE, points	Bias, points
Nine day	VIX9D	+0.918	0.810	3.97	+0.60
Thirty day	VIX	+0.877	0.818	3.15	+0.12
Three month	VIX3M	+0.740	0.762	3.43	-1.70

The two right-hand blocks cover different windows. The premium column is a median over 2016 to 2021, on which the pooled log-HAR volatility model was fitted before being frozen; the reconstruction columns are its out-of-sample test over 2021 to 2026, rebuilding each Cboe index from S&P 500 prices alone across n = 1,254 daily observations, with an R-squared of +0.642 at thirty days. The bias column is the one a short-dated seller should read carefully: the reconstruction runs +0.60 points rich at nine days and -1.70 points cheap at three months, so the modelled surface tilts in the direction that flatters short-dated selling and understates a calendar's back leg. Almeida and co-authors corroborate the gradient, finding the annualised 0DTE premium ranges from 1.54% to 2.95% against 0.55% at one day and 0.81% at 22 days[42]. Figure 4.1 plots the tenor profile; the mandated 7 to 12 day window sits between this study's nine-day and thirty-day measurements.



Sources: Yahoo Finance daily bars, Cboe volatility indices, and FRED macro series, all to 2026-08-18. Author's calculations.

Figure 4.1 — The measured index premium falls with tenor, standing furthest above expected realised variance at nine days and closest at three months.

A median conceals the episodes that matter most to a seller, because the premium inverts in crises. In October 2008 realised 21-day S&P 500 variance reached 0.6939 against a squared VIX high of 0.6410; in March 2020 realised variance reached 0.9064 against squared VIX of 0.6838[43]. OptionMetrics dates four breaches of a -0.2 threshold — the 2002 selloff, the 2008 crisis, the 2011 sovereign debt crisis and Covid-19 — bottoming on 17 April 2020 at -0.59[44]. Bekaert and Hoerova reject the benign reading that these are surges in risk appetite: when realised variances peak, the premium simply turns negative[45]. The fitting window used here contains March 2020, and negative readings grew more frequent at the shortest tenors after 2022[42].

4.2 The single-stock premium is a fraction of it

4.2.1 The correlation risk premium explains the gap: index options embed a payment for correlation that single-stock options do not

An index option is a diversified claim whose payoff depends on constituents moving together, so it hedges the loss of diversification in a crash. A single-stock option does not, and is not paid for it. Driessen, Maenhout and Vilkov formalise this: if individual variance risk is essentially unpriced,

index variance carries a premium only to the extent correlation risk is priced, and their model attributes the entire index premium to it[46]. The wedge is quantified — over 1996 to 2003 average option-implied S&P 100 correlation was 46.7% against realised 28.7%, so index options priced correlation 62.8% higher than history suggested[47]. The consequence for August 2026: the input driving the index premium is depressed, with mean pairwise correlation across the 40-name universe at 0.094 against a ten-year mean of 0.317, near the bottom of a range topping out at 0.694.

4.2.2 Liquid large caps sit at the bottom of the single-stock cross-section, with implied-to-realised ratios of roughly 1.03 to 1.07

The direction of the single-stock premium is agreed; its magnitude is not. Carr and Wu, on 35 actively traded stocks, find premia negative but smaller, with a stark significance gap: the mean log premium is significantly negative for 21 of 35 stocks, but in levels it is insignificant for all but three, against strongly significant index statistics[48]. Narrowing to liquid large caps narrows the premium further: delta-hedged equity option returns fall monotonically with idiosyncratic volatility, a bottom-minus-top quintile trade earning about 1.4% per month, and are more negative still when the stock or its options are less liquid[49]. The 40 names traded here are the most liquid, lowest-idiosyncratic-volatility segment of the cross-section, precisely where the premium is thinnest.

Some names carry no premium at all, or a negative one. González-Urteaga and Rubio, across 20 portfolios sorted from 181 S&P 100 components over January 1996 to February 2011, report average premia from 0.103 to -0.035 annually on a realised-minus-implied definition; at the positive extreme realised volatility exceeded implied across the whole sample, a standing loss for the seller[50]. Dispersion of that width is not a rounding error for a programme selling single-name premium one position at a time.

The disagreement over the level has a locatable cause, and it deserves stating rather than resolving by silent choice. Driessen, Maenhout and Vilkov compute model-free implied variance from options with delta above 0.15, omitting the far wings of the smile; on that basis the null of no premium is not rejected at the 5% level for 108 stocks, only 16 showing a significantly positive gap between implied and realised variance — which licenses attributing the whole index premium to correlation[51]. Duarte, Jones and Wang replicate that filter, reproduce the null, then show it depends on the filter: widening the strike range gives -1.6% in annualised variance at a monthly horizon, and an independent model-free dataset -1.4%. Their Fama-MacBeth estimates put the premium in individual equity options near that in S&P 500 index *calls*, both about -5 basis points per day, while index *puts* price volatility at -11 to -15[52]. The choice between zero and small-but-real is therefore about strike coverage and microstructure noise, not sample period. Weighting is difficult: the first is peer-reviewed in the *Journal of Finance* and the second unpublished, but the second diagnoses the first rather than merely contradicting it. No source consulted states a single-stock-to-index ratio as a headline number; every such ratio, this study's included, is a construction.

4.2.3 This study calibrates to a median modelled ratio of 1.060 against 1.364 for the index, and that choice governs every result that follows

The single-stock premium parameter kappa is set to 0.12, producing a median modelled implied-to-expected-realised ratio of 1.060 against 1.364 for the index, inside the 1.03 to 1.07 range defensible for liquid large caps. Two qualifications matter. These are modelled implied volatilities, not observed quotes — no historical option quotes were available to this study, so the index premium above is *measured* while the single-stock premium is *assumed*. And the 1.060 covers diffusive variance only: the earnings jump premium is applied separately, at a variance ratio of 0.23 with implied jump volatility 1.108 times realised announcement-day volatility. Figure 4.2 shows the dispersion behind that median.

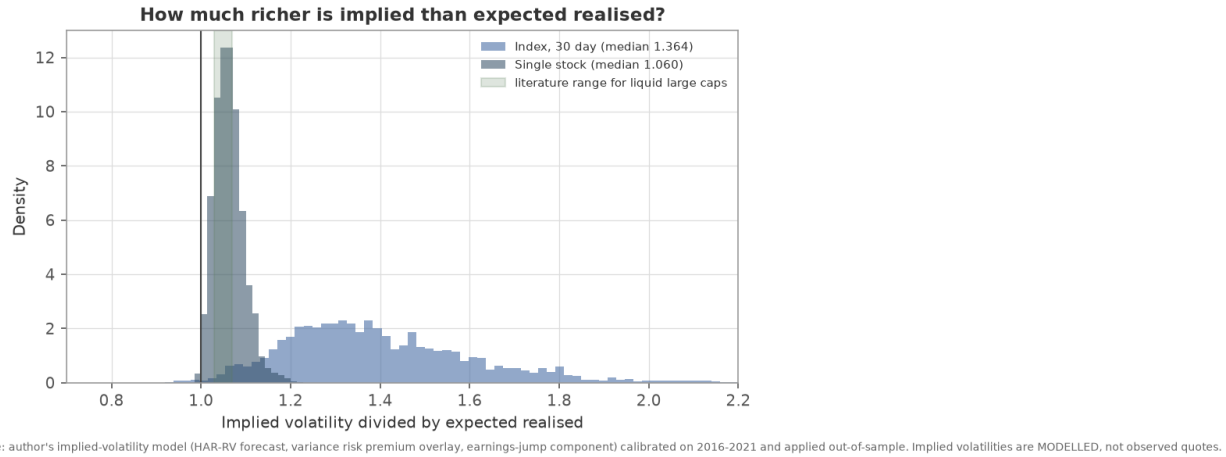


Figure 4.2 — The modelled implied-to-expected-realised ratio across the 40-name universe: a median of 1.060 against 1.364 for the index.

Assumed single-stock IV / expected RV	Iron condor	Calendar	Double calendar
1.061 (calibrated)	+4.12	+3.57	+2.48
1.128	+5.58	+2.50	+1.82
1.262	+11.85	+2.84	+0.79

Compound annual growth rate in percent under the managed rule; the grid labels the calibrated base row 1.061. This is the study's dominant sensitivity, and where the tension lands. The condor nearly triples, from +4.12% to +11.85%, without a single price path changing — the whole difference is the assumed richness of the options sold, so its headline result is less a measurement of an edge than a restatement of the calibration. Push the assumption the other way, to the zero that Driessen, Maenhout and Vilkov's stock-level tests cannot reject, and no wedge remains — only a round-trip cost of $18.88 against a median premium of $247, or 7.0% of

the credit. That is the honest reason the results that follow are thin rather than compelling — Sharpe ratios of 0.84, 0.61 and 0.52, and at twice the base slippage assumption a condor earning +0.85% while both calendars lose money. This programme is not selling the four-point index premium the literature documents; it is selling a fraction of it whose existence is disputed.

4.3 What each structure actually harvests

4.3.1 The condor is short variance outright and collects the premium directly through negative vega and positive theta

The iron condor's exposure is the simplest of the three and the most directly geared to that wedge. It is short variance outright — negative vega and positive theta throughout — with profit accruing as premium collected less variance actually realised. It therefore has no defence if the single-stock premium is absent, and is the only structure here whose return rises steeply with the assumed premium. The same sign governs its other exposures: trades spanning a scheduled report returned -2.47% of capital at risk against +2.97% for those that did not, because the event delivers realised variance — a mean absolute move of 5.48% against 1.38% on ordinary days — while its compensation is priced separately and only partly.

4.3.2 The calendar is long forward variance and short spot variance: it harvests the term structure, not the level

The calendar and double calendar are exposed to a different quantity. Short the front leg and long the back, the position is long forward variance and short spot variance: its edge comes from the slope of the term structure, not the height of the premium. The backtest confirms the sign directly: the calendar returned +10.7% mean return on risk when the curve was inverted at entry against -7.3% in steep contango, and the double calendar +6.8% against -4.5%. The same logic explains earnings behaviour that is the reverse of the condor's — +18.59% against -0.00% for the calendar, +17.13% against -2.22% for the double calendar. It also explains their flat response in the table above: raising the assumed premium enriches the back leg bought as well as the front leg sold, and for the double calendar the net effect is negative, from +2.48% to +0.79%. That asymmetry is the practical significance of this chapter's tension. If the single-stock premium is nearer zero than the calibration, the condor's case collapses while the calendars' does not, because a term-structure trade needs only the shape of implied variance to be wrong, not its level to be rich. What that shape looks like when the entry screen fires is the subject of the next chapter.

The abstract claim that one structure is short variance while the others are long it can be checked directly against the positions the backtest actually opened. Taking the median Greek across every trade at the moment of entry, the iron condor carries vega of -6.93 dollars per volatility point against theta of +13.43 dollars per day, while the calendar carries vega of +10.09 and the double calendar +19.74. All three are short gamma, at -3.06, -1.79 and -3.42 respectively. The condor is therefore paid to be short the level of implied volatility; the calendars pay to be long it, and must recover that outlay from the shape of the curve and from decay rather than from the premium

itself.

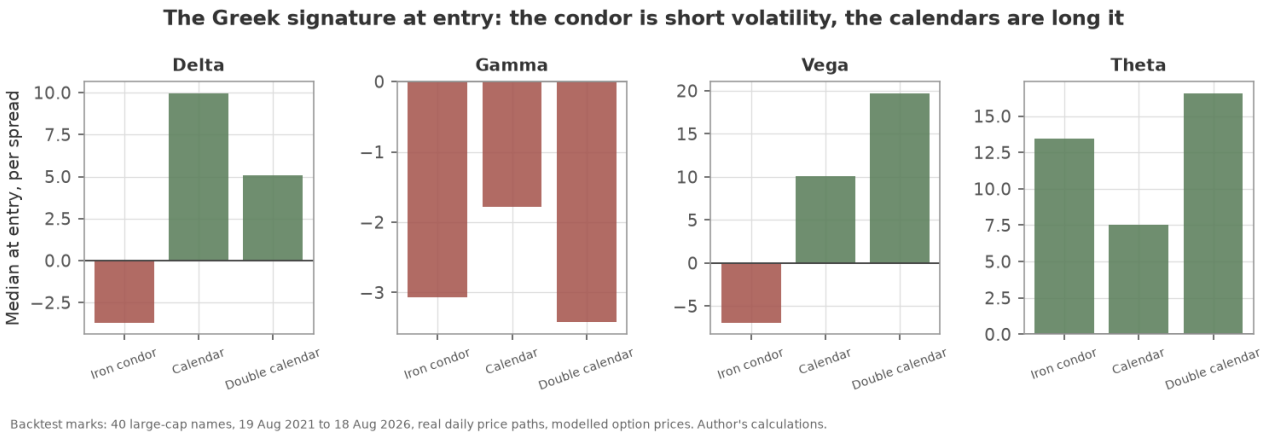


Figure 4.3 - The Greek signature at entry, median across every backtested trade. The sign of vega separates the condor from the two calendars and determines which of them benefits from the premium described in this chapter.

---

5. Skew, Term Structure and Their Evolution

The mandated specification fixes where on the volatility surface each trade sits, not what the surface looks like there. Two slopes set the entry price: the slope across strike, which prices a 30-delta option against the money, and the slope across time, which prices the front expiry against the back. The iron condor is exposed to the first, the calendar to the second, and every implied volatility in this chapter is modelled, not quoted.

5.1 The shape of the smile

5.1.1 Equity skew is a monotone decreasing smirk, and index skew is far steeper than single-stock skew: mean risk-neutral skew of -1.6 against -0.24

Post-1987 equity implied volatility falls monotonically as the strike rises: the curve is a decreasing smirk, not a symmetric smile, and out-of-the-money puts are dearer in volatility terms than calls. The slope "is significantly negative for both individual stocks and index options, and the slope is less negative for longer-term options"[53]. A negatively sloped curve maps one-to-one onto negative risk-neutral skewness, and measured that way index and constituents are not the same object: mean risk-neutral skew is about -1.6 for S&P 500 index options against -0.24 for individual stocks, and an individual stock's skew is the less negative of the two 92% of the time[54]. Bakshi, Kapadia and Madan find the same ordering on the OEX and 30 of its components[55]. Single names therefore carry roughly one-quarter to one-third of index steepness.

The cause is aggregation. Because the index is a weighted sum, "the difference between the index smile and the smiles for individual stocks is then determined entirely by the dependence structure among the stocks"[56]. Idiosyncratic skewness diversifies away under aggregation;

systematic downside jump risk does not, and correlation itself rises in the left tail. Mean pairwise correlation across the 40 names is 0.094 against a ten-year mean of 0.317, so on that mechanism the index-to-single-stock wedge should sit near its wider end — an implication of the argument, not a quantity measured here. The Cboe SKEW index closed at 143.6 on 18 August 2026, the 73rd percentile of its comparison sample against a mean of 136.1, but SKEW is a linear rescaling of SPX risk-neutral skewness, 100 minus ten times that skewness[57], and it measures exactly the quantity that does not carry across to single names. It is context for the index, not evidence about the strikes traded here.

5.1.2 The smile is parametrised in standardised moneyness, damped so the log-moneyness slope scales as T to the power -0.3 at the short end rather than -0.5

The surface is modelled, not quoted: no option-quote database enters the backtest. Each strike is priced from a parametric smile built on the volatility forecast of chapter 4, and every implied volatility quoted later inherits that construction. The smile is written in standardised moneyness:

$$\sigma(K,T) = \sigma_{\text{ATM}}(T)\left[\,1 - \beta(T)\,z + c\,z^2\,\right]$$

$$z = \frac{\ln(K/F)}{\sigma_{\text{ATM}}(T)\sqrt{T}}$$

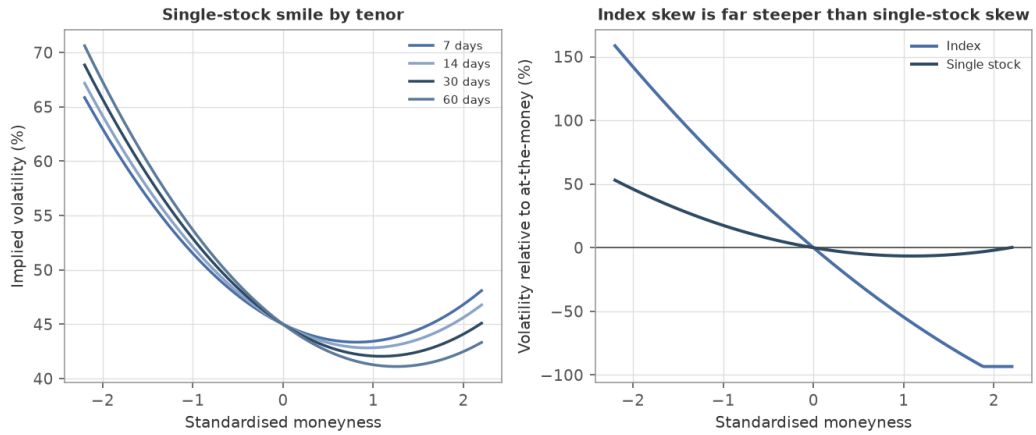
with F the forward, β the skew coefficient and c the curvature. The coefficient β is damped by tenor:

$$\beta(T) = \beta_{30}\left(\frac{T}{30}\right)^{0.20}$$

with T in days. The damping is the substantive choice. Differentiating at the money, the σ_{ATM} terms cancel and the log-strike slope is $-\beta(T)/\sqrt{T}$: a constant β would scale it as $T^{-0.5}$, and the damping combines the exponents as \$0.20 - 0.50\$:

$$\left.\frac{\partial\sigma}{\partial\ln(K/F)}\right|_{K=F} = -\frac{\beta(T)}{\sqrt{T}} \;\propto\; T^{-0.30}$$

The target exponent comes from the literature, not from a fit here. Bergomi's log-log fit of the 95/105 skew from three months to five years is "almost a straight line with slope about 1/2 (-0.51)"[58]. But the exponent breaks at the short end: using nine years of SPX data including weeklies, Delemotte, De Marco and Segonne confirm the -0.5 law "for longer maturities" while documenting "a change of regime in short-dated implied volatilities," the short end following T^{α_1} with $\alpha_1 \approx -0.3$ and "a separation threshold located around two months"[59]. The mandate trades 7 to 12 days to expiry, inside that regime, so propagating -0.5 into these expiries would overstate short-dated skew and inflate the premium credited to the condor's put wing. Figure 5.1 plots the resulting curves.



Source: author's implied-volatility model (HAR-RV forecast, variance risk premium overlay, earnings-jump component) calibrated on 2016-2021 and applied out-of-sample. Implied volatilities are MODELLED, not observed quotes.

Figure 5.1 — The modelled smile in standardised moneyness: implied volatility falls monotonically as the strike rises, and the tenor damping flattens the curve in z as expiry shortens.

An unbounded quadratic smile eventually prices a deep in-the-money option above its own intrinsic value, so the structure diagrams in chapters 7 to 9 cap standardised moneyness at plus or minus 2.5. That cap never touches the backtest, where every strike sits inside one standard deviation of the forward.

5.2 The term structure through an earnings cycle

5.2.1 Implied variance carries an earnings term that scales as one over time to the announcement, so the front expiry inverts above the back and then collapses

A scheduled report is a dated jump, not a change in diffusion. Total variance over an option's life is the diffusive part plus one jump variance for each scheduled report inside the window, annualised by dividing by time to expiry:

$$\sigma^2_{\text{imp}}(t,T) = \sigma^2_{\text{diff}} + \frac{1}{T-t} \sum_{i\,:\,t<t_i<T} \left(\sigma_J^{(i)}\right)^2$$

where $\sigma_J^{(i)}$ is the jump volatility attached to the report at date t_i. The numerator is fixed once the calendar is known; only the divisor moves. Implied volatility therefore "increases at rate $(T-t)^{-1}$ as the maturity decreases," and "will 'blow-up' if the maturity is just after the earnings announcement"[60]. The front expiry does not merely rise, it inverts above the back, while the same jump variance spread over a 60-day option is barely visible. The collapse afterwards is a discontinuity rather than a decay: it lands at the first post-announcement quote. The worked large-cap case is Intel, reporting after the close on 15 July 1997: the July at-the-money call implied 71.15% against 45.19% for August, a 25.96-point inversion, and fell to 42.96% the next session while August moved 1.27 points[60]. Cross-sectional estimates put implied earnings jump volatility "on the order of 10 percent," with a 75th percentile near 15%.

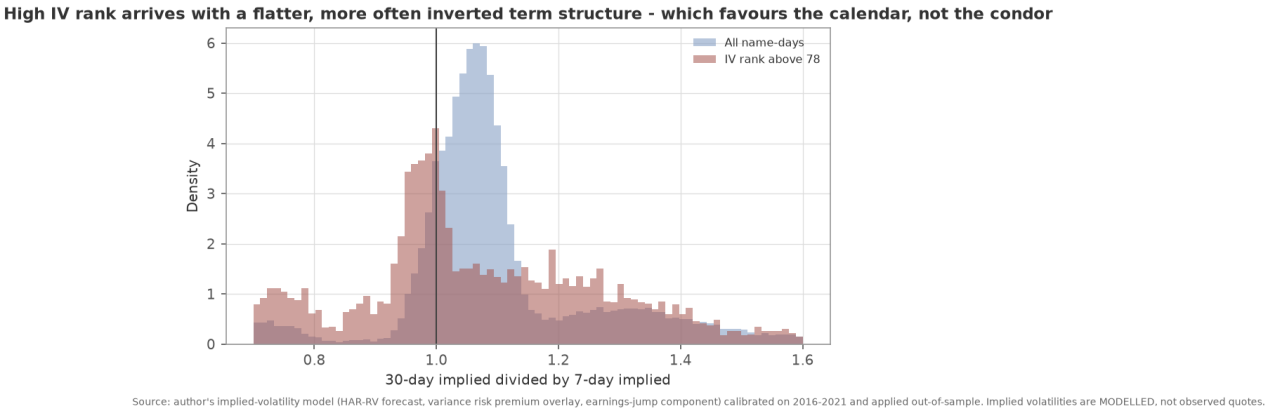


Figure 5.2 — The modelled term structure at trade entry, as the ratio of 30-day to 7-day implied volatility on high-IV-rank days against all name-days; readings below 1.0 are inverted curves. Quantified in 5.3.1.

5.2.2 The modelled event study reproduces this: the 7-day point rises hardest into the report and falls furthest the day after

The study's earnings dates are detected from returns, not imported from a vendor calendar. Detection identifies 1,577 reaction days across the 40 names, a median 1 day from the nearest SEC filing date. On those days the mean absolute move is 5.48% against 1.38% on ordinary days, a ratio of 3.96: the jump is not a refinement of the diffusive process but a different order of magnitude. The event premium is applied separately from the diffusive premium of chapter 4, at a jump variance ratio of 0.23 — an implied jump volatility of 1.108 times realised announcement-day volatility. The event carries a thin markup: the jump is priced close to what it delivers.

Feeding a fixed jump variance through the decomposition reproduces the shape set out in 5.2.1. The 7-day point rises hardest into the report because its divisor is smallest, inverts above the 30-day point, and falls furthest the day after, when the numerator loses the jump term outright; the 30-day point moves a fraction as far either way. Figure 5.3 shows the modelled path. The measured 5.48% mean absolute reaction-day move is close to the 5.5% implied earnings-day move reported for the average S&P 500 stock in July 2026[61], but with no matched implied-versus-realised pair in the sourced literature that is a coincidence of levels, not a test of calibration.

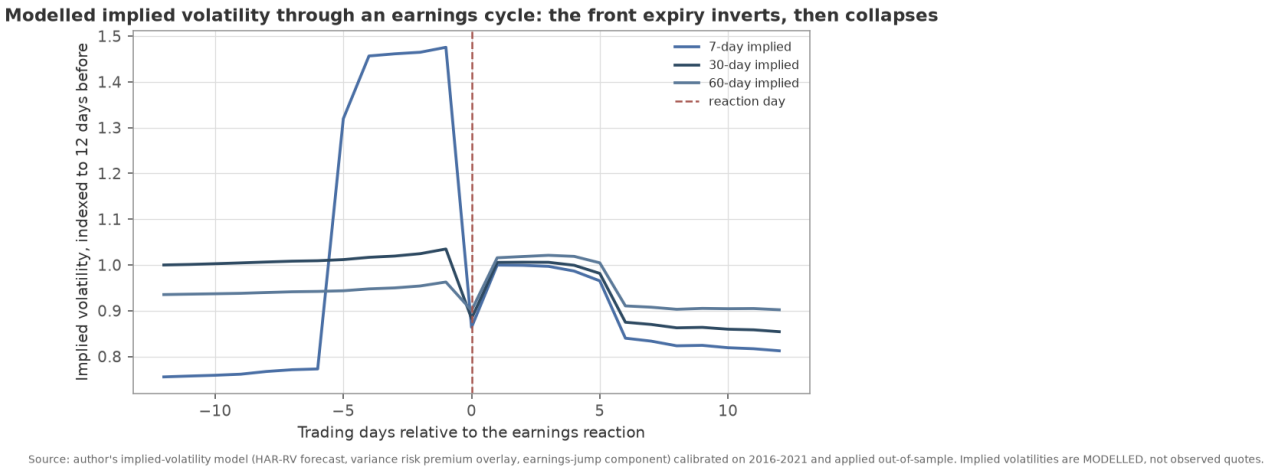


Figure 5.3 — The modelled event study: the 7-day point rises hardest into the report and falls furthest the session after, while the 30-day point barely moves.

5.3 The state of the curve when the screen fires

5.3.1 Median 30-day over 7-day implied is 1.023 when IV rank exceeds 78 against 1.068 unconditionally, and 43.3% of high-IV-rank days are inverted against 18.5% of all days

IV rank — the position of current 30-day implied volatility within its trailing 52-week range, constructed and criticised in chapter 6 — exceeds 78 on 8.85% of name-days. What those days look like on the surface, plotted in Figure 5.2, is not what a level screen advertises.

Measure at entry	IV rank above 78	All name-days
Median 30-day implied volatility	44.8%	33.5%
Median 30-day over 7-day implied	1.023	1.068
Share of days with an inverted curve	43.3%	18.5%
Scheduled report inside the next 12 days	32.6%	15.2%

Table 5.1 — The modelled surface when the screen fires, against all name-days in the backtest window.

Level and slope move in opposite directions. The screen lifts median 30-day implied volatility from 33.5% to 44.8% while flattening the curve from 1.068 to 1.023, and the inverted share rises from 18.5% to 43.3%. The last row explains both: a scheduled report falls inside the next twelve days

on 32.6% of high-IV-rank days against 15.2% of all days, an over-representation of 2.1 times, and by the decomposition of 5.2.1 a report inside the front window is what raises the near point above the far one. The index sits at the other extreme: VIX3M over VIX stood at 1.094 on 18 August 2026, and that cash pair was in contango on 92.3% of the last ten years of days — higher than the counts on the VIX futures curve, Cboe's "more than 80% of the time since 2010"[22] and OptionMetrics' 18.8% backwardated since 2006[23], because the further apart the tenors, the rarer the inversion. The pair here is 7-day against 30-day, the near end, which inverts first and most often; a contango share means nothing without its tenor pair attached.

5.3.2 That flattening is bad for a condor and good for a calendar, and the backtest confirms both signs

The condor sells one expiry; the calendar sells the front and buys the back, so a flat or inverted curve reprices them in opposite directions. The backtest carries both signs.

Strategy	Inverted at entry	Steep contango at entry	Trades spanning a report	With a report	Without
Iron condor	not measured	not measured	28.8%	-2.47%	+2.97%
Calendar	+10.7%	-7.3%	28.3%	+18.59%	-0.00%
Double calendar	+6.8%	-4.5%	28.3%	+17.13%	-2.22%

Table 5.2 — Mean return on risk conditioned on the entry slope and on earnings exposure. Term-structure columns are measured on positions held to expiry, steep contango being a ratio above 1.10; no slope split for the condor was measured.

For the calendars the split is large: entering into an inverted curve turns a losing structure into a profitable one, +10.7% against -7.3% of capital at risk for the calendar and +6.8% against -4.5% for the double calendar. The direction is the claim here — selling the expiry that is rich and owning the one that is not — and why that value accrues, including the vega term-structure exposure driving it, belongs to chapter 8. The condor has no slope split in the study's outputs, so its sign rests on the channel producing the flattening: 28.8% of condor trades span a scheduled report and average -2.47% of capital at risk against +2.97% for the rest. The condor is short the one expiry carrying that jump premium, priced at 1.108 times what it delivers. Index-level evidence agrees: Johnson finds the slope of the volatility term structure "negatively predicts future returns of 18 variance assets," significantly in 33 of 36 tests, so steep contango forecasts good returns to selling volatility and inversion forecasts poor ones[27]. Straddles bought three days before an announcement separately earn "a highly significant 3.34% return," an effect weakest in the liquid

large caps screened here[62].

The screen was specified as a level filter and behaves as a slope filter. The question for the next chapter is therefore not whether IV rank finds expensive volatility — it plainly does, 44.8% against 33.5% — but whether an instrument built from a 52-week high and low is the right way to select a term-structure state it never mentions.

Taken together, the smile and the term structure are one object, and it is worth seeing it whole. Reading across tenors at a fixed standardised moneyness gives the term structure; reading across moneyness at a fixed tenor gives the smile. The mandated 7 to 12 day entry window sits at the steep left edge of that object, where the surface changes fastest with respect to both arguments and where any error in the model is therefore most consequential for the prices this study computes.

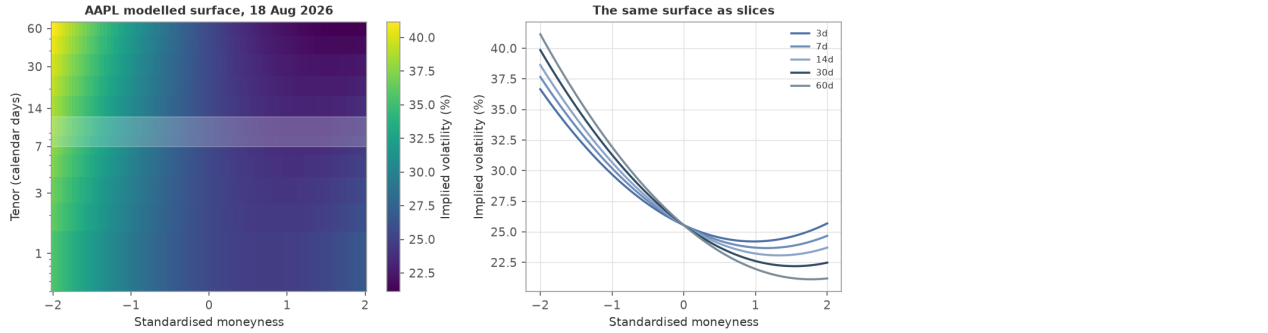


Figure 5.4 - The modelled surface for a single large-cap name on 18 August 2026, as a heat map and as slices by tenor. The shaded band marks the 7 to 12 day window in which every trade in this study is opened.

The same object seen through calendar time makes the earnings signature unmistakable. Rolling the surface forward day by day for a single name shows the short end lifting into each scheduled report and collapsing immediately after, four times a year, while the long end barely moves. Every inversion of the 30-day over 7-day ratio in the lower panel coincides with a reaction day. That regularity is what the entry screen is unknowingly selecting on, and chapter 18 measures what it costs and earns.

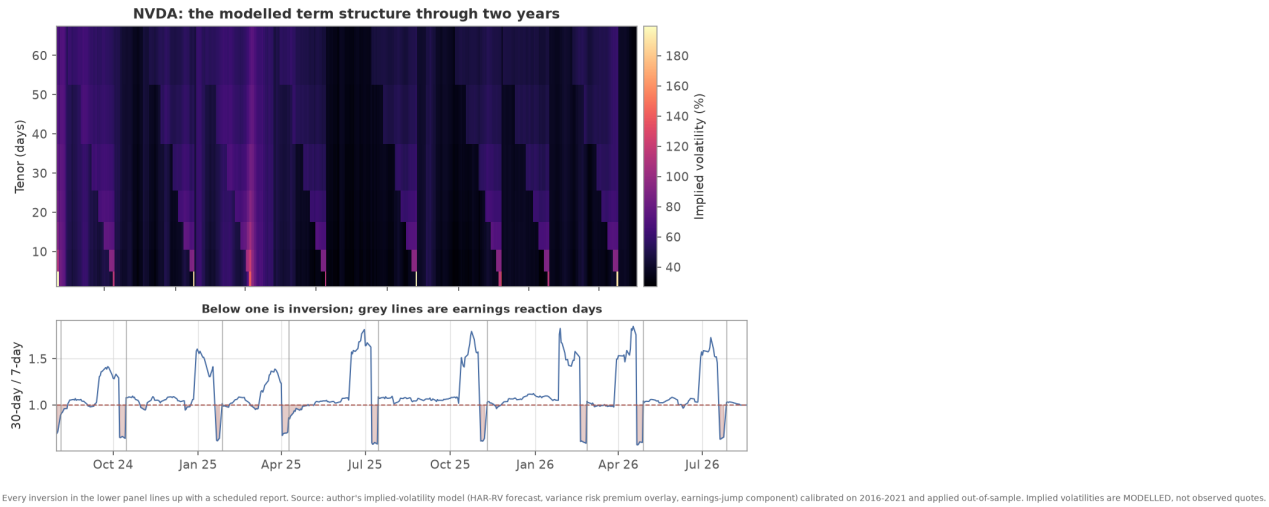


Figure 5.5 - Two years of the modelled term structure for one name. The lower panel plots the 30-day over 7-day ratio; values below one are inversions, and the grey verticals are detected earnings reaction days.

---

6. IV Rank: Construction, Criticism and What It Selects

6.1 Definition and its known defects

6.1.1 IV rank uses only the 52-week minimum and maximum and discards everything between; IV percentile uses the whole distribution

The mandated specification admits a trade only when IV rank exceeds 78, so the arithmetic of that one statistic determines which 822 iron condors and which 807 calendars the backtest ever sees. IV rank is a range position within a trailing 52-week window of implied volatility σ:

$$\text{IV rank}_t = \frac{\sigma_t - \min\limits_{52\text{w}} \sigma}{\max\limits_{52\text{w}} \sigma - \min\limits_{52\text{w}} \sigma} \times 100$$

IV percentile is a frequency over the same window:

$$\text{IV percentile}_t = \frac{\#\{s \in 52\text{w} : \sigma_s < \sigma_t\}}{\#\{s \in 52\text{w}\}} \times 100$$

The first expression reads exactly two observations off the year, the minimum and the maximum, and discards the roughly 250 between them; the second uses all of them.[63][64] Zero is the low print of the year and 100 the high print.[63]

The two therefore disagree, and Barchart's worked example shows how far. For a name that normally trades in the mid-40s to 60 but spiked briefly, leaving the trailing window's low at 40 and its high at 100, a current reading of 60 gives an IV rank of $(60-40)/(100-40) = 33\%$ while 200 of 252 days sat below it, an IV percentile of $200/252 = 79\%$.[64] Same series, same day, 33

against 79: one below every conventional low cutoff and the other above every conventional high one. The gap is an arithmetic identity rather than an empirical finding, and it has a direction — after any spike, IV rank reads lower than IV percentile for as long as the spike stays in the window.

Both statistics here are computed on modelled implied volatilities, not on observed option quotes: a pooled log-HAR forecast fitted on 2016 to 2021 and then frozen, with a variance risk premium overlay calibrated to a median modelled implied-to-expected-realised ratio of 1.060 for single stocks and a separately applied earnings jump component. Every IV rank below inherits that construction.

6.1.2 Both are censored at 0 and 100, and a single spike inflates the denominator for a full year afterwards

Neither metric can leave the range 0 to 100, which destroys information exactly where a short-premium programme most wants it. Once implied volatility exceeds the prior 52-week high, the current print becomes the new upper bound, so a reading of 100 says nothing about how much further volatility can travel.[63] The saturation is symmetric at zero, and it is visible in this universe: the histogram in Figure 6.1 carries a spike of mass against each bound rather than tapering into them.

The denominator is the more insidious defect. One spike raises $\max_{52w} \sigma$ and holds it there for a full year, compressing every subsequent reading toward zero. A book gated on IV rank therefore stops selling premium in the twelve months *after* a volatility event, precisely because the event is still inside the window — the opposite of what the premise intends.

The thresholds are conventions rather than estimates. One vendor calls IV rank elevated above 50 and extreme at 80 and above; another sets high above 70, mid at 30 to 70 and low below 30.[63][64] The first concedes elsewhere that "many volatility traders look at IVRs as low as 20 or 30 as relatively 'high' in this market".[65] A cutoff re-set by eye to suit the prevailing regime is not an out-of-sample rule. The 78 used here sits between the published conventions and was likewise not estimated; chapter 15 measures what moving it costs. That same vendor also describes IV rank in percentile language in one place while giving the min–max definition in another,[65][63] so two published "IV rank" results need not be measuring the same quantity.

6.1.3 The direct published test of high-IV-rank entry is vendor research with no published effect sizes

The only head-to-head test of high-IV-rank entry against unconditional entry located in the research for this report is vendor research that published no effect sizes. Sixteen-delta strangles were sold in SPX and RUT from 2005 and in GLD from 2009, sampled every five days, comparing all occurrences against IV rank above 50 and IV percentile above 50. The reported result is qualitative throughout: "Waiting for high IVR in SPX significantly raised each P/L metric although the win rate dropped slightly", RUT "similar results although to a lesser degree", and for GLD "IVR had a significantly less largest loss".[66] No magnitudes, trade counts, significance tests or cost

assumptions accompany those statements, and the publisher is a broker whose revenue rises with trade frequency. A companion study, SPX from 2013 to 2017 at 45 days to expiry, finds IV rank a better average-P/L filter than a skew rank — five years and one index, again without published effect sizes.[67] The direction of the evidence supports the premise; the magnitude is unmeasured.

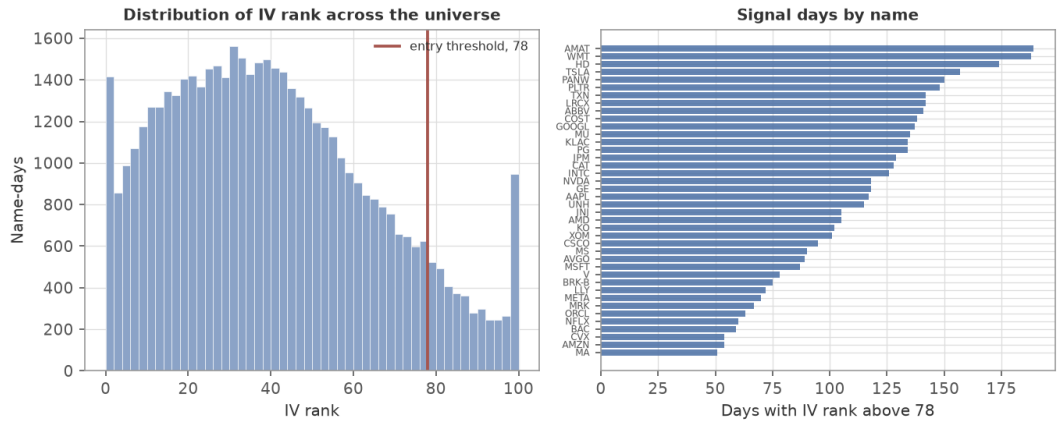
No peer-reviewed study testing IV rank specifically was located. The strongest peer-reviewed support for conditioning premium sales on volatility conditions on the VIX level, not on IV rank: Malkiel, Rinaudo and Saha report that the premium "is higher at elevated levels of implied volatility (as represented by the VIX index level)" and that from 1990 through 2018 the conditional strategy beat both the market and continuous option selling on absolute and risk-adjusted terms.[68] That signal is an absolute index level rather than a 52-week normalised rank, and it is measured on an index rather than on single names; the paper could not be retrieved beyond its abstract, so the cutoff used and the size of the outperformance are unknown here.

A well-known paper prescribes the opposite exposure at high volatility. Moreira and Muir find that "managed portfolios that take less risk when volatility is high produce large alphas", because "there is little relation between lagged volatility and average returns but there is a strong relationship between lagged volatility and current volatility", so the mean-variance trade-off weakens exactly when volatility is elevated.[69] Short premium is a levered short-volatility, long-beta exposure, which makes the implication a reduction in size when the screen fires. The paper studies realised-volatility scaling of equity factors rather than option writing, so it does not refute IV-rank entry directly, but nothing located reconciles the two prescriptions. Pointing the same way, Van Tassel finds one-month and three-month implied-volatility term premia spike *downward* on a VIX rise, becoming positive only two to three months after the shock — the "low premium response puzzle".[70] An IV-rank rule fires during the shock, not two months after it — though that connection is an inference drawn here, not a claim the paper makes. A recent preprint on 0–5 day SPXW options inverts the signal outright, cutting position size when VIX-rank is high, and reports annualised returns from 0.95% to 52.77% across parameterisations of that one signal — a spread that is itself evidence of fragility to specification.[71]

6.2 What the screen actually selects in this universe

6.2.1 IV rank exceeds 78 on 8.85% of name-days, giving 822 condor entries over five years across 40 names

Applied to the 40-name universe over the five-year backtest window, the screen fires on 8.85% of name-days. That yields 822 iron condor entries and 807 entries for each of the calendar and the double calendar. Entries are not evenly spread across names: Applied Materials (AMAT) contributes the most signal days at 189 and Mastercard (MA) the fewest at 51.



Source: author's implied-volatility model (HAR-RV forecast, variance risk premium overlay, earnings-jump component) calibrated on 2016-2021 and applied out-of-sample. Implied volatilities are MODELLED, not observed quotes.

Figure 6.1 — The left panel shows the distribution of modelled IV rank over all name-days with the 78 entry threshold marked; the right panel counts days above 78 for each of the 40 names.

The left panel makes the censoring of Section 6.1.2 concrete: the distribution does not taper into its bounds but stacks mass against both, so a material share of name-days carry a reading that says nothing about how extreme implied volatility actually is. The right panel shows the per-name count is dispersed rather than uniform, making the screen a name-selection rule as well as a timing rule.

6.2.2 Median 30-day implied is 44.8% when the screen fires against 33.5% unconditionally

On the level it is measuring, the screen does what it claims: median 30-day modelled implied volatility is 44.8% on days it fires against 33.5% unconditionally. What comes attached to that level is the more consequential part.

Property of the entry day	IV rank above 78	All name-days
Share of name-days	8.85%	100%
Median 30-day modelled implied volatility	44.8%	33.5%
Median 30-day over 7-day implied	1.023	1.068
Share with an inverted term structure	43.3%	18.5%
Scheduled report inside the next 12 days	32.6%	15.2%

Table 6.1 — Selection profile of the mandated entry screen, measured on this sample.

The screen selects a high level of implied volatility and, simultaneously, a flat or inverted term structure and a proximate earnings date. The median ratio of 30-day to 7-day implied falls from 1.068 unconditionally to 1.023 on signal days, and the share of days on which the front expiry prints above the back rises from 18.5% to 43.3% — chapter 5 sets out that curve state in detail. That matters because the structures are not exposed to the same thing: the condor is short variance outright, the calendars long forward variance and short spot variance, so a flatter curve at entry removes part of what a condor is paid for and supplies part of what a calendar needs. None of that was selected for; all of it is a consequence of the one threshold. The bottom row reaches furthest into the rest of the report.

6.2.3 A scheduled report falls inside the next twelve days on 32.6% of high-IV-rank days against 15.2% of all days, an over-representation of 2.1 times

Measured on this sample, a scheduled report falls inside the following twelve days on 32.6% of high-IV-rank name-days against 15.2% of all name-days, an over-representation of 2.1 times. This is one of the report's own findings rather than a result taken from the literature; the research surveyed for this chapter reached no source that tests what fraction of large-cap high-IV-rank episodes are explained by an upcoming print.

The mechanism is the pricing of the announcement itself. Pre-event annualised implied variance decomposes as a diffusive term plus an earnings jump term divided by the time remaining to the announcement, so implied volatility "increases at rate $(T-t)^{-1}$ as the maturity decreases" and the front expiry can print above the back before collapsing discontinuously at the first post-announcement quote.[60] A run-up of that shape is therefore a reliable route for a large-cap name to approach its 52-week implied high, and any statistic normalised on that high inherits the reporting calendar. On this sample the earnings effect on the underlying is large enough to matter: 1,577 reaction days were detected across the 40 names, with a mean absolute move of 5.48% against 1.38% on ordinary days.

6.3 The consequence for the mandate

6.3.1 Screening on IV rank is therefore partly an earnings screen in disguise, and the earnings exposure drives a large share of the tail

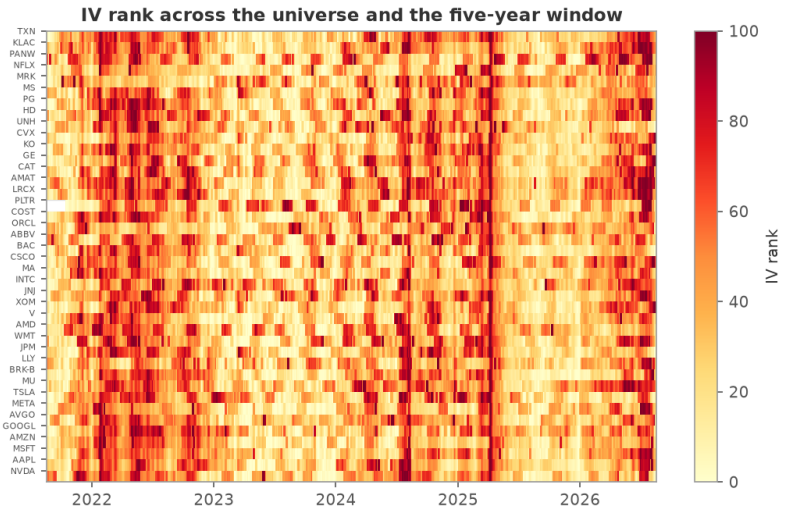
The over-representation carries through from signal days to executed trades. A scheduled report falls inside 28.8% of iron condor trades and 28.3% of trades in each calendar structure — close to three in ten positions, arrived at without any instruction to trade earnings.

Structure	Trades	Share spanning a scheduled report	Worst trade, report inside	Worst trade, no report

Iron condor	822	28.8%	-157%	-124%
Calendar	807	28.3%	-109%	-97%
Double calendar	807	28.3%	-107%	-100%

Table 6.2 — Earnings footprint of the screen at trade level; worst trade as a percentage of capital at risk, measured on this sample.

In all three structures the worst single loss is deeper in the earnings-spanning subsample than outside it, and the gap is widest for the iron condor, whose worst trade with a report inside reaches -157% of capital at risk against -124% without. The condor is the one structure of the three that is short variance outright, so an announcement is two offsetting exposures for it — the post-announcement collapse in implied volatility works in its favour, the gap in the underlying against it — and the net is chapter 18's measurement. The share column matters more than either tail figure: an exposure taken in 28.8% of positions is not a subsample curiosity but a description of what the mandate buys. Whether it is *paid* is a separate question, not settled here. The modelled surface prices the earnings jump at an implied jump volatility of 1.108 times realised announcement-day volatility, a thin margin, and the academic evidence on the run-up window runs against the seller: at-the-money straddles bought three days before an announcement earn a significant +3.34% into the announcement date, with the gain weakest in exactly the liquid large caps screened here.[62][72] Chapter 18 measures the split in mean return on capital at risk with and without a report inside the trade.



Source: author's implied-volatility model (HAR-RV forecast, variance risk premium overlay, earnings-jump component) calibrated on 2016-2021 and applied out-of-sample. Implied volatilities are MODELLED, not observed quotes.

Figure 6.2 — Modelled IV rank for each of the 40 names across the five-year window; dark horizontal dashes recurring within a row are name-specific episodes, and dark vertical bands crossing most rows are market-wide episodes when the screen fires almost everywhere at once.

Figure 6.2 shows the screen has two firing modes, and only one of them is the earnings calendar. The horizontal dashes recur within a row at roughly quarterly spacing; the vertical bands cross nearly the whole universe at once, and between them lie intervals of months in which almost nothing fires. Those bands are why the backtest reaches a maximum of 39 concurrent positions against a nominal cap of 12 on a median of 10 entries per month, a concentration chapter 19 takes up. The two modes need managing differently: the idiosyncratic mode delivers the earnings exposure quantified above, while the market-wide mode delivers a book with almost no diversification at the moment it is largest. What each threshold earns is chapter 15's measurement and what the earnings split earns is chapter 18's — but both are being asked about a signal that, on this sample, selects a proximate report 2.1 times as often as the unconditional rate and clusters its entries in time.

---

7. Iron Condor: Structure, Payoff and Greeks

The iron condor is the only structure in this study entered for a credit with a defined maximum loss on both sides, and its geometry is fully determined once three numbers are fixed: the short-leg delta, the tenor, and the wing width. Under the mandated specification — IV rank above 78, 7 to 12 days to expiry, short-leg delta 25 to 35 on each side — the mandate fixes the first two, and the third is a design choice, made here at a fixed 5% of spot.

7.1 Construction under the mandate

7.1.1 Short strikes at 25 to 35 delta both sides, 7 to 12 days, wings a fixed 5% of spot: on the representative trade, 144.46 and 156.28 short against 136.96 and 163.78 long

The representative trade used for the structure diagrams in this and the two following chapters is a $150 underlying with 10 days to front expiry and 38 to back, priced at 45% at-the-money implied volatility in the front month and 40% in the back, against a 3.87% risk-free rate and a 1.0% dividend yield. Those implied volatilities are modelled outputs of this study's volatility model, not observed market quotes, and should be read as such wherever they appear. The front-month level is close to the 44.8% median 30-day implied volatility recorded when the IV rank screen fires, against 33.5% unconditionally.

Leg	Strike	Type	Quantity	Role
1	136.96	Put	+1 (long)	Lower wing: caps the downside at the wing width

2	144.46	Put	−1 (short)	Downside premium leg, struck at 25 to 35 delta
3	156.28	Call	−1 (short)	Upside premium leg, struck at 25 to 35 delta
4	163.78	Call	+1 (long)	Upper wing: caps the upside at the wing width

Table 7.1 — The four legs of the representative iron condor, from this study's computed trade geometry.

The structure is symmetric in construction but not in strike distance. Both wings are the same fixed 5% of spot, so the maximum loss is identical on either side. The short strikes are not: 156.28 sits further above the underlying than 144.46 sits below it. That asymmetry is produced by the delta targeting rather than imposed on it, because the modelled smile assigns higher implied volatility to the downside, so a 25-to-35-delta put is reached at a smaller percentage move than a call of the same delta. The position is therefore not centred on spot, and a symmetric-looking payoff diagram conceals a directional tilt that the delta figures in Section 7.2 make explicit. The long wings exist only to convert an undefined-risk strangle into a defined-risk condor; at 5% of spot they contribute almost nothing to the credit.

7.1.2 Credit of $271 against $479 of capital at risk on a $750 wing, a credit-to-width ratio of 0.36

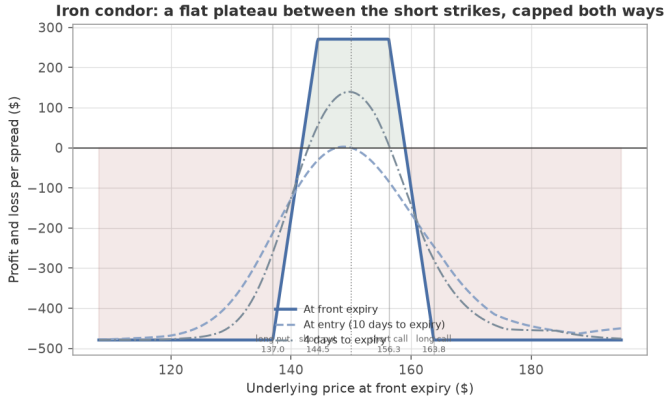
On the representative geometry the four legs net to a credit of $271 against a $750 wing width, leaving $479 of capital at risk — the term used throughout this report for the denominator of per-trade returns, defined as wing width less credit received. The credit-to-width ratio is 0.36. That single number carries most of the trade's economics: it is the fraction of the maximum loss collected up front, and, since the wings define the loss, it fixes the payoff asymmetry. Losing the full width costs $479; winning the full credit pays $271.

The realised backtest geometry is smaller in absolute terms. Across the 822 condors entered under the mandated specification, the median premium was $247 against a median capital at risk of $524. Round-trip transaction costs came to $18.88 per trade, which is 7.0% of the median premium — a figure Section 7.3.2 returns to, because tenor is settled by costs of this kind rather than by the Greeks.

7.2 The payoff and the Greeks

7.2.1 A flat plateau between the short strikes, capped at the credit and floored at width minus credit, with breakevens near 0.83 standard deviations

At expiry the condor pays a flat plateau. Anywhere between 144.46 and 156.28 every option finishes worthless and the position keeps the full $271 credit; outside either long wing the loss is pinned at the $479 of capital at risk. Between a short strike and its wing the value moves linearly at one unit per point. This is genuinely flat-topped and genuinely floored, worth stating because the calendar structures in Chapters 8 and 9 are routinely drawn the same way and are not: their value at front expiry is a curved function with asymptotic flanks.



Representative trade: $150 underlying, 45% front and 40% back at-the-money implied volatility, 10 and 38 days to expiry, 3.87% risk-free rate (3-month Treasury, 2026-08-17), 1.0% dividend yield. Prices are Black-Scholes on the report's modelled smile. Author's calculations.

Figure 7.1 — The condor's expiry payoff: a flat plateau at the $271 credit between 144.46 and 156.28, falling linearly to a floor of –$479 beyond each long wing.

The breakevens in Figure 7.1 sit where the plateau has been eroded by exactly the credit, which on the representative trade is near 0.83 standard deviations from spot in each direction. Under the representative parameters a one-standard-deviation move over the life of the trade is 7.5% of spot, so the profitable band is appreciably narrower than the strike placement alone suggests. This is the condor's central bargain in one line: a high probability of collecting $271, bounded by a low probability of losing $479.

7.2.2 Delta is near zero at inception and flips sign violently as either short strike is approached; gamma is most negative there

Delta is the first derivative of position value with respect to the underlying price, $\partial V/\partial S$ — the dollar change in the position for a one-dollar change in spot, quoted per share and multiplied by the contract multiplier for a whole position. **Gamma** is its derivative, $\partial^2 V/\partial S^2$, the rate at which delta itself changes per one-dollar move. A short option position carries negative gamma: its delta moves against the holder in both directions.

At inception the condor's delta is close to zero, because the short put and short call are struck at similar deltas of opposite sign and largely cancel. That near-neutrality is a property of one instant

and one spot price, not of the trade. As spot approaches either short strike the corresponding leg gains delta rapidly while its wing, still far out of the money, does not offset it, so the net delta swings hard towards the threatened side. The swing is violent rather than gradual because gamma is most negative in the neighbourhood of the short strikes, and gamma at a given strike grows as expiry nears.

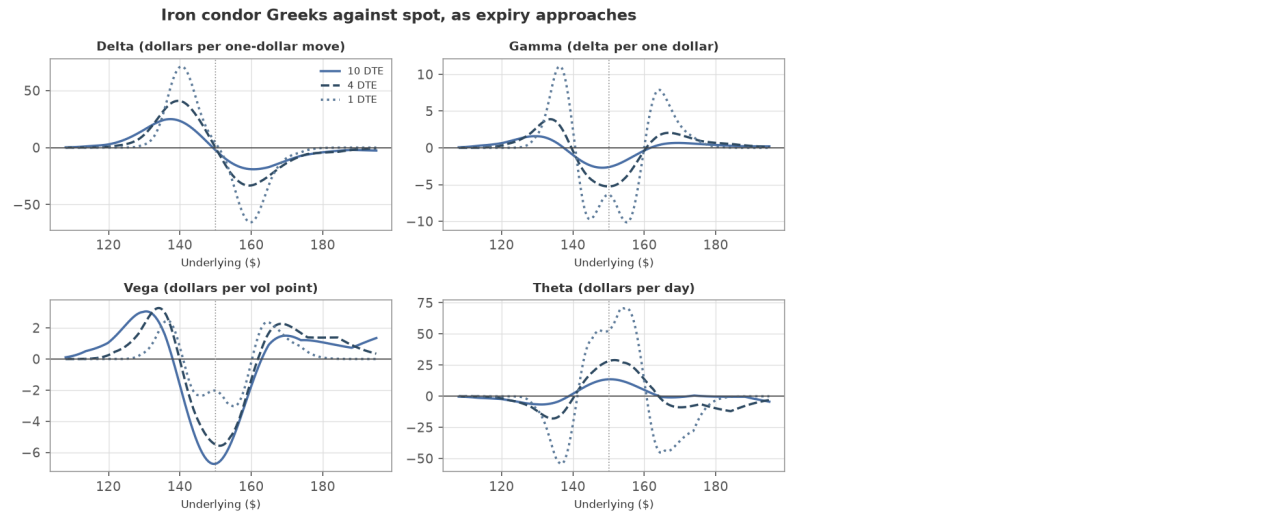


Figure 7.2 — The four Greeks across spot: delta crossing zero near the centre and reversing sharply at 144.46 and 156.28, where gamma reaches its most negative values.

7.2.3 Vega is negative and theta positive throughout, both concentrated between the short strikes

Vega is the sensitivity of position value to implied volatility, $\partial V/\partial \sigma$, conventionally quoted per one percentage point. **Theta** is the sensitivity to the passage of time, $\partial V/\partial t$, quoted per calendar day with spot and implied volatility held fixed. The condor is net short options at the strikes nearest spot, so it is short vega and long theta everywhere between the wings.

Greek	Definition	Sign at inception	Where it concentrates	What the exposure is
Delta	$\partial V/\partial S$, per $1 of spot	Near zero	Rises steeply approaching 144.46 or 156.28	Residual direction, not an intended exposure
Gamma	$\partial^2 V/\partial S^2$	Negative	Most negative at the two short strikes	The realised-move risk being sold

Vega	$\partial V/\partial \sigma$, per point of implied volatility	Negative	Between the short strikes	Short implied volatility level
Theta	$\partial V/\partial t$, per day	Positive	Between the short strikes	The compensation received for gamma

Table 7.2 — Signs and concentration of the four Greeks at inception on the representative trade. Point values are not tabulated in this study's computed facts, so none are quoted here; Figure 7.2 is the source for their shape across spot.

The pattern in Table 7.2 is the whole trade in four rows. Delta is incidental and unwanted, a by-product of delta-targeted strike selection rather than a view. Gamma and theta are the two sides of a single bargain: the position is paid to accept convexity risk, and Section 7.3.1 shows that the exchange rate between them is fixed by arbitrage. Vega is the exposure most often misattributed: a condor entered on a high IV rank reading is frequently described as a bet on volatility falling, but the negative vega decays with the same square root of time that governs everything else at this tenor, so at 7 to 12 days it is small relative to the gamma exposure. The short-dated condor seller is paid mainly for gamma, and only incidentally for a decline in implied volatility. Figure 7.3 shows how that structure hardens as expiry approaches.

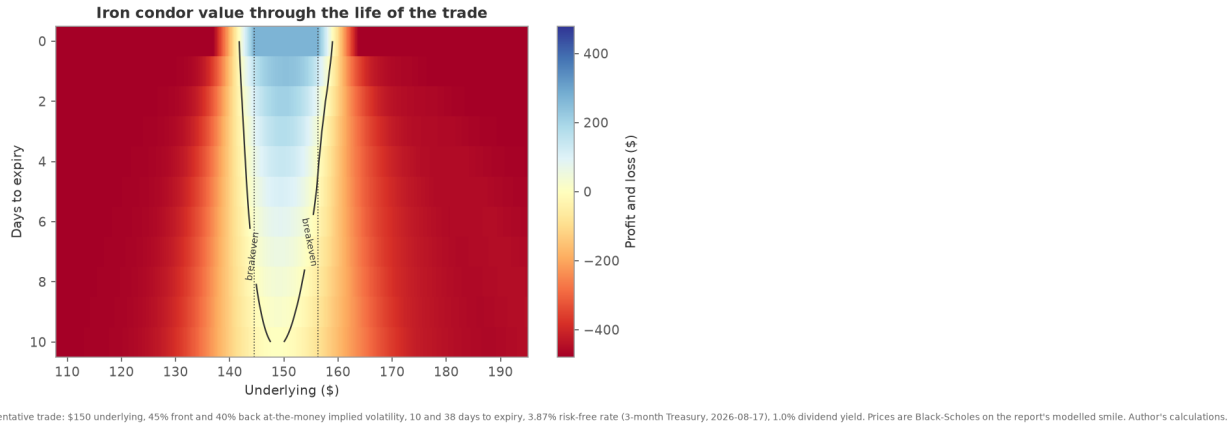


Figure 7.3 — The value surface across spot and time: the plateau sharpens into the expiry payoff as time runs out, and the gradient steepens at the short strikes rather than uniformly.

7.3 The days-to-expiry question

7.3.1 Theta and gamma both scale as one over the square root of time, so their ratio is invariant and there is no decay-to-risk sweet spot to find

The usual justification for a short tenor is that time decay accelerates into expiry. It does. The difficulty is that the risk being taken accelerates by exactly the same factor. In closed-form

Black-Scholes, gamma for a European option is

$$\Gamma = \frac{e^{-qT}\varphi(d_1)}{S\sigma\sqrt{T}}$$

and the leading term of theta is

$$\Theta \simeq -\frac{e^{-qT}S\varphi(d_1)\sigma}{2\sqrt{T}}$$

with both carrying the same $1/\sqrt{T}$ factor.[73] Combining them for a European call gives an exact identity:

$$-\Theta - \tfrac{1}{2}\sigma^2 S^2\Gamma = rKe^{-r(T-t)}N(d_2)$$

so that at zero rates $\Theta = -\tfrac{1}{2}\sigma^2 S^2\Gamma$ at every maturity.[74] The right-hand side is a discounted probability term that depends on the strike and the rate but not on any trade-off between decay and convexity. Theta is not an independent quantity to be maximised: it is gamma, priced at the variance rate.

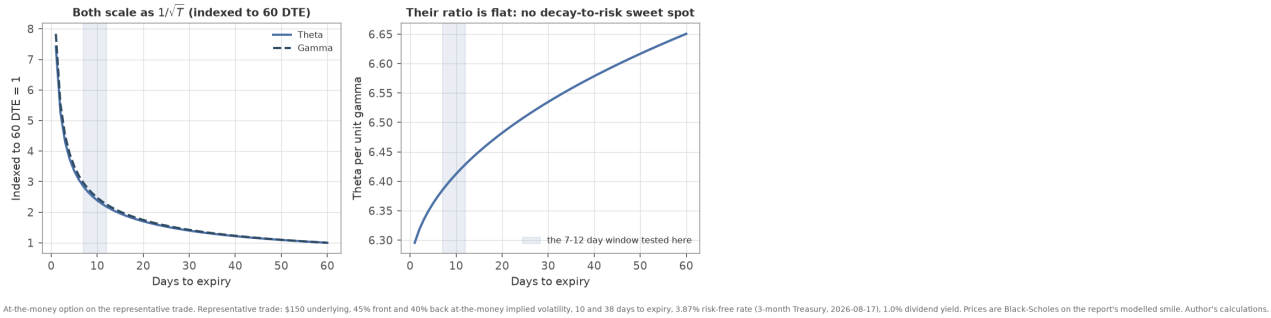


Figure 7.4 — Theta and gamma both scale as one over the square root of time to expiry, leaving their ratio effectively flat across the tenors considered.

Figure 7.4 plots that invariance across tenor, and the consequence for the mandate is direct. Moving from roughly 35 days to expiry to roughly 10 multiplies both theta and gamma by approximately $\sqrt{35/10} \approx 1.9$ times. That is a leverage decision, not an efficiency gain, and any edge available at the short tenor must come from the spread between implied and realised volatility at that tenor rather than from the shape of the Greeks.[73] There is no decay-to-risk sweet spot to locate, because the quantity that a sweet spot would optimise is constant.

The long-history condor benchmarks give some context for how little tenor alone has delivered. Cboe's CNDR index sells a monthly, roughly 20-delta SPX put and call, buys roughly 5-delta wings and holds one-month Treasury bills — a longer tenor, wider strikes and narrower wings than the mandate, on an index rather than single stocks. Its record to 18 August 2026 is +23.48% in total over twenty years, about 1.06% annualised, and +7.83% over ten years, about 0.76%.[64][75] Over the 29.5 years to end-2015, Cboe's own commissioned study ranked CNDR as the least volatile of six SPX option-selling benchmarks at 7.23% annualised standard deviation against 10.16% for the

PUT index, while the leading returns went to BXMD at 10.66% and PUT at 10.13%; CNDR's return does not appear in the sponsor's release.[76] Its worst drawdown was 19% against 47.1% for the iron butterfly benchmark and 51% for the S&P 500.[75] Both Cboe sources are sponsor documentation on Cboe's own indices. The distinguishing property is a truncated tail rather than a return, and whatever the mandate buys at 7 to 12 days is not bought from the tenor.

7.3.2 What shortening expiry actually changes is the premium collected per unit of time and the severity of the gap risk, not the trade-off

What a shorter tenor does change is throughput. Premium scales as the square root of time, so rolling weekly collects more gross premium per year than rolling monthly, but by a factor well below the naive multiple. Measured on Cboe's put-write benchmarks over 2006 to 2018, average annual aggregate premium was 22.1% for the monthly PUT index against 37.1% for the weekly WPUT — a ratio of about 1.68 times, not two and certainly not four, because at-the-money implied volatilities are typically in contango.[77] Roughly 1.7 times the gross premium is what four times the roll frequency buys.

That extra gross premium did not survive the round trip. Over the ten years to 2015 the monthly PUT index compounded at 6.59% a year against 5.61% for the weekly WPUT and 7.09% for the S&P 500, with annualised Sharpe ratios of 0.52, 0.50 and 0.46 respectively.[78] Collecting about 1.7 times the gross premium produced about one percentage point less compound return and an essentially identical risk-adjusted result: the shorter tenor was not rewarded on risk-adjusted terms. Against a round-trip cost of 7.0% of the median premium on the condors here, and an average hold of 4.8 days under the managed rule, the mechanism is easy to see. Frequency multiplies the cost line in full while multiplying the premium line only by a square root.

The gap-risk claim needs a similar qualification. Convexity risk does rise sharply into expiry, and short-put returns are empirically concentrated in the final days of the cycle, with back-month options generating almost no return — the strongest published support for a window of this kind, and explicit that the compensation is for jump exposure rather than a free harvest.[79] At the extreme, 0DTE strategies deliver a real variance risk premium but a return distribution so wide that mean returns are not statistically significant.[80] Yet at programme level the weekly put-writing index had the smaller drawdown, –24.2% against –32.7% for the monthly, because re-striking to at-the-money 52 times a year prevents any single multi-week gap from catching one strike.[81] Shorter tenor is worse per trade and not automatically worse per programme; the two statements are reconciled by rebalancing frequency, not by tenor.

Two limits should be stated. No located study isolates the 7-to-12-day band itself: published comparisons cluster at 0DTE, weekly, 21 days, 30 to 45 days and monthly, so any tenor-specific claim about the mandate is interpolation rather than a published result. And liquidity has migrated to the front of the curve — SPX options expiring the same day averaged 2.3 million contracts daily in 2025, 59% of total SPX volume, within a record 61 million contracts of average daily volume across US listed options[82] — so 7-to-12-day contracts are increasingly not the most liquid line on

the board. That evidence is index-level and should not be extrapolated to the single-stock universe traded here.

The tenor in the mandate therefore rests on the implied-to-realised spread available at 7 to 12 days rather than on any property of the Greeks, and the Greeks are the same trade viewed from two directions. Whether that spread is large enough to pay for four-leg costs at this frequency is a question the sensitivity work and the results chapters must answer, not one the structure can. The calendar spread in Chapter 8 attacks the same variance risk premium from the opposite side of the term structure, buying the tenor the condor sells, and its value at front expiry is not a payoff diagram at all.

A Greek profile drawn against spot at a frozen moment says what the position would be worth at prices it has not reached. What a trader lives with is the profile travelling through time as the underlying does whatever it does, and that can be read off the backtest directly. Aggregating every trade by day since entry shows the condor holding a small negative delta and a stable negative gamma for the first four sessions, then both deteriorating sharply as expiry approaches and the short strikes come within reach. Theta is collected steadily early and turns unreliable late. The comparison with the two calendars is included here because the contrast in vega is the cleanest single statement of what separates the three structures, and chapters 8 and 9 build on it.

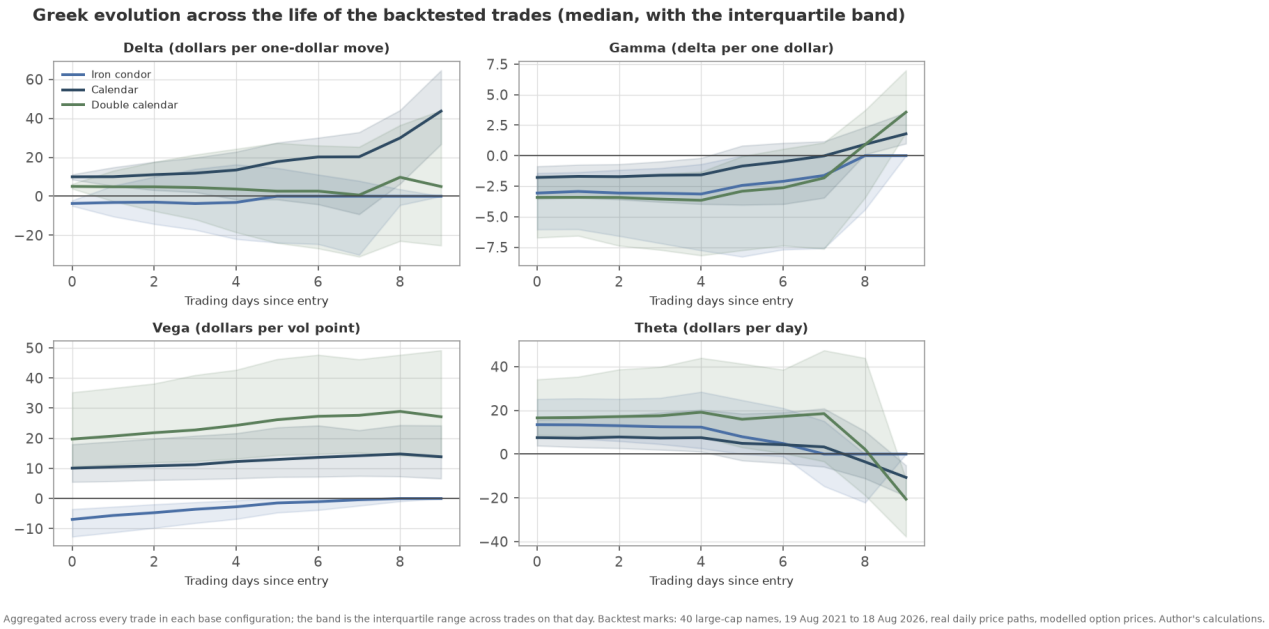


Figure 7.5 - Greek evolution across the life of the backtested trades: median by day since entry, with the interquartile range across trades. The condor's vega sits below zero throughout while both calendars sit above it.

---

8. Calendar Spread: Getting the Payoff Right

8.1 What the position is worth at front expiry

A long calendar spread sells the front-expiry option and buys the same strike in a later expiry. The representative trade carried through Chapters 7 to 9 is a $150 underlying, a 10-day front leg against a 38-day back leg, 45% front and 40% back at-the-money implied volatility from this study's modelled surface rather than from observed quotes, a 3.87% risk-free rate and a 1.0% dividend yield; the call calendar is struck at 154.76 for a $322 debit. What it is worth on the day the front leg expires is the one feature of the trade that published diagrams routinely draw wrong.

8.1.1 The value is an identity, not a payoff: the surviving back-month option priced at that moment, less the intrinsic surrendered on the expiring leg

At front expiry T_1 the sold option settles at intrinsic value and nothing else, so what remains is the back-month option with $T_2 - T_1$ of tenor to run, net of the intrinsic paid away. For calls,

$$V(S) = B(S;\, T_2 - T_1,\, \sigma_{\text{back}}) - \max(S - K,\, 0)$$

where B is the Black-Scholes value of the surviving back leg. This is an identity, not a payoff function, and the distinction is load-bearing. A payoff is a settled cashflow, determined once S is known; the right-hand side contains a model price, because B depends on σ_{back}, the back-month implied volatility prevailing *at* T_1, which is unobservable at entry. CME's worked example is this arithmetic: at the September expiry of a September/December 2440 calendar bought for a net 33.75, the September call expires worthless while the December call "is at the money with three months remaining. It still would be worth about 50.00 – minus the spread cost of 33.75, netting the trader a profit of 16.25."[83] The 50.00 is a valuation, not a settlement. Peak profit is therefore unknown at entry; at the forward it scales with the square root of the remaining back-month tenor and approximately linearly in volatility,[84][85] so an error in the assumed back-month volatility passes straight into the apex.

Three consequences follow, and they account for every geometric feature of the curve. The flanks are curved, because the first term is a Black-Scholes value and not a linear function of the underlying. The apex is a corner rather than a smooth maximum, because the second term is kinked at K while the first is differentiable there. And the two wings settle at different levels, because the subtraction leaves a residual on the in-the-money side that it does not leave on the out-of-the-money side.

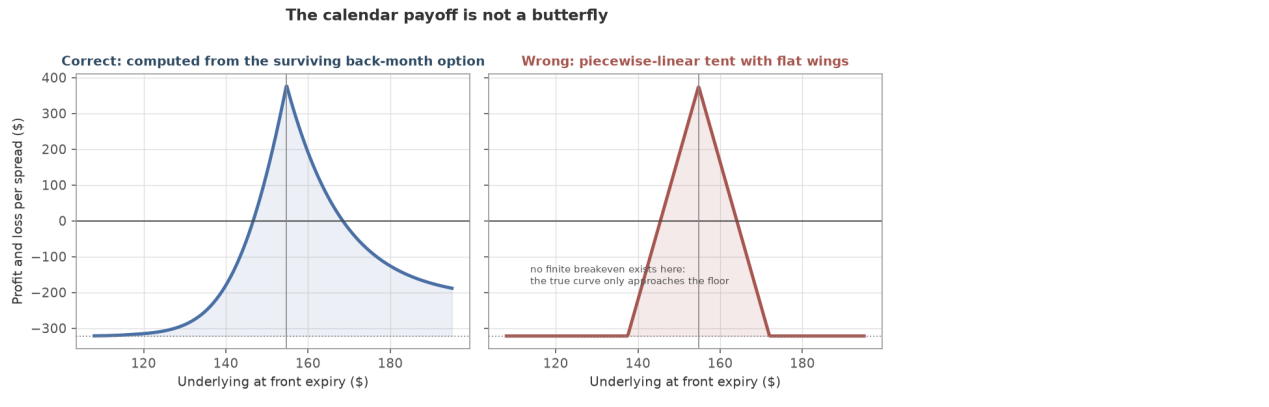


Figure 8.1 — The published tent against the profile the identity generates: the apex agrees, the flanks do not, and the wings do not finish level.

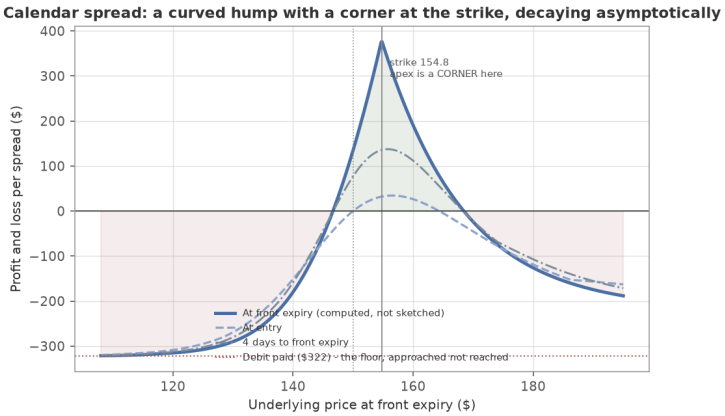
8.1.2 On the representative trade the curve peaks at +$379 at S = 154.75, exactly at the strike; the far downside settles at -$322, equal to the debit floor, and the far upside at -$299

Evaluating the identity on those parameters gives the profile plotted in Figure 8.2, whose three landmarks are measured off the priced curve rather than read from a sketch.

Location on the front-expiry curve	Underlying	Spread profit and loss	Governing term
Far downside limit	$S \to 0$	-$322	The debit floor; both legs decay to zero
Measured apex	154.75	+$379	Back-leg time value at the strike, less the $322 debit
Far upside limit	$S \to \infty$	-$299	The debit floor plus $23 of interest carry retained on the back leg

Table 8.1 — Landmark values of the representative call calendar at front expiry.

The apex sits at 154.75, which is the strike to the resolution of the evaluation grid; the coincidence is exact rather than approximate, for the reason given in 8.1.3. A peak of +$379 against a $322 debit is a gross profit slightly larger than the amount staked, which is the ordinary shape of a calendar. The downside limit of -$322 recovers the textbook claim that the debit caps the loss, and for a single call calendar that is correct as a limit; the upside limit does not, sitting $23 better at -$299. Both are quoted from a diagram in which standardised moneyness is capped at plus or minus 2.5, applied so that the modelled quadratic smile cannot price a deep in-the-money option above its own intrinsic value in the far wings. That cap governs the structure diagrams only, never the backtest, where every strike sits inside one standard deviation of the forward.



Representative trade: $150 underlying, 45% front and 40% back at-the-money implied volatility, 10 and 38 days to expiry, 3.87% risk-free rate (3-month Treasury, 2026-08-17), 1.0% dividend yield. Prices are Black-Scholes on the report's modelled smile. Author's calculations.

Figure 8.2 — The representative calendar at front expiry: a pointed maximum at the strike, curved flanks, and two floors that differ by $23.

8.1.3 The apex is a corner, not a smooth maximum: the slope jumps by exactly one at the strike

Differentiating the identity gives the result directly. Below the strike the subtracted intrinsic is identically zero, so $V'(S) = \Delta_{\text{back}}$, the delta of the surviving back leg, which is positive. Above it the subtracted term differentiates to one, so $V'(S) = \Delta_{\text{back}} - 1$, which is negative. The slope jumps by exactly one as the underlying crosses K, taking no intermediate value. At the strike the surviving back leg is at the money over its remaining 28 days, so its delta is in the region of one half: the curve approaches the apex at roughly plus one half and leaves it at roughly minus one half.[86] Because the derivative changes sign discontinuously and is bounded away from zero on both sides, the maximum cannot be interior to a smooth arc — it is a corner, located at the strike by construction rather than by numerical search. Cboe reaches the same location by another route, noting the long call "has the most time value when the index price equals the strike price."[87] Only the identity establishes that the point is a kink, and that step is a derivation: no exchange or industry publication reviewed here states it.

8.1.4 The flanks are curved and asymptotic, so no finite breakeven exists on the outside - the commonest error in published diagrams

The tent metaphor is half right — pointed apex, wrong sides. Away from the strike $V(S)$ is a Black-Scholes value less either zero or a straight line, and a Black-Scholes value is convex in the underlying, so the flanks curve and then flatten toward their floors without reaching them. Cboe's own limiting condition is incompatible with straight sides: if the index "moves far from the strike price of the calls in either direction (up or down), the amount paid for the strategy … is lost. That's because the price of both calls approach zero (if the index goes down) or parity (if the index goes up), a difference of zero in both cases."[87] The Options Industry Council states the same condition — "the maximum loss would occur should the two options reach parity"[86] — and parity is approached, not attained at finite spot.

Since $V'(S) > 0$ below the strike and $V'(S) < 0$ above it, the curve is monotone on each side and never flat. No finite underlying price exists at which the position stops losing and the maximum loss is realised; every further dollar of move costs something, in ever-diminishing amounts, without end. The zero-crossings are finite and unremarkable, but the outer structure a tent implies — a point beyond which the profile goes horizontal at the floor — does not exist. Flat wings meeting the floor at two finite points assert three false things at once: that the loss is attained rather than approached, that the flanks are linear, and that the two floors are level. Figure 8.1 overlays the two profiles, and they disagree most in exactly the region a risk manager reads.

8.2 Why the two wings settle at different levels

8.2.1 A deep in-the-money call calendar retains K times one minus exp of minus r delta-T in interest carry, which is why the upside wing sits $23 above the downside wing here

Deep out of the money both legs go to zero, the spread is worth nothing and the debit is lost in full. Deep in the money the back leg is priced by put-call parity at $S - Ke^{-r(T_2 - T_1)}$ while the expiring front leg surrenders $S - K$, and the difference does not vanish:

$$\lim_{S \to \infty} V(S) = K(1 - e^{-r(T_2 - T_1)})$$

The surviving long call is a deferred obligation to pay the strike, and the deferral is worth the interest on K across the gap between the two expiries. On the representative trade that gap is 28 days at a 3.87% rate against a strike of 154.76, and the pricing returns $23 per spread net of the 1.0% dividend yield, which pulls the other way. That is the whole of the $23 by which the upside limit of -$299 exceeds the downside limit of -$322: small against the debit staked, but systematic rather than random. Cboe's "a difference of zero in both cases" is the zero-rate simplification of the same limit.[87] The sign is specific to calls. For a put calendar the parity residual reverses and the deep in-the-money wing sits *below* the debit floor, which is why the double calendar of Chapter 9 can lose more than it staked while this single call calendar cannot. The residual is derived from the identity and from parity, and is reported as a derivation.

8.3 The Greeks and what the trade really is

8.3.1 Long theta and long vega from the back leg, short gamma from the front, and slightly long delta at an out-of-the-money call strike

The Greeks defined in Chapter 7 net out across the two legs by tenor rather than by direction. At-the-money theta and gamma both scale as one over the square root of time to expiry while vega scales as its square root, so the short 10-day leg dominates theta and gamma and the long 38-day leg dominates vega: positive theta from the decay differential, positive vega from the back leg, *negative* gamma from the front. The Options Industry Council states each leg's contribution — "the near-term call suffers more from time decay, i.e., has a greater Theta. Less intuitively, the

near-term call has a lower Delta but a higher Gamma (if the strike is at-the-money)", and "longer-term options have a greater sensitivity to changes in market volatility, i.e., a higher Vega."[86] Delta is small and positive at inception, because the 154.76 strike sits above the $150 spot at roughly 35 delta and the back leg's delta there slightly exceeds the front leg's. Figure 8.3 traces all four across the spot range.

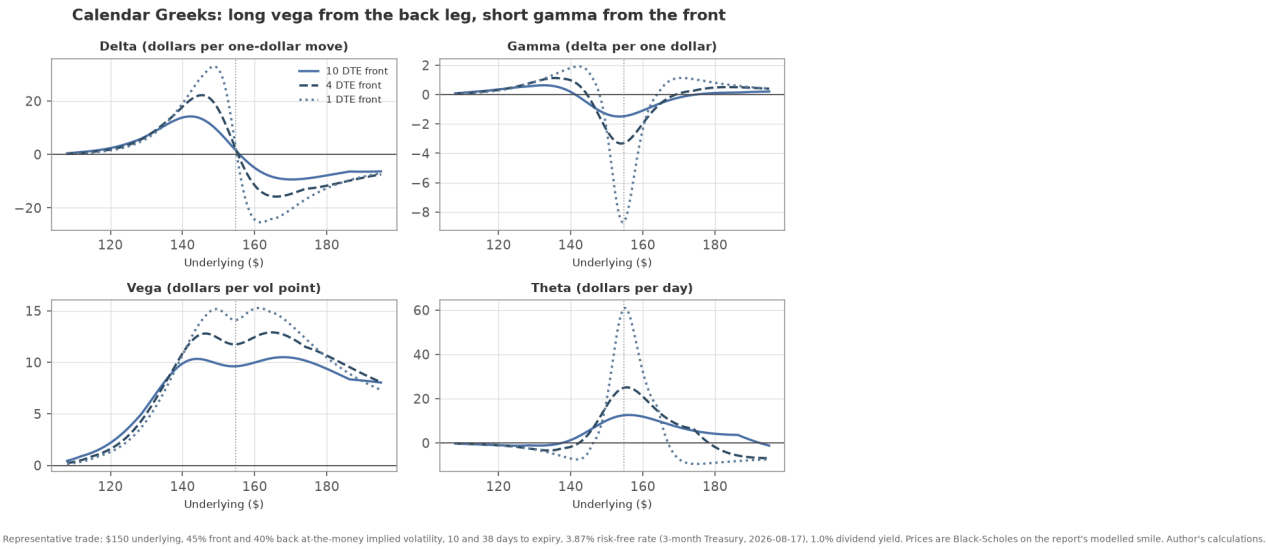


Figure 8.3 — The Greeks across spot: positive theta and vega, negative gamma, and a small positive delta at the $150 starting price.

8.3.2 The vega is term-structure exposure, not level exposure: a front-only volatility rise loses money despite the position being net long vega

Net long vega is a summary statistic that conceals the exposure it summarises. The two legs carry different implied volatilities and move independently — the Options Industry Council warns they "could and probably will trade at different implied volatilities"[86] — so what the calendar is long is the *shape* of the volatility term structure, not its level. A parallel rise in both tenors helps; a front-only rise, steepening backwardation, loses money despite the positive net vega, because the leg that reprices is the one that has been sold.[88] The representative trade is itself parameterised on an inverted curve, 45% front against 40% back, and that is not arbitrary: it is what the entry screen selects. The mandated specification — implied volatility rank above 78, a front leg 7 to 12 days from expiry, a short-leg delta of 30 to 40 — picks days on which the front month is bid relative to the back.

Measure, at trade entry	IV rank above 78	All name-days
Median 30-day implied volatility	44.8%	33.5%
Median 30-day over 7-day implied volatility	1.023	1.068

Share of days with an inverted term structure	43.3%	18.5%
Share with a scheduled report inside the next 12 days	32.6%	15.2%

Table 8.2 — The modelled term structure faced at entry under the mandated specification, against the unconditional distribution across the 40 names.

The screen fires on 8.85% of name-days, and what it selects is a flatter and more frequently inverted curve, not merely a higher one. The median ratio of 30-day to 7-day implied volatility falls from 1.068 unconditionally to 1.023, and the inverted share rises from 18.5% to 43.3%, so a signal day is well over twice as likely to present an inverted curve. The last row supplies much of the cause: a scheduled report falls inside the next 12 days on 32.6% of signal days against 15.2% of all days, and a dated event inside the front window is the standard reason front volatility exceeds back. All four rows come from this study's modelled surface; none is an observed quote. The index-level curve behaves differently — the three-month volatility index VIX3M over the VIX closed at 1.094 on 18 August 2026 and the cash curve was in contango on 92.3% of the last ten years of days, against Cboe's figure of more than 80% for the VIX *futures* curve since 2010[22] — so the definitions are not interchangeable, and the wider the tenor pair, the rarer the inversion.

The outside evidence for shape over level is strong. Johnson finds the shape of the VIX term structure "conveys information about the price of variance risk rather than expected changes in the VIX, a rejection of the expectations hypothesis," with the slope principal component negatively predicting returns on 18 variance assets — six maturities each of synthetic variance swaps, VIX futures and S&P 500 straddles — significantly in 33 of 36 tests, next-day returns across extreme slope quintiles spanning 29 to 181 basis points; the level factor does not predict, contradicting most option-pricing models.[27] Simon and Campasano find the futures basis "does not have significant forecast power for the change in the VIX spot index from 2006 through 2011 but does have forecast power for subsequent VIX futures returns," a -0.79 point estimate on the lagged basis reflecting "a risk premium that can be harvested."[28] Both results concern pure volatility-selling assets at index level, and a calendar is not one — short the front and long the back, it takes the opposite side of the slope. No source located for this study quantifies calendar returns conditional on slope, so the mapping from index-level slope to single-name calendar P&L is offered as a structural argument, not a cited finding.

The exposure has a hard edge. A floor of minus the debit is a property of the value at front expiry, not of the mark at every intermediate moment, and the managed rule closes at two days to expiry rather than holding through. If front implied volatility spikes against a static back leg the spread's mark can fall below zero: the worst single calendar trade in the 807-trade backtest lost 468% of its capital at risk, with the mean of the worst 5% at -208%. Figure 8.4 decomposes the net vega into the two legs that produce it.

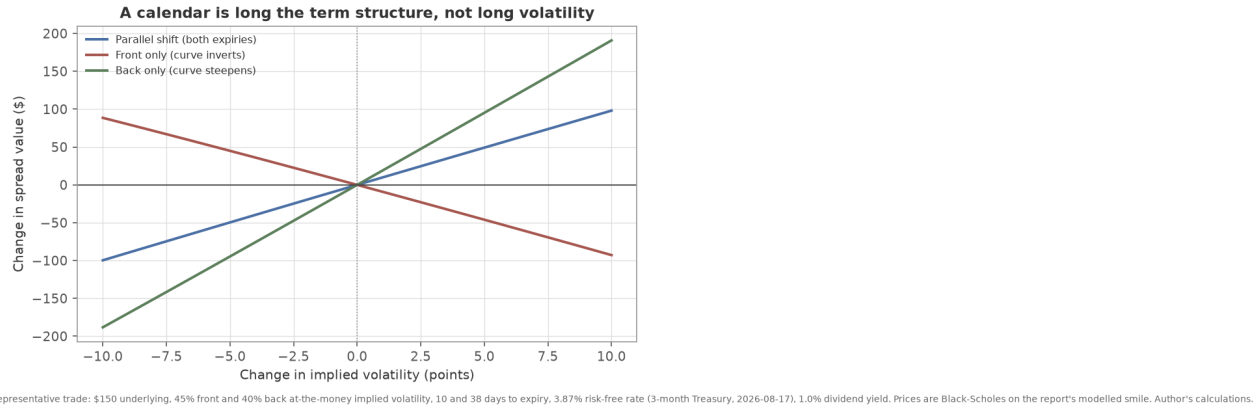


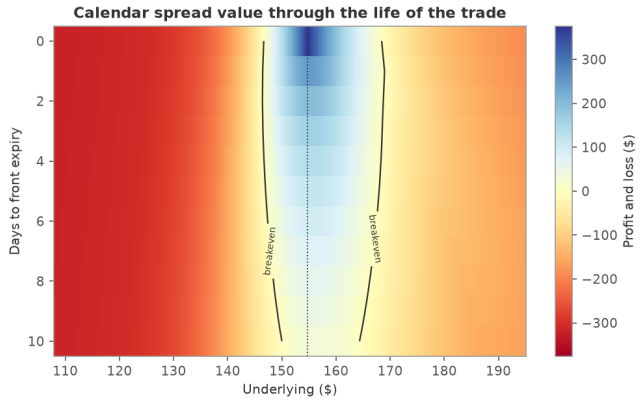
Figure 8.4 — Vega resolved into its two tenors: the position gains on a parallel rise and loses when only the front leg reprices.

8.3.3 The position is long implied and short realised - long forward variance, short spot variance

Because variance is additive across time, selling the 10-day option and buying the 38-day option at the same strike buys the forward variance between the two dates and sells the spot variance to the first, at an implied forward volatility of

$$\sigma_{\text{fwd}} = \sqrt{\frac{\sigma_2^2 T_2 - \sigma_1^2 T_1}{T_2 - T_1}}$$

which prices below back-month at-the-money volatility whenever the curve is inverted and above it in contango.[89] That resolves a disagreement between two exchange descriptions of the same trade. Cboe presents the calendar as a decay harvest — "this strategy uses time decay or Theta, as a mechanism to capture gains" for an investor expecting "minimal price movements … in the near-term but also anticipates longer-term volatility"[87] — while CME presents it as a volatility purchase: "If you believe the market will be stable followed by high volatility … provided that the higher volatility is after the expiry of the front month."[90] Both hold under stated conditions. Decay harvest is correct only with the surface held fixed, in which case the accretion is the mechanical unwinding of a forward-variance position and nothing more. Volatility purchase is correct as soon as the surface moves, which covers any horizon the backtest actually holds — 5.2 days on average. Calling it pure theta collection hides the negative gamma and the vega, and those are where the losses come from. Figure 8.5 traces the position's value across price and calendar time, and the corner of Figure 8.2 is visible in it only as front expiry is approached.



Representative trade: $150 underlying, 45% front and 40% back at-the-money implied volatility, 10 and 38 days to expiry, 3.87% risk-free rate (3-month Treasury, 2026-08-17), 1.0% dividend yield. Prices are Black-Scholes on the report's modelled smile. Author's calculations.

Figure 8.5 — Value across price and time: the ridge along the strike sharpens into the front-expiry corner only in the final days, which is where the managed rule exits.

Chapter 9 doubles this geometry: two strikes, two corners, the same asymmetric wings, except that the put-side residual carries the opposite sign and lets the maximum loss exceed the debit. It also makes the term-structure conditioning reported in Chapter 12 an expected result rather than a curiosity — a position whose economics are a bet on slope should pay differently on an inverted curve than in contango, and on this sample it did, by a wide margin.

---

9. Double Calendar: Two Peaks and a Sag

9.1 Construction and geometry

9.1.1 A put-side calendar below spot and a call-side calendar above it, same strike within each pair: 146.14 and 154.76 for a $623 debit

The double calendar sells a front-expiry put and buys a back-expiry put at a strike below spot, then does the same with calls above it, holding the strike fixed within each pair.[91] Four legs, two strikes, two expiries - a short front-expiry strangle against a long back-expiry strangle struck identically.[91] The back legs outlast the fronts sold, so it opens for a debit.

The representative trade uses a $150 underlying, modelled at-the-money implied volatility of 45% at the front expiry and 40% at the back, 10 and 38 days to expiry - a 28-day gap - a 3.87% risk-free rate and a 1.0% dividend yield. Implied volatility is modelled throughout this study, not quoted. Within the mandated 30 to 40 delta band, 35 delta each side gives strikes of 146.14 and 154.76 for a $623 debit.

Table 9.1 - The eight legs of the two four-leg structures at the representative parameters: this chapter's double calendar and, for contrast, Chapter 7's iron condor.

Structure	Leg	Direction	Type	Strike	Expiry	Days to expiry

Double calendar	1	Short	Put	146.14	Front	10
Double calendar	2	Long	Put	146.14	Back	38
Double calendar	3	Short	Call	154.76	Front	10
Double calendar	4	Long	Call	154.76	Back	38
Iron condor	5	Long	Put	136.96	Front	10
Iron condor	6	Short	Put	144.46	Front	10
Iron condor	7	Short	Call	156.28	Front	10
Iron condor	8	Long	Call	163.78	Front	10

The two structures spend four legs very differently. The double calendar uses two strikes across two expiries; the condor uses four strikes in one. Everything in the condor settles on a single date, which is why it opens for a $271 credit against $479 of capital at risk on a $750 wing width, whereas the double calendar's long legs outlive its shorts by 28 days and nothing settles at front expiry but the two fronts. Its short strikes also sit closer to spot, 146.14 and 154.76 against 144.46 and 156.28, despite the higher mandated delta. That proximity produces the geometry below.

9.1.2 Twin peaks at +$311 and +$434 with a genuine trough of +$305 at S = 147.60 between them - never a single hump, but only $6 deep at 35-delta strikes

Chapter 8 established the front-expiry identity - the surviving back-expiry option priced at that moment, less the intrinsic surrendered on the expiring leg - with its corner at the strike and curved asymptotic flanks. The double calendar applies it twice, once per side. Each back leg holds its greatest extrinsic value when spot sits at its own strike,[87] so the curve takes a local maximum at each and sags between them. No tier-one source specific to the double calendar was located for this study; the twin-peak result is the single-calendar mechanism applied twice, with tier-two confirmation only.[92]

The turning points are measured off the priced curve, not eyeballed. The put-side peak is +$311 at 146.15, the call-side peak +$434 at 154.75, and the trough between them +$305 at 147.60. The

trough is real and this structure is never a single hump, but it is shallow - only $6 below the lower peak - because at 35 delta the strikes are 5.8% of spot apart against a 7.5% one-standard-deviation move. Humps that close overlap heavily.

That is tighter than the literature assumes: it anchors two-sided short strikes at one standard deviation or more[93] and works in a 15 to 20 delta band.[94] Accounts of a "dual tent" or a "valley of death"[92] describe a wider structure than the mandate builds. The peaks are unequal because an at-the-money option's value scales with the level of the underlying as well as with volatility and time.[84]



Figure 9.1 - Peaks of +$311 at 146.15 and +$434 at 154.75 with a measured trough of +$305 at 147.60: at mandate spacing the humps very nearly merge.

9.1.3 Strike width governs the depth of the trough: at 1.6 times the base width the peaks are +$209 and +$386 with the middle at +$180, and at 2.3 times, +$175 and +$390 with the middle at +$72

Trough depth is a property of the spacing, not of the structure. Moving both strikes outward proportionally, everything else fixed, the put-side peak falls from +$311 at the base width to +$209 at 1.6 times it and +$175 at 2.3 times, while the call-side peak holds near its level at +$434, +$386 and +$390. The middle does something else: +$305, then +$180, then +$72. The peaks decline modestly and the trough collapses.[92]

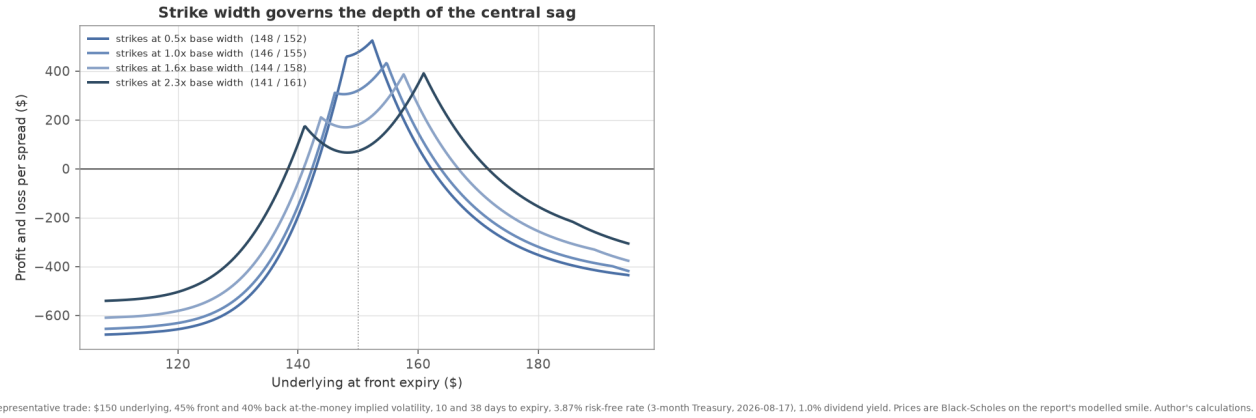


Figure 9.2 - The same structure at four strike widths, from half the base spacing to 2.3 times it: the sag deepens to +$72.

At all three widths the middle stays above zero, so the structure keeps two breakevens rather than four. The four-breakeven case, where the sag cuts the debit line, is an inference corroborated only by a low-tier guide[95] and is reached at no spacing tested here.

9.2 How it differs from its neighbours

9.2.1 Against the single calendar: twice the debit for a second peak and a broader profitable band

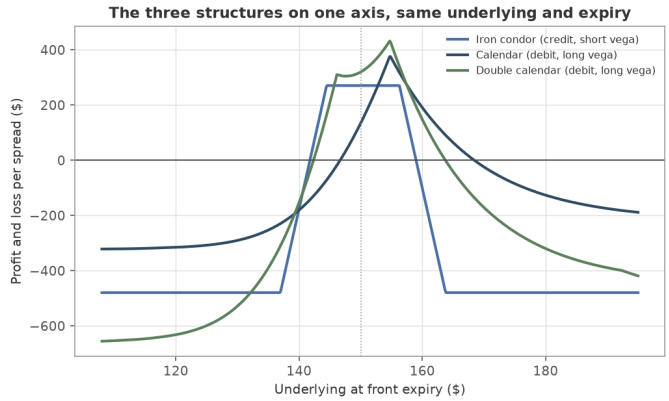
The single calendar at the same 154.76 strike costs $322 against $623, roughly twice the outlay for a second peak and a broader profitable band.[96] Both are positive theta, positive vega and near delta-neutral across that band,[96] so the second pair changes breadth rather than character. At the shared strike the double calendar is worth more, +$434 against +$379, each net of its own debit: at S = 154.75 the surviving back-expiry put at 146.14 still carries extrinsic value the single calendar does not own.

9.2.2 Against the double diagonal: identical strikes within each pair rather than longs placed further out of the money

Two incompatible definitions are in live circulation, and the ambiguity is worth flagging rather than settling silently. One places the long legs further out of the money than the shorts, on a ladder A < B < C < D.[96][97] Fidelity's buys a longer-dated straddle against a shorter-dated strangle sold, longs inside the shorts.[98] The first matters here: "in the case of double calendars, the strikes of the short and long contracts are identical. In a double diagonal, the strikes of the long contracts are placed farther OTM than the short strikes."[96] Cheaper longs mean a smaller debit or a credit and a slower response - the double calendar "responds favorably more rapidly to increases in IV while the double diagonal responds more slowly."[96] The diagonal's vega sign is disputed, so Table 9.2 leaves it unsigned.

9.2.3 Against the iron condor: a debit and long vega against a credit and short vega, and a maximum profit that is unknown at entry

The condor puts four legs in one expiry for a credit, capping profit at that credit and loss at the wing width less the credit[93] - $271 against $479 here. Three differences follow. Shape: a flat plateau between the short strikes, not twin peaks. Vega: condors short, calendars long,[99] so the two want opposite things from implied volatility. Determinacy: it "is impossible to know for sure what the maximum profit potential is, because it depends on the price of the long straddle ... subject to the level of volatility which can change."[98]



Representative trade: $150 underlying, 45% front and 40% back at-the-money implied volatility, 10 and 38 days to expiry, 3.87% risk-free rate (3-month Treasury, 2026-08-17), 1.0% dividend yield. Prices are Black-Scholes on the report's modelled smile. Author's calculations.

Figure 9.3 - The three structures on one axis at front expiry: the condor's plateau, the calendar's one peak, the double calendar's two peaks and sag.

The condor's loss is bounded by construction, the calendars' by carry - as Table 9.2 records and 9.3.1 explains.

Table 9.2 - The four structures on premium sign, Greek signs, shape, profit determinacy and maximum loss, at the representative parameters.

Structure	Premium at entry	Vega	Gamma	Peaks at front expiry	Maximum profit known at entry	Maximum loss
Double calendar	Debit, $623	Positive	Negative	Two, +$311 and +$434	No	-$661 far downside, below the debit
Single calendar	Debit, $322	Positive	Negative	One, +$379	No	-$322 downside, -$299 far upside
Double diagonal	Debit or credit, configuration-dependent	Disputed in the sources	Negative (inferred, not sourced)	Two	No	Not established by the sources
Iron condor	Credit, $271	Negative	Negative	None; a flat plateau	Yes, the credit of $271	$479, the wing width less the credit

Three patterns matter more than the individual cells. Every structure listed is short gamma, so none is a long-volatility position in the realised sense; they differ in where that gamma sits and what is bought with the proceeds. The vega column is the real fork: the condor and the calendars want opposite moves in implied volatility, making them complements rather than substitutes. The determinacy column is the most understated. A condor's maximum profit is a settled cashflow known at entry, while a calendar's terminal value is a model price set by back-expiry implied volatility at front expiry - modelled here rather than observed, so the peaks above inherit the model's error.

9.3 Settlement

9.3.1 Inside the strikes both front legs lapse and a long back-month strangle survives, which can be rolled into an iron condor; and contrary to the usual claim, the maximum loss of -$661 exceeds the $623 debit by about 6% because of interest carry on the in-the-money put calendar

Settling between the strikes at front expiry leaves both short legs worthless and the trader holding a long back-expiry strangle.[98] It can be closed, held, or sold against: write a new front strangle at the same strikes into the back cycle and "you've gotten yourself into a good old iron condor."[97] Outside them the tested short is assigned, leaving a long call, a long put and short stock.[98] Practice closes beforehand, because an in-the-money front leg gains value far faster than the back legs in "gamma week."[97] The managed rule closes at two days to front expiry, which is that convention codified.

The usual statement is that the debit paid is the maximum loss.[91] On this study's own pricing of the representative structure it is not. The far downside reaches -$661 against a $623 debit, exceeding it by about $38, roughly 6% of the debit.

The mechanism is interest carry on the put-side calendar. Far below both strikes the call side is worthless in both legs, but the put side does not go to zero. Both put legs are deep in the money: the expiring front put settles at full intrinsic value, while the surviving back-expiry put is worth roughly the strike discounted over the remaining 28 days, less the discounted spot. What is left is a negative residual,

$$V_{\text{put}} \to -K_p(1 - e^{-r\Delta T})$$

because the holder of the deep in-the-money back put receives the strike 28 days after the front put has paid it away, and that deferral costs money. At 3.87% over the 28-day gap the term is about $43 per spread; the 1.0% dividend yield offsets part of it, which is why the measured excess is about $38 rather than the full $43.

The call side confirms the mechanism by symmetry. A deep in-the-money call calendar retains a positive residual of the same form, so the far upside settles at -$600, $23 better than the debit floor - the same $23 by which the single calendar at the identical 154.76 strike beats its own $322

floor, at -$299. Same strike, same gap, same carry. The wings therefore straddle the debit line rather than resting on it: one above by $23, the other below by about $38. Sources stating that the debit caps the loss are giving the zero-rate simplification. Cboe's far wings, where both legs "approach zero ... or parity ... a difference of zero in both cases",[87] and the Options Industry Council's "maximum loss would occur should the two options reach parity",[86] hold only at a zero rate.

Three qualifications bound the finding. The structure diagrams cap standardised moneyness at plus or minus 2.5, because an unbounded quadratic smile would price a deep in-the-money option above its own intrinsic value in the far wings; that cap applies to the diagrams alone, never to the backtest, where every strike sits inside one standard deviation of the forward. The residual is a rate effect scaling with the strike and the gap between expiries. And it holds under European exercise; American early assignment realises the same economics sooner.

The consequence runs into the results. The backtest records the double calendar's capital at risk as equal to its debit, a median of $737 on both, so returns on risk are struck against a denominator roughly 6% short of the floor reachable at front expiry - one reason a double calendar can lose more than the capital recorded against it. Chapter 13 measures what this geometry produced across 807 trades on 40 names.

---

10. Backtest Methodology and Its Limits

10.1 What is real and what is modelled

10.1.1 Underlying price paths, the VIX complex, rates and SEC filing dates are observed; option prices are not, because no historical option quotes were available to this study

Not one historical option quote enters this study. Every price path, every volatility index and every interest rate used below is an observed market print; every option price, every implied volatility and every IV rank is produced by the model described in this chapter. The inventory below is set out in full so that a reader can locate the exact point at which the study stops measuring and starts assuming.

Series	Source	Date range	Observed or modelled
Daily open, high, low, close, adjusted close and volume, 40 US large caps	Yahoo Finance daily bars	19 Aug 2016 – 18 Aug 2026	Observed

S&P 500 and Nasdaq 100 daily bars	Yahoo Finance daily bars	19 Aug 2016 – 18 Aug 2026	Observed
VIX	Cboe	19 Aug 2016 – Aug 2026	Observed
VIX9D and VIX3M	Cboe	19 Aug 2016 – 17 Jul 2026	Observed, then carried to the sample end by holding each tenor's ratio to the VIX fixed
Cboe SKEW, VVIX	Cboe	19 Aug 2016 – Aug 2026	Observed
Three-month Treasury yield, DGS3MO	FRED	Jan 2015 – 17 Aug 2026	Observed
Baa credit spread, financial conditions, 10y–2y, policy uncertainty, initial claims	FRED	to Aug 2026	Observed; inputs to the fair-value VIX model of chapter 3, not to the surface
10-Q and 10-K filing dates, 37 of the 40 names	SEC company filings	company history to Aug 2026	Observed; used only to validate, never as a model input
Dividend yield, one constant per name	Adjusted-close against unadjusted close over the sample	whole sample	Derived
Yang-Zhang realised volatility, 5/10/21/63/252-day windows	Computed from the daily bars	2016–2026	Derived from observed bars
Diffusive realised volatility, earnings gaps damped	Same estimator on gap-adjusted bars	2016–2026	Modelled

Earnings reaction days, 1,577 across 40 names	Detected from price gaps	2016–2026	Modelled, validated independently
Per-name earnings jump variance	Expanding mean of that name's past squared earnings-day returns	2016–2026	Modelled, past data only
At-the-money term structure, smile, option prices, greeks, IV rank	This study's volatility model	2016–2026	Modelled throughout

Three rows in that inventory deserve to be read twice. The filing dates are held out: they are never an input to detection, only a check on it, which is what makes the validation in 10.2.2 independent rather than circular — and three of the forty names have no filing history on file, so their reaction days rest on the price rule alone. The dividend yield is a single constant per name inferred from the divergence between adjusted and unadjusted closes across the whole decade, so a name that initiated or cut a dividend inside the window carries the wrong carry at both ends. And the two outer Cboe tenors stop on 17 July 2026, a month before the sample ends; over that final month the term-structure shape is held fixed while the level follows the VIX. Everything in the last four rows is model output. Where those quantities are called implied volatilities in this report, they are modelled implied volatilities, not quotes.

10.1.2 Implied volatility is reconstructed as a pooled log-HAR forecast of diffusive variance, plus an earnings jump, times a variance risk premium

The pipeline runs in a fixed order, and each stage is separable from the next.

Realised volatility is first estimated with the Yang-Zhang range estimator on real open, high, low and close bars. Yang-Zhang is drift-independent and handles the opening gap explicitly, which for single stocks is not a refinement but the main event; its efficiency over close-to-close is roughly 7 in practice against a theoretical peak of 14.[100] That gain is smallest exactly where it is most wanted: when variance is dominated by opening jumps the efficiency "will reduce to almost one".[100]

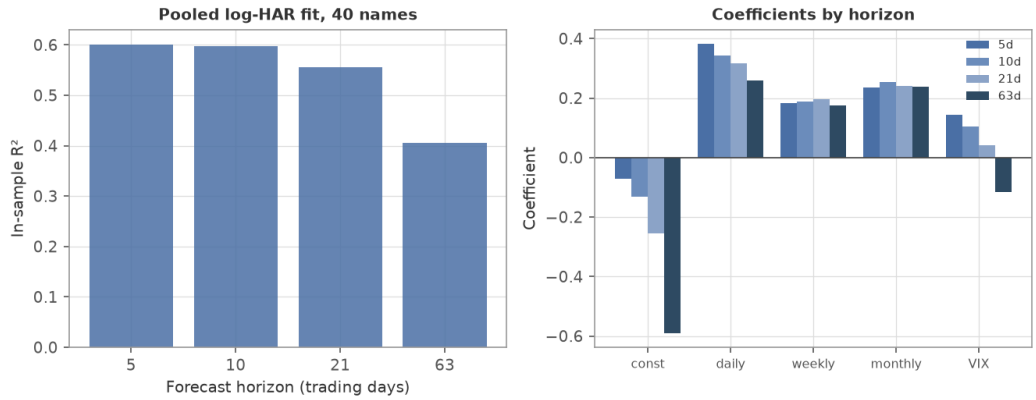
Earnings reaction days are then located by one uniform rule applied to every name — walking forward in quarterly steps and taking, inside a window 60 to 125 days after the last accepted event, the day with the largest earnings-type move, measured as the overnight gap with a partial weight on the intraday move for names that report during the session. No filing calendar is consulted at this stage.

A second realised-volatility series is built for the same names with those days' gaps shrunk to the name's median absolute gap, holding the intraday range. This is the diffusive series: it carries the continuous part of variance with the announcement gap removed, so that the jump can be added back explicitly at pricing time without being counted twice. The damping is imperfect, and in a direction worth naming — only the overnight gap is shrunk, so a report released during the session, or a reaction that continues into the next day, leaves residual announcement variance inside the diffusive leg.

That diffusive series is the target and the input of a pooled log-HAR regression, in the manner of Corsi: log realised volatility over the forecast horizon on the logs of three trailing windows plus the log VIX, one regression per horizon, estimated across all 40 names at once.[101] The three windows stand in for Corsi's daily, weekly and monthly terms; a range estimator needs several bars before it is defined at all, so the short leg is a 5-day window rather than a single day. Searching for better lags does not pay — with flexible lag selection on tick data for the DJIA constituents, "flexible lags do not improve in-sample or out-of-sample fit".[102]

Forecast horizon	Regressors	Fitted R^2
5 trading days	constant, ln RV(5d), ln RV(21d), ln RV(63d), ln VIX	0.600
10 trading days	constant, ln RV(5d), ln RV(21d), ln RV(63d), ln VIX	0.597
21 trading days	constant, ln RV(5d), ln RV(21d), ln RV(63d), ln VIX	0.555
63 trading days	constant, ln RV(5d), ln RV(21d), ln RV(63d), ln VIX	0.405

The fit decays monotonically with horizon, from 0.600 at a week to 0.405 at a quarter, which is the expected shape and is close to the published benchmark: Corsi reports out-of-sample Mincer-Zarnowitz R^2 of 0.551 at one day, 0.609 at one week and 0.570 at two weeks, against an AR(1) that collapses to 0.038 by two weeks.[101] Two qualifications belong beside that comparison. The published triple is the favourable panel of a three-asset table whose asset mapping could not be confirmed, and a second row of the same table shows 0.264 at one day, so it should be read as the good case rather than a guaranteed index number.[101] Separately, regression R^2 is mechanically depressed when a noisy proxy stands in for realised volatility, so a study built on daily open, high, low and close bars should expect measured fit below a high-frequency benchmark for proxy-noise reasons rather than model failure.[103] The horizons the mandate actually trades — 7 to 12 days — sit at the top of this table, where the model is at its strongest.



Source: author's implied-volatility model (HAR-RV forecast, variance risk premium overlay, earnings-jump component) calibrated on 2016-2021 and applied out-of-sample. Implied volatilities are MODELLED, not observed quotes.

Figure 10.1 - The pooled log-HAR across all 40 names: fit falling with horizon on the left, and on the right the weight the regression places on each input, with the VIX term carrying a material share at every horizon.

Total variance over an option's life is then assembled. Writing n for the trading days to expiry, $m(t,T)$ for the number of scheduled reports inside the window, J_t for the per-event jump variance and $\hat{\sigma}_{\text{diff}}$ for the interpolated HAR forecast, the physical expectation is

$$V^P(t,T) = \hat{\sigma}^2_{\text{diff}}(t,n)\,\tfrac{n}{252} \;+\; m(t,T)J_t$$

Each component is then charged its own premium — $\lambda_t(n)$ for diffusive variance, scaled from index level to single-stock level by a factor κ, and λ_J for the jump:

$$V^Q(t,T) = \hat{\sigma}^2_{\text{diff}}(t,n)\,\tfrac{n}{252}\,(1+\kappa\lambda_t(n)) \;+\; m(t,T)J_t\,(1+\lambda_J)$$

$$\sigma_{\text{ATM}}(t,T) = \sqrt{\tfrac{252}{n}\,V^Q(t,T)}$$

$J_t$is the expanding mean of that name's own past squared earnings-day returns, using past data only, so a name's first few events are priced off a universe-wide median and every later one off its own history to that date. The premium $\lambda_t(n)$ is measured, not assumed: at index level it is the ratio of observed Cboe variance to the model's own forecast at three tenors, whose medians over the fitting window are +0.918 at nine days, +0.877 at thirty and +0.740 at three months, log-interpolated across tenor. A multiplicative form is the defensible one — relative variance risk premia, unlike differences, are reported to be independent of the level of variance.[104] The jump premium is set separately and is larger in volatility terms: a variance ratio of 0.23, equivalent to an implied jump volatility 1.108 times realised announcement-day volatility, following Dubinsky and Johannes.[60]

Strike-by-strike, at-the-money volatility is tilted by a quadratic smile in standardised moneyness z, with slope β from the Cboe SKEW index scaled down for single stocks and convexity c:

$$\sigma(K) = \sigma_{\text{ATM}}\big(1-\beta z + c\,z^2\big), \qquad z = \frac{\ln(K/F)}{\sigma_{\text{ATM}}\sqrt{T}}$$

Options are then priced Black-Scholes on that surface, with **the same smile volatility used for the put and the call at a given strike**, so put-call parity holds exactly by construction and no synthetic arbitrage can leak into the results. Costs are charged on entry and again on exit.

10.1.3 The model is fitted on 2016 to 2021 and frozen, so the entire five-year backtest window is out of sample

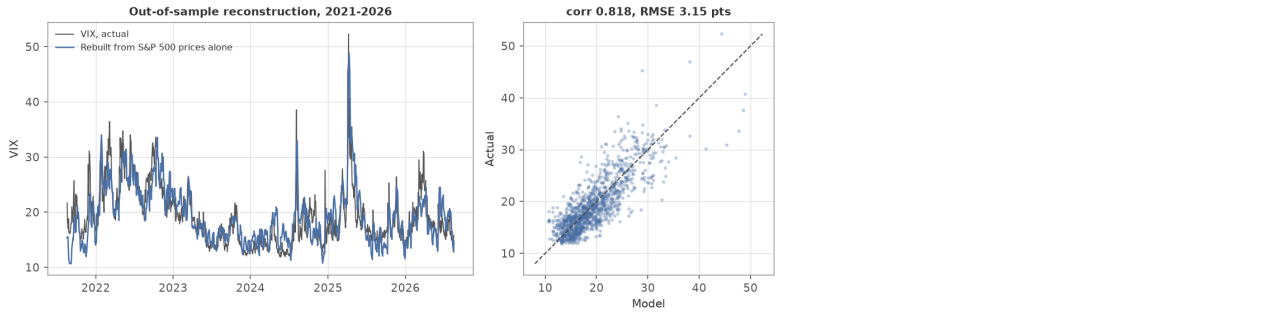
The HAR coefficients are estimated only on observations whose forecast window closes before 18 August 2021, and the index premium is frozen at its median over that same window. The backtest then runs from 19 August 2021 to 18 August 2026 without refitting. IV rank is computed on a trailing 252-day window of the model's own 30-day point using past data only, and each name's jump variance expands forward through time, so no conditioning variable in the study ranks a past observation against data that had not yet occurred — the specific look-ahead channel that hides in full-sample percentiles rather than in timestamps.[105]

Two things are honestly not frozen, and saying otherwise would overstate the case. The single-stock regression takes the log VIX as a regressor, and the single-stock premium is scaled from the *contemporaneous* index premium rather than from its frozen median. Both are observable on the day and introduce no look-ahead, but they mean the surface is not built from stock prices alone: it reads the market's own volatility level each day and redistributes it. The index model, by contrast, deliberately excludes the VIX, which is what makes the test in 10.2.1 a reconstruction rather than a regression of the VIX on itself.

10.2 Validation

10.2.1 Rebuilding the VIX from S&P 500 prices alone reproduces it out of sample with a correlation of 0.818, an RMSE of 3.15 points and a bias of +0.12 points

The same machinery, pointed at the index and denied any volatility-index input, must reproduce a quantity that is observed. Over 2021 to 2026, 1,254 days entirely outside the fitting window, the reconstruction attains a correlation of 0.818, an R^2 of +0.642, a root mean squared error of 3.15 VIX points and a bias of +0.12 points. The outer tenors are worse and in opposite directions: VIX9D at correlation 0.810, RMSE 3.97 and bias +0.60, VIX3M at 0.762, RMSE 3.43 and bias −1.70. The near-zero bias at thirty days is the meaningful result, because a level bias in a premium-selling backtest is a direct subsidy or a direct tax; an RMSE of 3.15 points is not small against a VIX that closed the sample at 15.83.



The volatility model is fitted on 2016-2021 and frozen; this window is entirely out of sample. Source: author's implied-volatility model (HAR-RV forecast, variance risk premium overlay, earnings-jump component) calibrated on 2016-2021 and applied out-of-sample. Implied volatilities are MODELLED, not observed quotes.

Figure 10.2 - The reconstruction tracks the level and the timing of volatility episodes out of sample but compresses their peaks, which is where the 3.15-point error is concentrated.

10.2.2 Earnings reaction days detected from price gaps move 5.48% on average against 1.38% on ordinary days, and sit a median of one day before the corresponding SEC filing

The detector finds 1,577 reaction days across the 40 names. Those days move 5.48% in absolute terms against 1.38% on ordinary days, a ratio of 3.96 — a separation large enough that the rule is clearly finding announcements rather than arbitrary large moves. The independent check is the one that matters: measured against 10-Q and 10-K filing dates reported to the SEC, and never used in detection, the median offset from a detected reaction day to the nearest filing is 1 day. That is the expected sign and size for a market that reacts before the document is lodged. Where the rule will fail is on unscheduled announcements, pre-announcements and reports whose reaction is split across two sessions, and on the three names without filing history, where nothing corroborates it.

10.2.3 The single-stock premium is calibrated to a ratio of 1.060, inside the published range for liquid large caps

The scaling factor K is the study's most consequential free parameter, and it is calibrated to an external target rather than fitted. The resulting median modelled implied-to-expected-realised volatility ratio is 1.060 for single stocks against 1.364 for the index, and the published range for liquid large caps is 1.03 to 1.07. The single-stock figure sits inside that range, in its upper half; the index figure is far above it, which is the correlation premium index options carry and single stocks do not. Chapter 15 shows this assumption to be the dominant sensitivity in the entire study, which is the correct place for it: a ratio-form premium is defensible in functional form, but the scalar itself cannot generate skew and is biased whenever risk-neutral skewness moves, since the bias of implied against physical volatility loads on the third-order risk-neutral cumulant.[106]

10.3 The known biases, stated plainly

10.3.1 A modelled surface cannot reproduce idiosyncratic volatility demand, so real high-IV-rank episodes will include causes this study cannot see

The honest verdict on this class of proxy is that a realised-volatility forecast scaled by a variance risk premium is defensible for a broad index at horizons of a week or more, where the forecast is at its most accurate and the ratio form is stable,[101][104] and **not** defensible for short-dated single-stock options near earnings, where true implied volatility carries an event term the proxy cannot contain and the market charges for a jump the proxy prices at the diffusive level.[60] That criticism is aimed squarely at what this study trades, and the response is structural: the jump is modelled explicitly, from each name's own announcement history, and priced with its own separate and larger premium, which is precisely the omission the criticism identifies. Including earnings jumps has been shown to cut option pricing errors around announcements by more than half.[60]

It does not dispose of the objection. Modelling the jump answers the scheduled, quantifiable part of event risk and says nothing about the rest. A real surface also prices takeover speculation, litigation, product launches, guidance withdrawn between reports and dealer positioning — demand for optionality with no signature in past price ranges. High modelled IV rank in this study can only ever be a high forecast of diffusive volatility, a scheduled report inside the window, or a high VIX. Real high-IV-rank episodes are a strictly larger set, and the extra members are systematically those where the market's price is high for a reason a backward-looking model cannot see. The smile compounds this: it is one quadratic slope scaled off the index SKEW reading rather than a per-name measured skew, so cross-sectional variation in single-stock skew is absent by construction. Any conclusion in chapters 11 to 13 about *which* trades were selected, as opposed to how a given structure behaves, inherits this limit.

10.3.2 Mid-price fills are the largest documented source of overstatement, which is why 65% of a quoted half-spread is charged both ways

Pricing an options backtest at the midpoint is the best-documented way to manufacture an edge. Goyal and Saretto's straddle portfolio returns 22.5% per month at the midpoint and 7.5% when the effective spread equals the quoted spread, and delta-hedged calls fall from 2.1% to statistically insignificant.[107] Worse than the level shift, mid-pricing inverts the cross-section: illiquid names look best at the mid and worst after costs.[107] The correction is not to charge the full quoted spread either, because orders genuinely fill inside the quotes — measured effective half-spreads run at 76% of quoted, an adjusted average at 60%, and execution timers at 42%.[108] This study charges 65% of the quoted half-spread, inside that band but nearer the 60% adjusted average than the 76% effective figure.

Component	Calibration basis	Charge, per contract per leg, one way

Quoted spread		$0.09 average full quoted spread, 374 bps of option price, on the top 15% most active US option classes[109] $0.045 quoted half-spread
Share of the half-spread crossed		65%, inside the 60%–76% measured band[108] ≈ $3.00 on the 100 multiplier

Quoting-increment floor		Cboe Rule 5.4: penny increments only for Penny Interval Program classes, eligibility capped at underliers below $200, nickels otherwise[110]
Commission		Tiered institutional rate, $0.15–$0.25 at systematic volumes[111]
Clearing and regulatory fees		OCC clearing $0.025 per contract, plus Section 31 and ORF[112]
Total, one way		**about $3.25**

$0.04
$0.25

Round trip, iron condor		four legs in $18.88, 7.0% of the $247 median premium and out
Round trip, calendar		two legs in $22.46, 4.4% of the $371 median premium and out
Round trip, double calendar		four legs in $43.95, 4.4% of the $737 median premium and out

The percentage column is where the design of the mandate shows through. The condor pays 7.0% of its premium to costs against 4.4% for either calendar, not because its legs are dearer but because it collects the least — a $247 median credit against $371 and $737 — while crossing four spreads. Only the slippage component is scaled in the cost sensitivity of chapter 15; commissions and fees are fixed, because they are known to within a few cents while the fill assumption is not. Two caveats sit against this table. The spread anchor is a January 2020 market-wide average, with no 2026-vintage measurement available, and spreads widened 256% in dollar terms while posted depth fell 53% in the March 2020 stress — a static charge understates a crisis. More pointedly for this mandate, short-dated contracts are strictly *less* liquid than longer-dated ones at every moneyness, with relative quoted spreads of 34.02% for 20-to-70-day out-of-the-money calls against 22.67% at 71 to 180 days.[113] The trades here are shorter still, at 7 to 12 days, so the charge applied is more likely too low than too high.

10.3.3 Thirty-eight configurations were tried against a minimum-backtest-length guidance of about forty-five for five years of data

Thirty-eight configurations were run, each against fifteen management rules, for 570 configuration-rule cells. The published guidance is that with five years of data no more than about forty-five independent model configurations should be tried before a skill-less strategy is expected to show an in-sample Sharpe of 1 with zero out-of-sample Sharpe; at two years the limit is seven.[114] Thirty-eight sits just inside forty-five, and the comparison flatters the study in both directions at once. The guidance counts *independent* trials, and configurations differing by one delta bucket or one wing width are heavily correlated, so the effective count is below thirty-eight; against that, the rule grid multiplies the true search, and 570 cells is not forty-five by any reading. The configuration search is therefore defensible and the rule search is not, and any single

best-performing rule picked off that grid must be discounted rather than reported at face value. The formal correction — the deflated Sharpe ratio, which needs the trial count, the dispersion of Sharpe ratios across trials, the sample length and the skew and kurtosis of the selected strategy — is applied in chapter 20, alongside what would falsify the study.[115] One further limit is structural rather than statistical: five years contains few genuine volatility crises, and a non-parametric resampling of observed trades "will never" produce an event absent from the sample.[116]

Chapters 11 to 13 report what this pipeline produced. Those figures are not observations of a traded strategy; they are the output of a frozen forecast, an explicitly priced jump, a calibrated premium and a cost charge. The most exposed number among them will be the behaviour of trades spanning a scheduled report — the one place where the model does the most work and the market is most likely to know something it cannot.

---

11. Iron Condor Results

11.1 The base case

11.1.1 822 trades, a 61.4% win rate held to expiry, +2.76% mean return on risk, and a Sharpe of -0.04 before management

The mandated specification — implied volatility rank above 78, 7 to 12 days to expiry, short-leg delta 25–35 — produced 822 iron condor trades over the backtest window. Held to expiry with no intervention of any kind, 61.4% of them made money. Mean return on risk, the per-trade profit or loss expressed as a percentage of capital at risk, was +2.76%. The profit factor, gross winnings divided by gross losses, was 1.08.

Every one of those numbers is positive, and the programme still failed to compound. The same 822 trades, run as a portfolio and held to expiry, returned -0.76% a year with a Sharpe ratio of -0.04, a Sortino ratio of -0.04 and a maximum drawdown of -22.0%. That is the central result of this chapter and the thing the rest of it has to explain: a strategy that won 61.4% of the time and carried a positive average trade still delivered a negative compounded return and a negative risk-adjusted one. The gap is path. A positive arithmetic mean survives compounding only if the losses are small enough and scattered enough not to consume the base they compound from, and at a profit factor of 1.08 there is almost no margin for them not to be. The distribution examined in section 11.2 shows why they were neither.

Two qualifications belong here before the comparison. The option prices underlying every figure in this chapter are modelled rather than observed market quotes; chapter 10 sets out that construction and its known biases, and the results below inherit all of them. And the window is five years, which is short enough that the ordering of a few large losses materially moves the annualised figures.

Against published condor programmes, a 61.4% win rate is unremarkable and the near-zero unmanaged return is close to type. The only long-history listed benchmark, Cboe's CNDR index — a monthly roll that sells an approximately 20-delta S&P 500 put and call, buys approximately 5-delta wings and holds one-month Treasury bills — returned +23.48% in total over twenty years, about 1.06% annualised, and +7.83% over ten years, about 0.76% annualised, to 18 August 2026.[87][64] An independent journal test of S&P 500 zero-days-to-expiry condors over 244 trades recorded win rates of 57.8% at 15-delta and 61.5% at 10-delta.[117] This study's held-to-expiry win rate of 61.4% sits inside that range, and its unmanaged annualised return of -0.76% sits just below CNDR's roughly 0.76% a year over ten years and roughly 1.06% over twenty. The managed +4.12% a year exceeds both, but it is a five-year single-stock result set against a two-decade index programme, and the sponsor's own long-window study begins in 1986 and includes the high-premium 1990s — any headline condor return is dominated by its start date.[76] Selection, in other words, was not the differentiator — the entry screen delivered exactly the kind of high-frequency, low-margin win rate the literature reports, and exactly the same failure to convert it into return.

11.1.2 Under the managed rule the same trades return +4.12% a year with a Sharpe of 0.84 and an 8.2% maximum drawdown

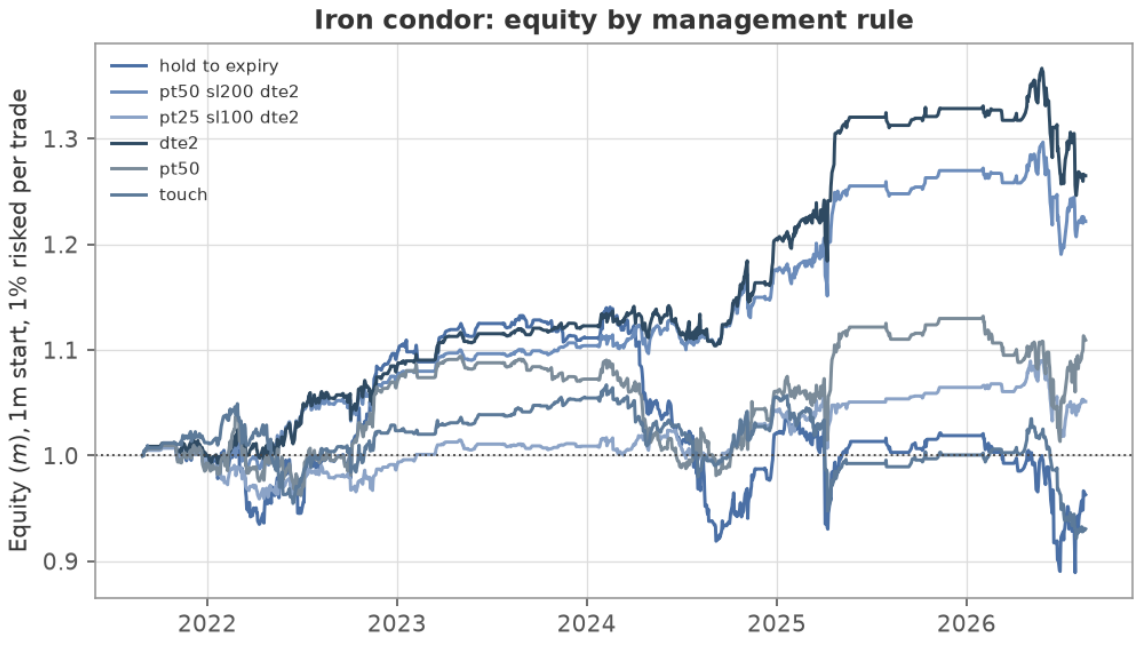
The managed rule — a 50% profit target, a 200% stop and a close at two days to expiry, whichever comes first — was applied to the identical 822 entries. Nothing about selection changed: same names, same days, same strikes, same specification. Only the exit changed.

Metric	Held to expiry	Managed rule
Trades	822	822
Win rate	61.4%	67.2%
Mean return on risk	+2.76%	+4.99%
CAGR	-0.76%	+4.12%
Sharpe	-0.04	+0.84
Sortino	-0.04	+0.96
Maximum drawdown	-22.0%	-8.2%
Profit factor	1.08	1.26

Every headline metric improves, but not in equal proportion, and the pattern of the improvement is the finding. The win rate rises from 61.4% to 67.2% — a modest gain, because the strategy was

already winning often. Mean return on risk rises from +2.76% to +4.99%, and the profit factor from 1.08 to 1.26. The compounded return moves from -0.76% to +4.12% a year and the Sharpe ratio from -0.04 to +0.84, while the maximum drawdown is cut from -22.0% to -8.2%, more than a halving. The Sortino ratio, which penalises only downside deviation, improves from -0.04 to +0.96 — a larger move than the Sharpe, which is what a rule that truncates losses rather than boosting gains should produce. The shift in win rate is far too small to carry the swing in annualised return on its own; the drawdown and Sortino figures identify where the difference came from, which is the left tail rather than the body.

This runs against part of the published evidence. An independent nineteen-year test of S&P 500 strangles found a managed rule set returned a 4.44% compound annual growth rate at a Sharpe of 0.52, against 5.46% and 0.67 for a passive alternative with no profit target.[118] An eleven-year strangle backtest found the 200%-of-credit stop to be the worst-performing approach of those tested.[119] Vendor research points the other way on win rate, reporting 90% success managing one-standard-deviation strangles at 50% of maximum profit against 82% held to expiration.[120] On this sample and this structure, management is not a marginal refinement — it is the difference between a programme that loses money and one that does not. Figure 11.1 shows where along the path that difference accrued.



Backtest: 40 large-cap names, 19 Aug 2021 to 18 Aug 2026, real daily price paths, modelled option prices, costs charged both ways. Author's calculations.

Figure 11.1 — The managed equity curve compounds to a +4.12% annual rate while the held-to-expiry curve ends below where it started, with the two separating at the drawdowns rather than across the whole path.

11.2 The distribution

11.2.1 Negative skew of -0.74 with excess kurtosis of +5.01; the worst single trade lost 469% of the capital posted and the mean of the worst 5% was -191%

The per-trade return distribution has a skew of -0.74 and an excess kurtosis of +5.01. Neither figure is dramatic in isolation, but both point the same way: a tall body of small gains, and a left tail long enough to offset them. The worst single trade lost 469% of the capital posted for it, and the mean of the worst 5% of trades was -191%. Both are multiples of the position's own risk budget, so in this backtest the realised loss tail was not bounded by the nominal defined-risk figure.

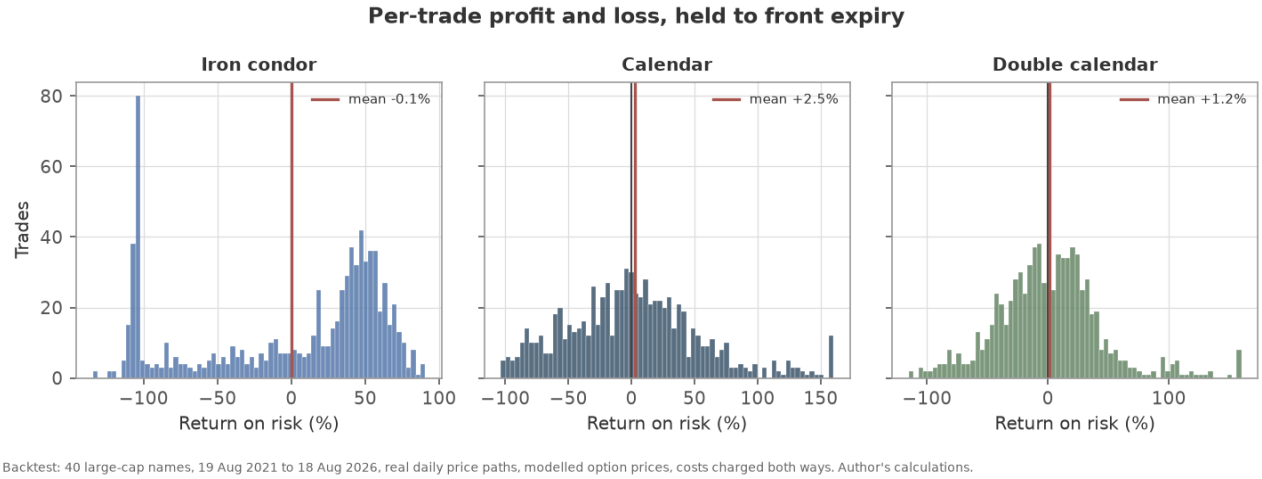


Figure 11.2 — The iron condor panel shows the shape the headline statistics describe: a dense cluster of small positive outcomes with a thin left tail running far past the body. The other panels belong to chapter 14.

That shape, visible in the iron condor panel of Figure 11.2, is the whole reconciliation of the chapter's opening puzzle. Winning 61.4% of the time is compatible with losing money because the average loss is a large multiple of the average win, and the worst losses are a large multiple of the average loss. The published evidence anticipates this precisely. An eleven-year test of S&P 500 strangles found roughly half of trades reached a 50% loss, 25% reached 100% and 16% reached 200%, with about 66% of those hitting 100% going on to hit 200%; its worst observed loss was 434% of premium, in February 2018, and it noted that 18.2 profitable trades would have been needed to recoup that one trade unmanaged.[119] Separate vendor research reports that 79% of all strangle trades went underwater at some point before expiration even though 83% ultimately won.[121] A high win rate and a dangerous loss distribution are the normal joint condition of this trade, not a contradiction.

The wings do work. Cboe's own cross-benchmark study puts CNDR's worst drawdown at 19%, against 47.1% for the iron butterfly and 51% for the S&P 500, and ranks CNDR lowest of six option-selling benchmarks on annualised standard deviation at 7.23%.[87][76] This study's unmanaged maximum drawdown of -22.0% is close to CNDR's 19%, which is consistent with wing width truncating the programme-level tail while leaving individual trades capable of the -469% outcome recorded here.

11.2.2 Round-trip cost of $18.88 against a $247 median credit is 7.0% of the premium, on an average hold of 4.8 days

Trade economics	Value
Median premium collected	$247
Median capital at risk	$524
Round-trip transaction cost	$18.88
Cost as a share of premium	7.0%
Average hold	4.8 days

The condor is a four-legged structure, and every leg crosses the spread twice over an average hold of 4.8 days, so the resulting $18.88 consumes 7.0% of the $247 median credit before the position has done anything. Against a median capital at risk of $524, the median credit is a little under half the posted risk. The cost is not a rounding error at this tenor: it is charged per leg and is largely insensitive to how much premium the leg carries, while premium itself scales roughly with the square root of time to expiry, so a short-dated programme collects less per trade against the same fixed friction.[77] The 4.8-day average hold, against a 7 to 12 day entry window, indicates most positions were closed before expiry rather than carried to it.

The tenor choice is a leverage decision rather than an efficiency gain. Both theta and gamma scale as the inverse square root of time, so shortening from roughly 35 to roughly 10 days multiplies both by about 1.9 times without changing their ratio.[73] Published work finds short-put returns concentrated in the final days before expiry and attributes that to convexity risk rising sharply into maturity — the short-dated seller is paid for jump exposure, not for a free harvest.[79] No located study isolates the 7 to 12 day band specifically, so that literature is directional support rather than a like-for-like benchmark.

11.3 What drove it

11.3.1 Trades spanning a scheduled report returned -2.47% against +2.97% for those that did not

Metric	Scheduled report inside the trade (28.8% of trades)	No scheduled report
Mean return on risk	-2.47%	+2.97%
Worst trade	-157%	-124%

28.8% of trades spanned a scheduled earnings report, an over-representation that follows mechanically from screening on high implied volatility rank. Those trades averaged -2.47% of capital at risk; the rest averaged +2.97%. The sign flips entirely across the split, and the worst trade is worse on the earnings side, at -157% against -124%. For a short-premium structure this is the expected direction: the position is short the gap, and an event that reliably produces one is priced but not, on this sample, priced richly enough to compensate a delta-neutral condor.

Two cautions apply to the table. The study does not state the basis on which the earnings split was computed, so it should not be read as decomposing the +2.76% headline mean or as commensurate with the -469% worst trade from the full unmanaged distribution. And the base case applies no earnings filter at all; what a screen excluding these trades would have returned is a change to the entry rule, which chapters 15 and 16 examine, not a result of the specification reported here.

The earnings composition therefore identifies a real drag without accounting for the unmanaged Sharpe of -0.04: a penalty on 28.8% of trades is small beside a tail whose worst member lost 469% of its own capital at risk. Neither the entry screen nor the earnings mix moved this programme from -0.76% to +4.12% a year — the exit rule did, applied to an unchanged set of entries. That places the burden of the result on a single, crude, three-clause rule rather than on the selection logic the mandated specification is built around, and whether such a dependence is a property of the condor or of short-dated premium selling generally cannot be settled from one structure. The debit structures examined next are the first test of it.

: Stated as an evidence gap in the assigned research file: published DTE comparisons cluster at 0DTE, weekly, 21 DTE and 30–45 DTE, and no study isolating the 7 to 12 day band was located.

---

12. Calendar Results

12.1 The base case

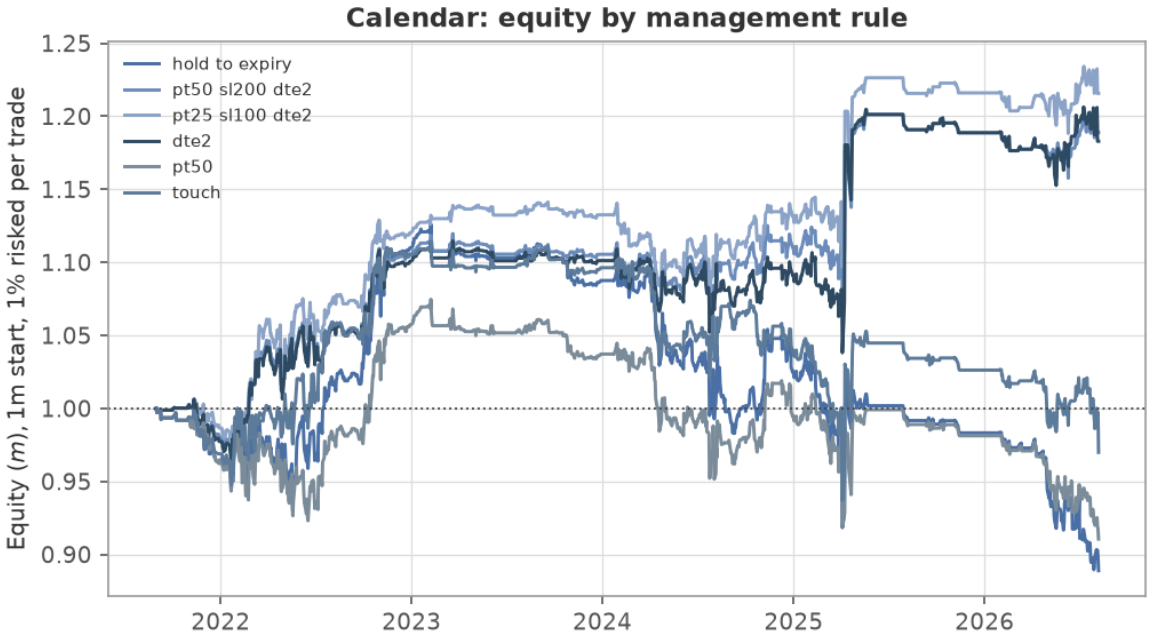
Calendar entries come from the same screen as the condor's — implied volatility rank above 78, a front leg 7 to 12 days from expiry — with the short-leg delta band at 30 to 40 rather than 25 to 35. It produced 807 calendar entries across the 40-name universe over five years, against 822 condor entries. Every implied volatility behind them is modelled rather than quoted, so the term-structure states below are properties of this study's surface, not observed markets.

12.1.1 807 trades, a 49.4% win rate and -4.64% mean return on risk held to expiry; +3.57% a year with a Sharpe of 0.61 after management

Metric	Held to front expiry	Managed rule
Trades	807	807

Win rate	49.4%	59.7%
Mean return on capital at risk	-4.64%	+4.34%
CAGR	-2.36%	+3.57%
Sharpe ratio	-0.25	+0.61
Sortino ratio	-0.33	+0.77
Maximum drawdown	-21.1%	-5.6%
Profit factor	0.85	1.21

Held to front expiry the calendar is not a strategy: it wins slightly less than half the time, surrenders 4.64% of capital at risk on the average trade, and compounds at -2.36% a year through a 21.1% drawdown. The managed rule — a 50% profit target, a 200% stop and a close at two days to expiry, whichever comes first — moves every column, flipping the mean trade positive and cutting the drawdown to 5.6%. The fifteen-rule grid of Chapter 16 identifies which component did it: a 200% stop alone returns -2.36%, identical to holding to expiry; a 50% profit target alone -1.88%; closing at two days to expiry alone +3.46%, against +3.57% for the combined rule. The time stop is doing essentially all of the work, as Figure 12.1 shows, and the stop's inertness is structural rather than a quirk of the sample, for the reason 12.1.2 gives.



Backtest: 40 large-cap names, 19 Aug 2021 to 18 Aug 2026, real daily price paths, modelled option prices, costs charged both ways. Author's calculations.

Figure 12.1 — The same 807 trades under two exit disciplines: held to front expiry the equity line drifts down through a 21.1% drawdown, while the managed rule compounds at +3.57% a year with

a 5.6% worst drawdown.

12.1.2 A debit structure caps the loss at the premium paid, which produces a very different distribution from the condor

The calendar is bought, not sold. Median premium is $371 and median capital at risk is the same $371, because for a debit structure the two are one quantity. Chapter 8 established what that buys: at front expiry the value is floored at the debit, approached asymptotically on both flanks and never attained at finite spot, the two wings settling at different levels. A stop at 200% of premium therefore asks for a dollar loss twice the stake on a structure whose dollar loss is bounded near the stake itself, which is why the 200% stop leaves the held-to-expiry return untouched at -2.36%.

The distribution inherits the same shape. Trade returns are positively skewed at +0.35 against the condor's -0.74: many small losses against fewer larger gains. Excess kurtosis is nonetheless higher, at +9.34 against +5.01, so the tails stay heavy with the skew reversed. Round-trip cost is $22.46, or 4.4% of the premium, against 7.0% for the condor — two legs rather than four — on an average hold of 5.2 days.

Two figures sit awkwardly with the floor: the worst single trade lost 468% of the capital posted and the mean of the worst 5% was -208%. The reconciliation is arithmetic rather than geometric — capital at risk for a calendar is the debit, the debit shrinks as the curve inverts, and a modest dollar loss against a near-zero stake registers as a large multiple of it. The backtest does not decompose that trade, so this is likely arithmetic rather than measured attribution.

12.2 The term structure is the whole trade

Chapter 8 argued the calendar is a term-structure position: long forward variance, short spot variance, with a net vega describing the shape of the curve rather than its height. That was a structural argument; this section is the measurement, and it settles what the argument could not.

12.2.1 Mean return on risk is +10.7% when the curve is inverted at entry against -7.3% when it is in steep contango

Term structure at entry, modelled 30-day over 7-day implied volatility	Calendar	Double calendar
Inverted — back-month implied below front	+10.7%	+6.8%
Steep contango — ratio above 1.10	-7.3%	-4.5%
All trades, unconditional	-4.64%	-3.49%

All six are mean return on capital at risk held to front expiry, unmanaged. That basis matters: the slope condition alone changes the sign of the strategy, with no exit rule helping. Unconditionally the unmanaged calendar loses 4.64% per trade; on the inverted subsample it gains 10.7%; in steep contango it loses 7.3%, worse than its unconditional average. Figure 12.2 plots the trade-level relationship the two conditional means condense. The double calendar of Chapter 13 shows the same ordering at a smaller amplitude, +6.8% against -4.5%. That two structures order the same way on the same 807 entry dates argues the effect belongs to the entry condition, not to one payoff geometry.

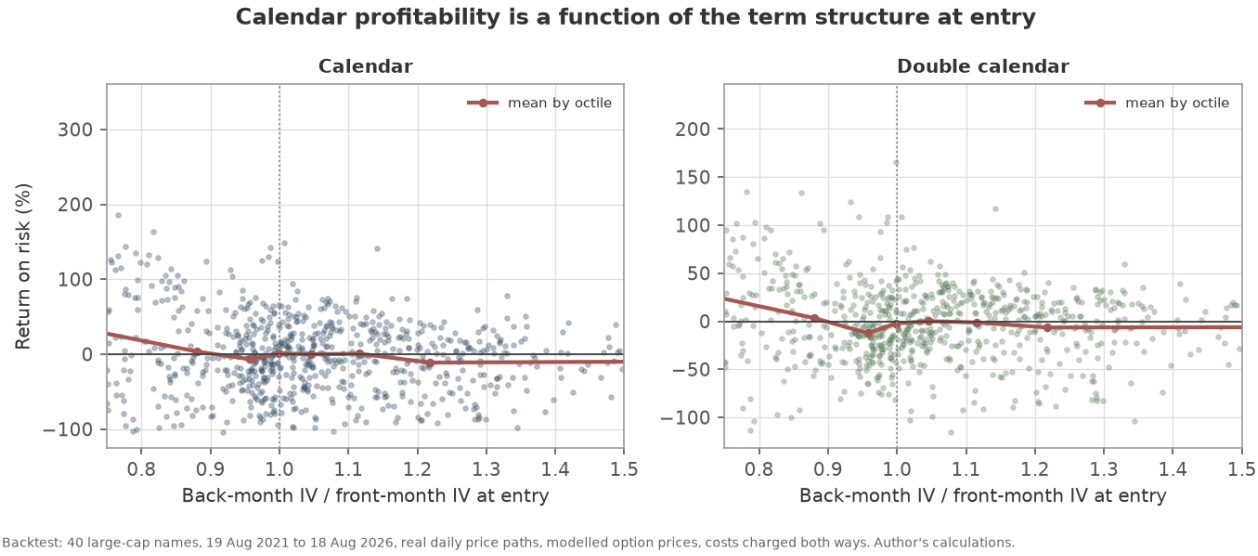


Figure 12.2 — Trade outcomes against the term-structure ratio at entry: profitable trades concentrate on the inverted left-hand side and losses on the steep-contango right, which the two conditional means summarise.

12.2.2 That is the opposite of the intuition that calendars want contango, and it follows directly from selling a rich front against a cheap back

The sign was genuinely uncertain before it was measured, because two forces run against each other.

The pricing force is straightforward. A long calendar sells the front expiry and buys the back, so an inverted curve means selling the expensive leg and buying the cheap one, while contango means paying up for the back leg and collecting little for the front. Through the forward-variance identity set out in Chapter 8, the implied forward volatility the position acquires prices below back-month at-the-money volatility whenever the curve is inverted and above it in contango,[89] so entry price alone favours inversion unambiguously.

The countervailing force is the one the calendar-mechanics literature identifies and does not resolve. Whatever inverts a term structure is usually anticipated realised movement, and realised movement is precisely what the short front leg's negative gamma is exposed to. Buying the entry price cheaply and buying it dangerously are the same transaction. Published sources are

unreconciled: the same term-structure treatment that lists backwardation among the headwinds a calendar faces elsewhere calls the short-front calendar the purest expression of selling event premium.[89] The Options Industry Council warns only that the legs "could and probably will trade at different implied volatilities",[86] and the exchange descriptions specify a volatility path, not a curve shape.[87][90] The peer-reviewed results concern volatility-selling assets at index level — Johnson's slope factor negatively predicts returns on eighteen variance assets, significantly in 33 of 36 tests, while the level factor does not predict at all[27] — and a calendar, short the front and long the back, sits on the other side of that slope. No source located for this study quantifies calendar returns conditional on slope.

What this study measures is that the pricing force wins. On this universe, on this modelled surface, over 807 trades, the entry-price advantage of an inverted curve more than pays for the gamma risk it carries, and by enough to reverse the sign of an otherwise losing strategy. The qualification worth keeping is the mechanics literature's: inversion driven by a dated event resolves on the date, while inversion driven by stress persists as realised movement, and the two cases need not share a sign. This sample's inversions are predominantly the first kind, which is the subject of 12.2.3, so the result should be read as the event-driven case rather than as a universal law about curve shape.

12.2.3 Trades spanning a scheduled report returned +18.59% against 0.00% for those that did not, for the same reason

A scheduled earnings report inside the front expiry is the standard reason a single-stock term structure inverts: the earnings variance term loads on the front expiry and not the back, as Chapter 5 sets out, so the front prices above the back and collapses once the event passes. The earnings split and the slope split are two views of one condition. On 28.3% of calendar trades a scheduled report fell inside the position, and those returned a mean of +18.59% of capital at risk against -0.00% — indistinguishable from zero — for those without one.

The same split run on the iron condor in Chapter 11 has the opposite sign: -2.47% with a report inside the trade against +2.97% without. Both were entered on the same screen and largely the same dates, and both are short front-expiry volatility, so both collect the post-announcement collapse. What separates them is the price move. Earnings reaction days in this universe carry a mean absolute move of 5.48% against 1.38% on an ordinary day, and the condor meets it with short strikes defended only by a wing further out. The calendar meets it with a long back-month option at the same strike as the leg it sold, which bounds its loss near the debit however far the underlying travels. The calendar keeps the collapse and survives the gap; the condor keeps it and pays. Chapter 18 takes the exposure apart properly; what matters here is that the calendar made its money on exactly the trades that cost the condor money.

12.3 Implication for the mandate

12.3.1 An IV rank screen at 78 is past the optimum for calendars: 65 returned +5.37% and 50 returned +4.55%

Minimum IV rank at entry	Calendar CAGR, managed rule
No screen	+1.75%
50	+4.55%
65	+5.37%
78 (mandated)	+3.57%
90	-0.39%

Screening helps; screening this hard does not. The function is single-peaked — unscreened entry returns +1.75% a year, a threshold of 65 returns +5.37%, and the mandated 78 gives back a substantial part of that at +3.57% — and at 90 the strategy stops working altogether, at -0.39%. The mechanism is in what the screen selects: raising the threshold raises the share of entries taken on an inverted curve, the median ratio of 30-day to 7-day implied volatility falling from 1.068 unconditionally to 1.023 when the screen fires and the inverted share rising from 18.5% of all days to 43.3% of signal days. By 12.2.1 those are the profitable entries. Past some point it selects on something else — the most extreme volatility episodes of a name's year, where inversion is likelier to be stress-driven than event-driven and the front leg's negative gamma is paid for rather than harvested. The 90 threshold also thins the sample, so its -0.39% carries the widest error bar. Chapter 15 runs the full grid, where the condor peaks elsewhere.

The practical reading is that the mandated specification screens on the wrong variable and reaches the right one by accident. Implied volatility rank measures the level of the surface; the calendar is paid for its slope, and the correlation between the two is a by-product of what raises single-stock implied volatility in the first place. A screen written directly on the term-structure ratio, or on the presence of a dated event inside the front expiry, would target the exposure the structure actually carries. Whether that is a book anyone should want to run — one deliberately concentrating single-stock earnings risk into a holding period averaging 5.2 days — is the question Chapter 18 has to answer.

---

13. Double Calendar Results

13.1 The base case

13.1.1 807 trades, +2.48% a year with a Sharpe of 0.52 and a 7.4% maximum drawdown under management

The double calendar ran on the same entry screen as the single calendar — the mandated specification of an implied volatility rank (IV rank) above 78, 7 to 12 days to front expiry and a short-leg delta of 30 to 40 — and so fired on the same 807 name-days, holding an average of 5.3 days. Held to front expiry it lost money. Under the managed rule, a 50% profit target, a 200% stop and a close at two days to expiry, whichever comes first, the same entries returned +2.48% a year.

Metric	Held to front expiry	Managed rule
Trades	807	807
Win rate	49.9%	57.4%
Mean return on risk	-3.49%	+1.13%
CAGR	-2.09%	+2.48%
Sharpe	-0.30	+0.52
Sortino	-0.42	+0.73
Maximum drawdown	-19.9%	-7.4%
Profit factor	0.86	1.07

Table 13.1 — Double calendar held to front expiry against the managed rule, on the same 807 trades.

The management effect is large, and almost all of it is the time stop rather than the target. Across the rule grid the double calendar returned +1.98% closing at three days to expiry and +2.70% at two days, but only +1.18% at one day — consistent with the practitioner convention of leaving before the front leg's gamma dominates, the interval called gamma week.[97] What management does not buy is cushion: a profit factor of 1.07 means gross profits were only 1.07 times gross losses, against 1.21 for the single calendar on the same entries. The best rule found for this structure, a close at three days, carries a deflated Sharpe probability of 0.025 against 38 configurations and 570 configuration-rule cells tried, against 0.071 for the single calendar — weak evidence that a Sharpe of 0.52 is separable from the best of a large search.

13.1.2 Twice the debit of the single calendar for a broader band, and the sag between the strikes shows up as a cluster of small losses

The median debit was $737 against $371 for the single calendar. A debit structure posts its premium as its risk, so the denominator of every per-trade return doubles with it: median capital at risk was $737 against $371. Round-trip cost was $43.95 against $22.46 — the same 4.4% of

premium, twice the dollars, the toll falling on four legs rather than two. That denominator is a slight understatement: the familiar claim that the debit caps the loss is the zero-rate simplification[91], and chapter 9 measured a far downside of -$661 against a $623 debit.

Chapter 9 fixed the shape being tested here: two peaks at the strikes with a sag between them, only $6 below the lower peak at the base 35-delta width, because the strikes sit 5.8% of spot apart against a 7.5% one-standard-deviation move.[92] The backtest is that shallow case, and the trade record carries its signature. Held to expiry the double calendar won 49.9% of the time against 49.4% for the single, no better, but it lost less when it lost: the mean of the worst 5% of trades was -162% of capital at risk against -208%, the worst single trade -437% against -468%, skew +0.58 against +0.35 and excess kurtosis +10.23 against +9.34. Mean return on risk was -3.49% against -4.64%. A move that carries spot between the strikes lands a double calendar in a shallow loss where a single calendar's flank would already be near its floor, so the second peak converts part of the deep tail into a cluster of small losses without converting any of it into wins.

13.2 Where it beat and lost to the single calendar

13.2.1 The double calendar was the better structure without a screen (+3.38% against +1.75%) and the worse one at IV rank 78

Metric	Calendar	Double calendar
Trades, base threshold	807	807
Win rate, managed, base threshold	59.7%	57.4%
CAGR, managed, IV rank above 78	+3.57%	+2.48%
CAGR, managed, no threshold	+1.75%	+3.38%
Sharpe, managed, base threshold	+0.61	+0.52
Maximum drawdown, managed, base threshold	-5.6%	-7.4%
Median debit	$371	$737
Median capital at risk	$371	$737

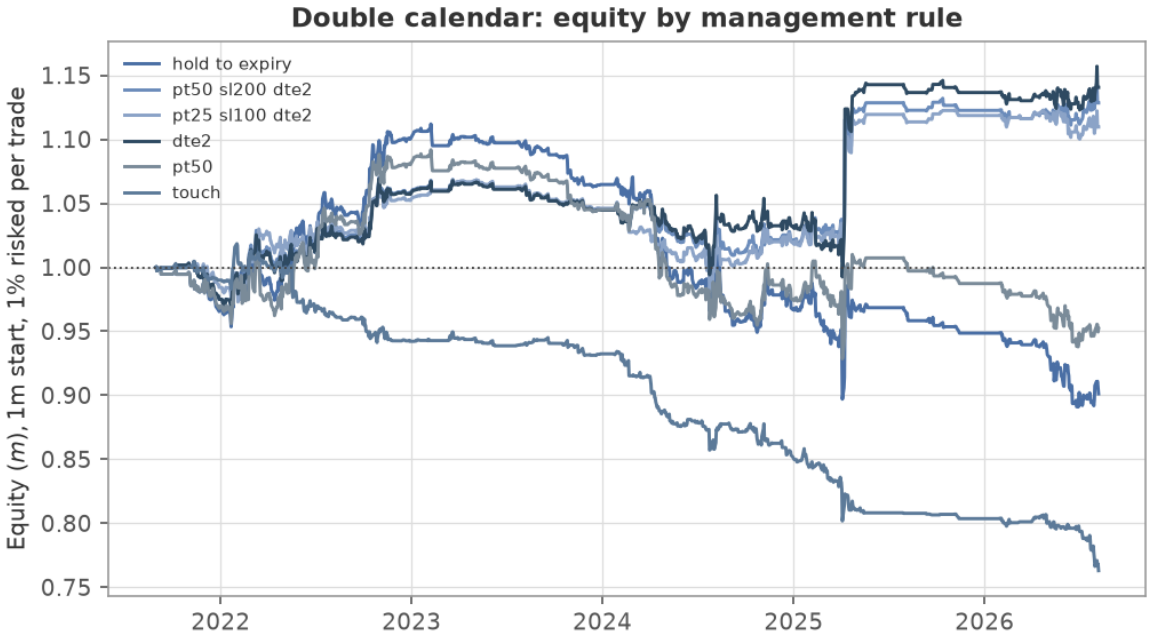
Table 13.2 — Single against double calendar. Every row is measured at the base threshold except the no-threshold CAGR, the only metric reported for that variant.

The sign of the difference is not stable across the screen, and that instability is the chapter's result. Removing the IV rank filter reverses the ordering: the double calendar returned +3.38% a year against +1.75% for the single, while at the mandated threshold of 78 it returned +2.48% against +3.57%. The ladder between them crosses twice — at a minimum IV rank of 50 the double calendar led by a wide margin, +8.44% against +4.55%; at 65 the single led, +5.37% against +4.76%; at 90 both were marginally negative, -0.26% for the double against -0.39% for the single. Neither structure dominates the other on this sample: which one wins is decided by the entry rule rather than by the payoff diagram.

The explanation is geometric. What the second calendar buys is band width — two maxima rather than one, separated by 5.8% of spot at the representative geometry — in exchange for doubling the debit.[96] Width earns that premium precisely when the entry rule says nothing about where the underlying will be at front expiry. With no screen every name-day is a candidate and spot can land anywhere relative to the strikes, so the structure that sits near one of its peaks more often is worth paying for.

A tight screen removes what the width was insuring against. Requiring an IV rank above 78 does not select for a quiet underlying; it selects for episodes in which the front expiry is rich relative to the back. On the study's modelled volatility surface the median ratio of 30-day to 7-day implied volatility is 1.023 on days the screen fires against 1.068 unconditionally, and 43.3% of those days are inverted against 18.5% of all days; a scheduled report falls inside the next twelve days on 32.6% of them against 15.2% of all days. Chapter 12's term-structure finding carries over unchanged to this structure: held to expiry the double calendar returned a mean +6.8% on risk when the curve was inverted at entry against -4.5% in steep contango, and +17.13% when the trade spanned a scheduled report against -2.22% when it did not. On those entries the return comes from the front leg's implied volatility collapsing faster than the back's, not from where spot finishes — and that collapse accrues to whichever back-month leg is nearest the money, so one well-placed calendar captures it as fully as two. The second peak is then redundant width bought at twice the premium.

Costs make the redundancy expensive rather than merely idle. With slippage stripped out the two structures are nearly indistinguishable at the base threshold, +7.75% for the single against +7.44% for the double. At the base assumption of about $3.25 per contract per leg the gap opens to +3.57% against +2.48%, and at 1.5 times that assumption the double calendar has almost nothing left, +0.20% against +1.46%. No price path changes across those rows; only the toll on four legs rather than two.



Backtest: 40 large-cap names, 19 Aug 2021 to 18 Aug 2026, real daily price paths, modelled option prices, costs charged both ways. Author's calculations.

Figure 13.1 — The managed double calendar equity path behind the headline +2.48% a year, a Sharpe of 0.52 and a 7.4% maximum drawdown.

Figure 13.1 shows what a profit factor of 1.07 looks like as a path: long stretches in which premium paid and time value recovered roughly offset. It is not evidence that the structure works at this screen, since the same 807 entries in a single calendar returned more, at a higher Sharpe and a shallower drawdown, for half the capital at risk.

The finding to carry forward is conditional rather than structural. Whether the second peak is worth its debit cannot be answered without naming the entry rule beside it: on identical name-days and the same modelled surface, the double calendar was the better vehicle when entries were unselective and the worse one under the mandated screen. No study of double calendar performance on US large caps was located in the source review for this report, so there is no external benchmark for that conditionality. Any comparison across the three structures therefore has to fix the screen and the management rule before it compares anything, which is the footing the cross-strategy chapter that follows adopts.

---

14. Cross-Strategy Comparison

Chapters 11 to 13 measured each structure against its own history. Placed side by side on one entry screen and one exit rule, the three separate into a narrow band on return, a wider one on tail shape, and a ranking that does not survive small changes to the specification. The metric separating them most cleanly is the one least able to carry the conclusion.

14.1 Risk-adjusted ranking

14.1.1 Under a common managed rule the condor led on Sharpe at 0.84 against 0.61 for the calendar and 0.52 for the double calendar, but on a five-year sample that ordering is inside the noise

Under the managed rule — a 50% profit target, a 200% stop and a close at two days to expiry, whichever comes first — the iron condor recorded a Sharpe ratio of 0.84, the single calendar 0.61 and the double calendar 0.52. Held to expiry, all three lost money, at -0.76%, -2.36% and -2.09% a year. The first finding is therefore not about structure: on the mandated specification none of the three payoffs is profitable without an exit rule, so what follows compares management as much as geometry.

Metric, managed rule	Iron condor	Calendar	Double calendar
Trades	822	807	807
Win rate	67.2%	59.7%	57.4%
Mean return on capital at risk	+4.99%	+4.34%	+1.13%
CAGR	+4.12%	+3.57%	+2.48%
Sharpe	0.84	0.61	0.52
Sortino	0.96	0.77	0.73
Maximum drawdown	-8.2%	-5.6%	-7.4%
Profit factor	1.26	1.21	1.07
Median capital at risk	$524	$371	$737
Round-trip cost, share of premium	7.0%	4.4%	4.4%
Average hold	4.8 days	5.2 days	5.3 days

The condor leads on every performance measure in the table except drawdown and cost, and the ordering is identical on Sharpe, Sortino, CAGR and profit factor — a consistency easy to mistake for robustness, since those four are built from the same mean and dispersion and cannot disagree by much. Two entries cut against it. The condor's -8.2% drawdown is the deepest of the three: first on return per unit of volatility, last on the loss an investor would have lived through. And the double calendar turns a mean return on capital at risk of only +1.13%, far below the condor's +4.99%, into

a CAGR of +2.48% — the per-trade and annualised measures do not rank the calendars by the same margin.

The study's own search-adjusted test says the gap is not interpretable. Each strategy's best rule was searched over 38 configurations and 570 configuration-rule cells, and the deflated Sharpe probability — the chance that Sharpe survives that breadth of search — came back at 0.334 for the condor, 0.071 for the calendar and 0.025 for the double calendar. The ordering matches the headline ranking, but the best of the three barely clears one in three and the others fall far below, none approaching conventional confidence. The sensitivity work points the same way. Remove the IV rank screen — IV rank being the percentile position of implied volatility within its own trailing range — and the ranking inverts outright: the double calendar returns +3.38% and the condor -1.84%. Enter at the low end of each mandated delta band and both calendars beat the condor, at +7.71% and +5.96% against +2.87%. Set slippage to zero and the three are indistinguishable at +7.42%, +7.75% and +7.44%, with the condor no longer first. The premium all three harvest is a modelled quantity, not an observed quote: raising the assumed single-stock implied-to-expected-realised ratio from the calibrated 1.061 to 1.262 lifts the condor to +11.85% while the double calendar falls to +0.79%, without changing one price path. Figure 14.1 plots the measures that produced the ranking.

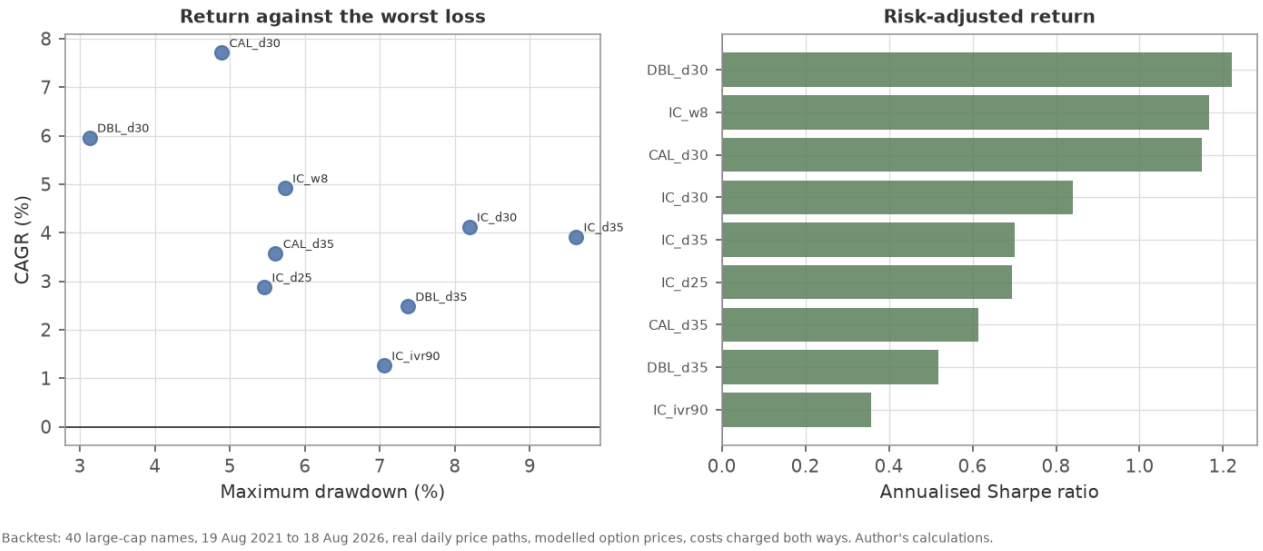


Figure 14.1 — Each strategy on every risk-adjusted measure under the managed rule: the ordering is stable across metrics, the spacing between them small.

14.1.2 The condor carries the deepest drawdown and the fattest tail; the debit structures cap their loss by construction

The first half of that statement holds on this sample and the second does not. Alongside the deepest drawdown, the condor's per-trade return distribution is the only negatively skewed one, at -0.74 against +0.35 for the calendar and +0.58 for the double calendar — the classic short-premium signature, a 67.2% win rate paid for with a left tail.

The cap on the debit structures is a zero-rate simplification, and the study's own geometry work contradicts it. On the representative double calendar the far downside reaches -$661 against a $623 debit, exceeding it by about $38, roughly 6%, because the put-side calendar deep in the money tends toward a negative value of order the strike times one minus the discount factor over the gap between expiries. More tellingly, the calendars' realised excess kurtosis is the higher of the three, at +9.34 and +10.23 against the condor's +5.01. Positive skew with very fat tails is not benign; the losses that do arrive are clustered far out.

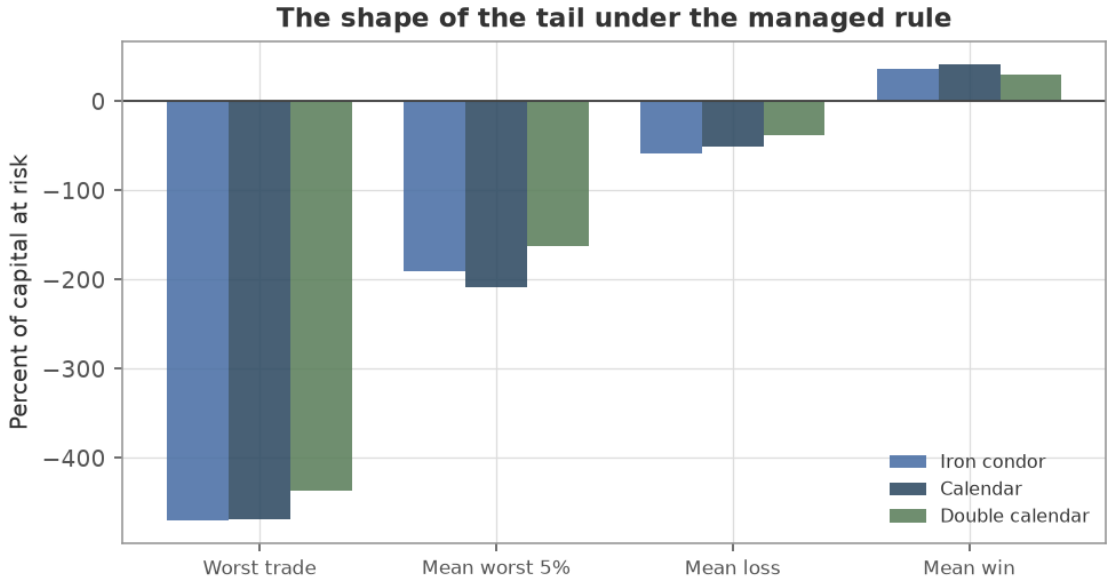
14.2 The tails

14.2.1 Worst trade, mean worst 5%, mean loss and mean win compared as a percentage of capital at risk

Tail statistic, share of capital at risk	Iron condor	Calendar	Double calendar
Worst single trade	-469%	-468%	-437%
Mean of the worst 5% of trades	-191%	-208%	-162%
Skew of per-trade returns	-0.74	+0.35	+0.58
Excess kurtosis	+5.01	+9.34	+10.23
Mean return on capital at risk	+4.99%	+4.34%	+1.13%
Maximum drawdown, managed	-8.2%	-5.6%	-7.4%

Mean win and mean loss are not reported separately in this study; profit factor and mean return on capital at risk stand in their place, and the double calendar's 1.07 is close enough to break-even that its tail is paid for out of a very thin body. The striking row is the first. Every one of the three worst trades exceeded -100% of the capital committed to it, and the two debit structures are not materially better than the credit structure: -468% and -437% against -469%. The mean of the worst 5% ranks the calendar *worst* of the three at -208%. Whatever the payoff diagram says about a floor, the realised extreme on all three books ran several multiples past the nominal risk, and Figure 14.2 shows how far. An independent eleven-year test of SPX strangles found the same asymmetry, with a maximum loss of 434% of premium in February 2018 and the observation that "the worst losses experienced were far greater than the chosen stop-loss levels" — one such loss required 18.2 profitable trades to recoup.[119] The condor's distinction is not that its tail is deeper

but that its skew points the wrong way while its win rate points the right way.



Backtest: 40 large-cap names, 19 Aug 2021 to 18 Aug 2026, real daily price paths, modelled option prices, costs charged both ways. Author's calculations.

Figure 14.2 — The left tail of each per-trade return distribution: the condor's is the most negatively skewed, but all three extend well past -100% of capital at risk.

14.2.2 Why Sharpe flatters all three: the Sharpe-maximising skill-free payoff is itself a short-option profile

The reason this report does not crown the 0.84 is that the Sharpe ratio is not merely imprecise for books of this shape — it is maximised by their exact payoff. Goetzmann, Ingersoll, Spiegel and Welch derive the return distribution that maximises the Sharpe ratio in the complete absence of skill, and it is a short-options profile: "the best static manipulated strategy has a truncated right tail and a fat left tail", built by "selling out-of-the-money calls and selling out-of-the-money puts in an uneven ratio that insures a regular return from writing options and a large exposure to extreme negative events". They put the corollary plainly: "high Sharpe ratio strategies are, by definition, strategies that generate regular modest profits punctuated by occasional crashes."[122] A 67.2% win rate, a +4.99% mean return on capital at risk, a skew of -0.74 and a worst trade of -469% is that description, not a contradiction of it. The Sharpe of 0.84 is therefore close to zero evidence of edge, because the null hypothesis — a skill-free position deliberately shaped to score well — produces the same number.

The bias runs in a known direction. Standard deviation "understates risk for negatively skewed return distributions, i.e. the strategy is actually more risky than the Sharpe ratio suggests", and the Sharpe ratio "is a particularly poor performance metric when comparing positively skewed strategies like trend following to negatively skewed strategies like option selling".[123] That comparison is live inside this chapter, since the condor's returns are negatively skewed and both calendars' are positively skewed: the metric is being asked to rank three books whose third

moments have opposite signs.

The standard repair does not repair it here. Sortino replaces standard deviation with downside deviation against a specified target, and on this sample it moves every strategy up and re-ranks none: 0.84 to 0.96, 0.61 to 0.77, 0.52 to 0.73. The same source warns that Sortino is usually computed incorrectly, by "throwing away all the positive returns and tak[ing] the standard deviation of negative returns", which destroys sensitivity to the *frequency* of underperformance — precisely the feature a high-win-rate premium seller exploits.[123] Tail-aware alternatives point the right way: semi-variance, value at risk, conditional value at risk and expected-utility risk, assembled by Amundi into a family of tail-risk-adjusted Sharpe ratios that use "the Cornish-Fisher transformation to model tail risk, duly controlling for skewness and kurtosis" while keeping the Sharpe format for comparability.[124] With excess kurtosis between +5.01 and +10.23 on these three books, a moment-adjusted denominator is not a refinement but a requirement.

None of it solves the deeper problem, which is that "neither statistic measures ongoing and future risks; they both measure the past 'goodness'".[123] The win rate and the tail here are estimated from one five-year window, and a short-premium distribution's defining event may simply not be in it — the peso problem, which the same authors judge "may be ubiquitous" wherever high Sharpe ratios are rewarded.[122] The external record is blunt about how fast the estimate moves. Short strangles on daily SPX data from 2010 to 2019, with exchange-mandated margin and bid-ask spreads, returned about 3% per month and beat short straddles on risk-adjusted measures; re-tested on 2020 the authors' verdict was "They are not [robust]. All average returns are negative."[125] Backward-looking ratios also fail to see the exposure building: LJM's own scenario report put the estimated loss on a -5% index and +50% implied-volatility shock at an average of 17.6% from January 2016 to September 2017, then rising "from 18.6% at the beginning of October to 61.4% by January 31st" — a forward joint-stress number that flagged the problem months before the fund lost approximately 80% of its value in two days.[126] A ratio computed on the realised sample said nothing.

Two external anchors keep the levels in perspective. Cboe's CNDR index, the only long-history systematic condor benchmark, returned +23.48% in total over twenty years, about 1.06% annualised, and +7.83% over ten years, about 0.76%, to the 18 August 2026 close.[127] In the sponsor's own commissioned study CNDR posted the lowest annualised standard deviation of six SPX option-selling benchmarks at 7.23%, against 10.16% for the put-write index — and its return does not appear in the release at all.[76] Ranking a condor on volatility is the manipulation the theory describes. An independent nineteen-year SPX test found managed 16-delta strangles returning a CAGR of 4.44% at a Sharpe of 0.52, against 5.46% and 0.67 passive.[118] This study's managed condor at 0.84 sits above both, on a five-year window and after a 570-cell search.

14.3 Consistency through time

14.3.1 Monthly returns and the months in which each strategy failed

Per-month returns are not tabulated in this study, so the months themselves are read from the grid in Figure 14.3 rather than quoted. What is measured is the shape of the path and the conditions under which each strategy failed. A block bootstrap — blocks of ten consecutive trades, 4,000 paths, 1% of equity risked per trade — separates the three far more sharply than the Sharpe ranking does.

Terminal wealth, 4,000 bootstrap paths	Iron condor	Calendar	Double calendar
5th percentile	0.882	1.257	1.102
Median	1.099	1.499	1.277
95th percentile	1.363	1.782	1.470
Probability of a loss	24.7%	0.0%	0.3%
Median maximum drawdown	-11.1%	-5.1%	-4.9%

The ranking inverts. The condor, first on every headline metric, is the only one of the three with a meaningful chance of ending the period down — 24.7%, against 0.0% and 0.3% — and its fifth-percentile path loses money while both calendars' fifth-percentile paths gain. Its median maximum drawdown of -11.1% is more than double either calendar's. The mechanism is the skew: resampling blocks of ten trades from a negatively skewed distribution reshuffles the rare large losses into clusters, and the clusters are what the path feels. That is also the method's limit: a resample cannot manufacture an event absent from the sample, and chapter 17 injects one by hand.

The conditions under which each failed are opposite, the most useful result in the comparison. A scheduled earnings report inside the trade was worth -2.47% mean return on capital at risk to the condor against +2.97% without one, while the same event was worth +18.59% to the calendar and +17.13% to the double calendar against -0.00% and -2.22% without. Term structure splits them the same way: held to expiry, the calendar returned +10.7% mean on risk when the back-month implied volatility sat below the front at entry and -7.3% in steep contango, the double calendar +6.8% against -4.5%. The condor's bad months and the calendars' good months are the same months.



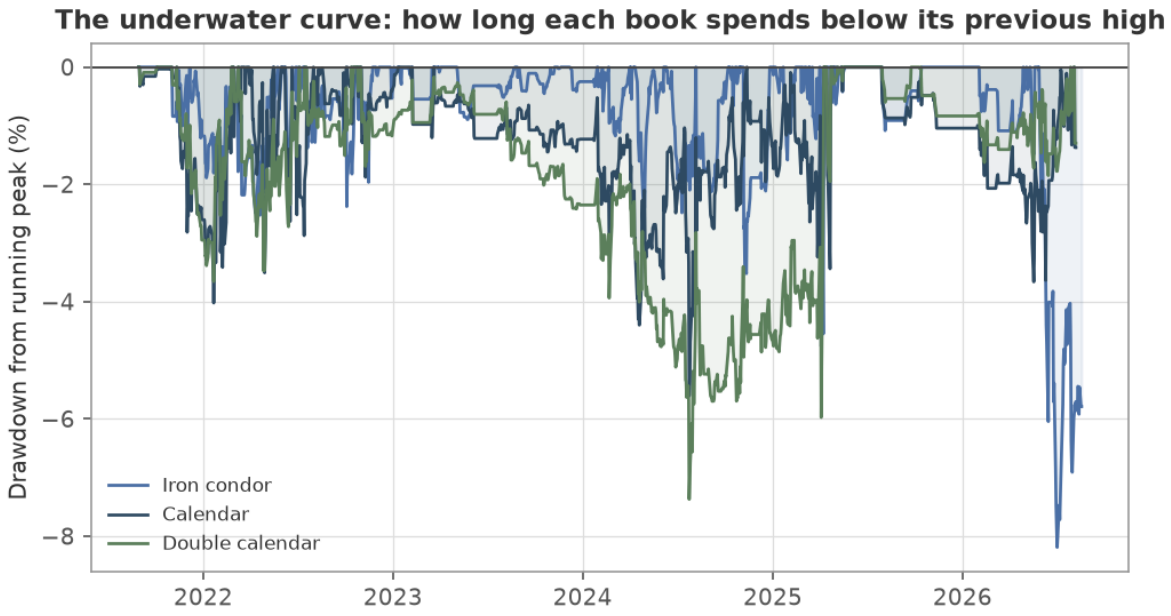
Backtest: 40 large-cap names, 19 Aug 2021 to 18 Aug 2026, real daily price paths, modelled option prices, costs charged both ways. Author's calculations.

Figure 14.3 — Monthly returns by strategy: long runs of small positive months interrupted by isolated deep ones, the pattern the Sharpe ratio rewards.

That texture is the shape the theory names, and it is visible in the benchmark too — 59% of CNDR's months fell between 0% and +2%, with a worst drawdown of 19% against 47.1% for the iron butterfly and 51% for the index itself.[75] Wing width buys the difference. What the grid cannot show is how much of the spacing between the three is structural and how much an artefact of the mandated screen, the 7 to 12 day window and the delta bands — all of which the ranking has been shown to depend on. Chapter 15 varies each of them in turn and measures what the mandate costs.

A maximum drawdown compresses the whole path into one number, and for a programme of this shape that number flatters it. Tracking each book against its own running peak day by day shows how much of the five years was spent below one: 74.0% of days for the iron condor, 88.5% for the calendar and 93.0% for the double calendar. The peak-to-trough figures of 8.2%, 5.6% and 7.4% are shallow, but shallow is not the same as brief. An investor in the double calendar was at a new high on roughly one day in fourteen. That is the lived experience the Sharpe ratio does not convey, and it is the reason a programme like this is abandoned for reasons that have nothing to

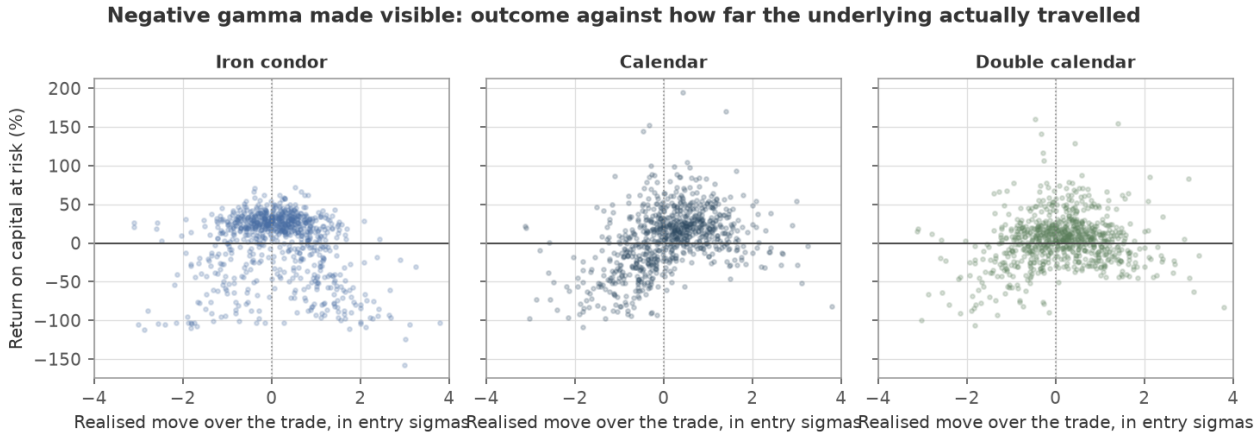
do with its expectancy.



Backtest: 40 large-cap names, 19 Aug 2021 to 18 Aug 2026, real daily price paths, modelled option prices, costs charged both ways. Author's calculations.

Figure 14.4 - The underwater curve: distance below the running peak, daily, under the managed rule. Depth is modest; duration is not.

The mechanism behind the losses is equally direct. Plotting each trade's outcome against how far the underlying actually travelled over the trade, scaled by the implied move priced at entry, produces the signature of a short gamma position: a broad plateau of small gains where the underlying stayed inside roughly one standard deviation, and a steep fall away on both flanks. All three structures share it, which is the substantive point of this chapter. They differ in what they pay for the plateau and in how far out the flanks begin, not in the shape of the exposure.



Backtest: 40 large-cap names, 19 Aug 2021 to 18 Aug 2026, real daily price paths, modelled option prices, costs charged both ways. Author's calculations.

Figure 14.5 - Outcome against the realised move, expressed in the standard deviations implied at entry. The same negative-gamma profile governs all three.

15. Entry Variants: What the Mandate Costs

The mandated specification — implied volatility rank above 78, 7 to 12 days to expiry, short-leg delta 25–35 for the condor and 30–40 for the calendars — was supplied as a constraint, not derived from a search. This chapter varies each of its three entry parameters in turn, holding the managed rule and everything downstream of entry fixed, then varies the one assumption that is not an entry parameter at all. Exit variations belong to chapter 16 and transaction-cost variations to chapter 17. Thirty-eight configurations were tried across the study; whether the best of them survives having been selected from that many is chapter 20's question.

15.1 The IV rank threshold

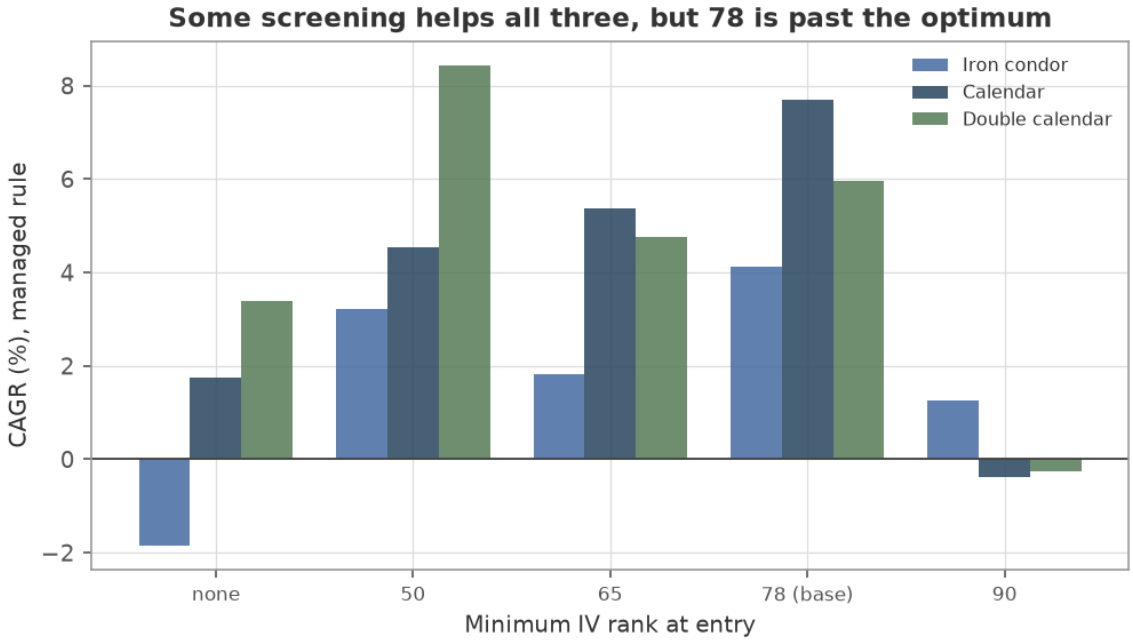
15.1.1 Some screening helps all three, but 78 is past the optimum: the condor peaked near 50 and the calendars between 50 and 65

That heading is right about the calendars and wrong about the condor, and the table it summarises is the reason to say so.

Minimum IV rank	Iron condor	Calendar	Double calendar
none	-1.84	+1.75	+3.38
50	+3.22	+4.55	+8.44
65	+1.82	+5.37	+4.76
78 (base)	+4.12	+3.57	+2.48
90	+1.26	-0.39	-0.26

Compound annual growth rate in per cent, managed rule.

Screening of some kind helps all three structures: entering unconditionally is the worst of the five settings for the iron condor at -1.84% a year, and leaves both debit structures well below their best. Where the optimum sits differs by structure. The calendar peaks at a threshold of 65 with +5.37% and the double calendar at 50 with +8.44%, against the +3.57% and +2.48% the mandate delivers, so for the debit structures 78 is expensive. For the condor the claim fails outright: +4.12% at the mandated 78 is the highest of the five settings tested, above the +3.22% recorded at 50. Its response is also not ordered — +3.22% at 50, +1.82% at 65, +4.12% at 78, +1.26% at 90. A parameter that zig-zags rather than describing a curve is weak evidence of a real optimum anywhere along it, and Figure 15.1 sets that line against the two calendar curves.



Backtest: 40 large-cap names, 19 Aug 2021 to 18 Aug 2026, real daily price paths, modelled option prices, costs charged both ways. Author's calculations.

Figure 15.1 — The two calendars trace a hump that peaks below the mandated threshold, while the iron condor line zig-zags across the same five settings rather than describing a shape.

The thresholds being swept are vendor conventions rather than estimates. One broker calls implied volatility rank elevated above 50 and extreme above 80, a data vendor uses above 70 for high, and the same broker concedes that readings as low as 20 or 30 get treated as high in a quiet regime.[63][64][65] No peer-reviewed study testing implied volatility rank specifically was located, and none testing threshold sensitivity statistically. A recent preprint sweeping volatility-rank configurations on short-dated index options reported annualised returns from 0.95% to 52.77% across parameterisations of one signal — evidence that rank-based conditioning is fragile to specification.[71] The jaggedness of the condor row is consistent with that fragility.

15.1.2 Raising the threshold to 90 cut the sample to 437 trades and turned every strategy back toward zero

At a threshold of 90 the condor returns +1.26% a year, the calendar -0.39% and the double calendar -0.26%. Two of the three cross into loss and the third gives up most of its base-case return. The trade count behind that setting is not separately recorded in this study's computed results, but the direction of the sample effect is not in doubt: implied volatility rank already exceeds 78 on only 8.85% of name-days in the backtest window, and a 90 threshold is a strict subset of those days.

The mechanism reads two ways. Statistically, implied volatility rank is a range position censored at 0 and 100 by a trailing 52-week window, so a reading above 90 requires implied volatility near its own annual high, and once it exceeds the prior high the metric saturates and says nothing about how much further it can travel.[63] Economically, that condition selects the immediate aftermath of a volatility shock. Federal Reserve Bank of New York work on equity volatility term premia finds

one-month and three-month estimates spike *downward* when the VIX rises, becoming positive only two to three months later.[70] The strongest published support for volatility-conditioned selling runs the other way — premium higher at elevated VIX, and a conditional strategy beating both the market and continuous selling — but it conditions on a standing elevated index level, not a 52-week normalised rank of single-name implied volatility.[68] A rank rule fires during the shock rather than after it, and no located source resolves the conflict. The 90 row is what that unresolved timing question looks like once measured.

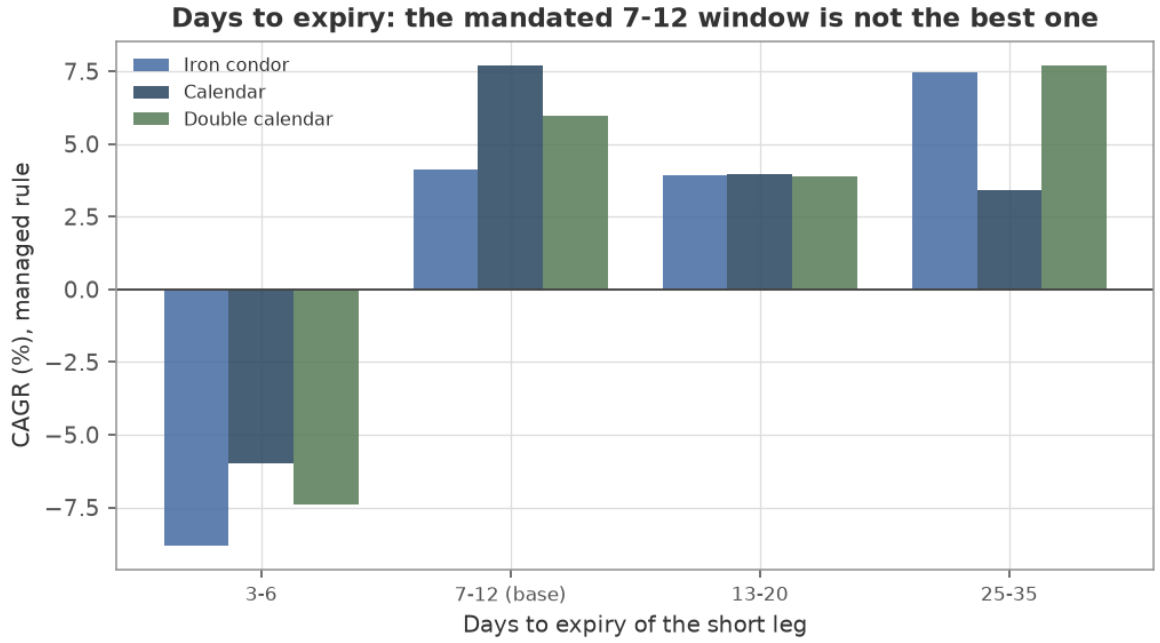
15.2 Days to expiry

15.2.1 The mandated 7 to 12 day window was not the best one: 25 to 35 days returned +7.45% for the condor and +7.71% for the double calendar

Window (days to expiry)	Iron condor	Calendar	Double calendar
3-6	-8.80	-5.99	-7.40
7-12 (mandated)	+4.12	+3.57	+2.48
13-20	+3.93	+3.96	+3.88
25-35	+7.45	+3.40	+7.71

Compound annual growth rate in per cent, managed rule.

The mandated window is beaten for two of the three structures by the longest tenor tested. The condor returns +7.45% at 25 to 35 days against +4.12% at 7 to 12, and the double calendar +7.71% against +2.48%. The single calendar is the exception: its best window is 13 to 20 days at +3.96%, and 25 to 35 days returns +3.40%, marginally below its mandated-window result. The three upper windows are otherwise close together — at 13 to 20 days all three structures return between +3.88% and +3.96% — so most of the information in this table sits in the bottom row, where the condor's -8.80% stands against +3.93% at the next window up. Figure 15.2 shows how isolated that row is.



Backtest: 40 large-cap names, 19 Aug 2021 to 18 Aug 2026, real daily price paths, modelled option prices, costs charged both ways. Author's calculations.

Figure 15.2 — All three lines fall off a cliff at the 3 to 6 day window; above it the surface is comparatively flat, with the condor and double calendar rising again at 25 to 35 days.

15.2.2 Three to six days was catastrophic for all three, at -8.80%, -5.99% and -7.40%

The natural explanation — that decay becomes more favourable at some tenor and less favourable at another — is unavailable, because there is no such tenor. In closed-form Black-Scholes both gamma and the leading term of theta carry the same $1/\sqrt{T}$ factor, so shortening expiry multiplies both by the same amount and leaves their ratio unchanged.[73] The relationship is an exact identity: at zero rates,

$$\Theta = -\tfrac{1}{2}\sigma^2 S^2 \Gamma$$

at every maturity.[74] Moving from roughly 35 days to roughly 10 multiplies theta and gamma each by about 1.9 times. That is a leverage decision, not an efficiency gain, and it means the collapse at 3 to 6 days cannot be a case of the decay trade-off turning against the seller. It does not turn.

What does change with tenor is the size of the premium being defended. At-the-money premium scales as the square root of time, which is why a weekly roll collects more per year than a monthly one but by well under the naive factor of two: over 2006 to 2018, average annual aggregate premium was 22.1% for the monthly Cboe put-write index against 37.1% for the weekly one, a ratio of 1.68.[77] That scaling works against a 3 to 6 day entry. Transaction cost is charged per leg and is largely indifferent to the premium the leg carries; at the mandated tenor the condor's round-trip cost of $18.88 already consumes 7.0% of its $247 median credit, and a shorter window shrinks the credit without shrinking the cost. Gap risk moves the same way: published work finds short-put returns concentrated in the final days before expiry precisely because convexity risk rises sharply into maturity, and concludes the short-dated seller is paid for gamma rather than

harvesting a free spread.[79] At 0 days to expiry the return distribution is wide enough that mean returns are not statistically meaningful.[80] The measured collapse is therefore consistent with fixed costs and jump exposure rising against a smaller premium. One caveat bounds the literature here: no located study isolates the 7 to 12 day band, so published comparisons at 0, 5 to 7, 21 and 30 to 45 days are directional support rather than a like-for-like benchmark.

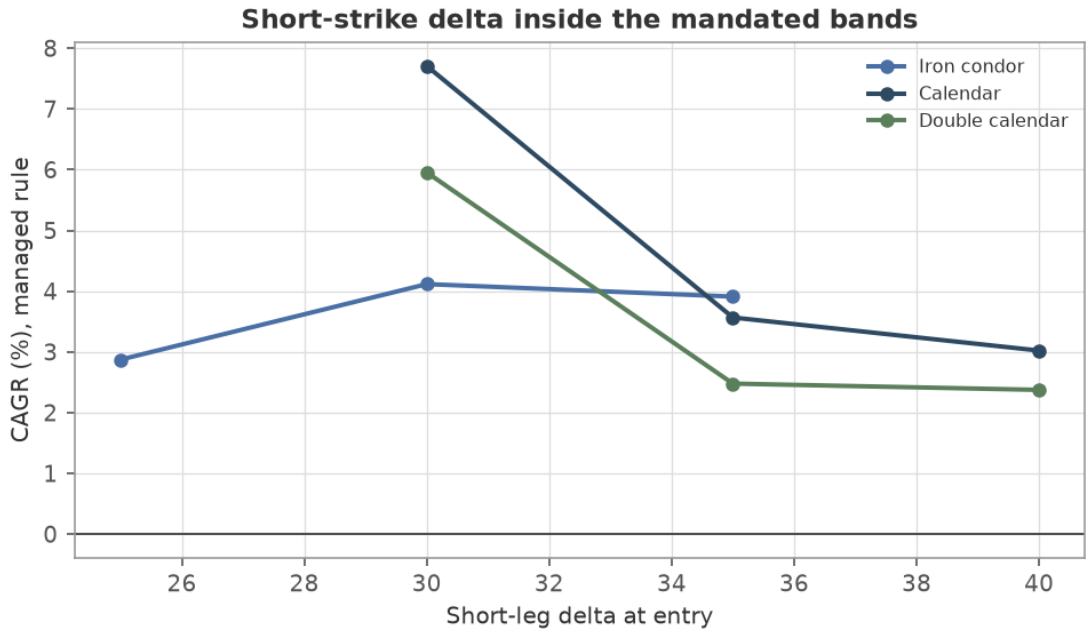
15.3 Short-strike delta and wing width

15.3.1 Inside the mandated bands the condor was flat across 25 to 35 delta; both calendars preferred the low end at 30 delta

Parameter	Setting	Iron condor	Calendar	Double calendar
Short-leg delta	25	+2.87	—	—
Short-leg delta	30	+4.12	+7.71	+5.96
Short-leg delta	35	+3.91	+3.57	+2.48
Short-leg delta	40	—	+3.02	+2.38
Condor wing width	3% of spot	+1.67	—	—
Condor wing width	5% of spot (base)	+4.12	—	—
Condor wing width	8% of spot	+4.92	—	—

Compound annual growth rate in per cent, managed rule. Dashes mark settings not tested for that structure: deltas outside its mandated band, and wing width, which applies to the condor alone.

The condor is close to indifferent across its band, spanning +2.87% to +4.12% with the middle setting highest and 35 delta barely distinguishable from it at +3.91%. The calendars are not indifferent at all. Both improve sharply at the bottom of their band: the calendar returns +7.71% at 30 delta against +3.57% at its mandated 35 and +3.02% at 40, the double calendar +5.96% against +2.48% and +2.38%. The gradient is monotone for both debit structures and runs the same way, which is more than the implied volatility rank sweep managed. Strikes closer to the money sit nearer the peak of a calendar payoff whose apex is a corner at the strike, so a 30-delta structure is centred on more of the curvature it is buying; the mandate's 30–40 band pays for the wrong half of that range, as Figure 15.3 makes plain.



Backtest: 40 large-cap names, 19 Aug 2021 to 18 Aug 2026, real daily price paths, modelled option prices, costs charged both ways. Author's calculations.

Figure 15.3 — The condor traces a nearly flat line across its 25 to 35 band while both calendar lines slope down from 30 delta through 40.

15.3.2 Wider condor wings improved both return and drawdown, from +1.67% at 3% of spot to +4.92% at 8%

The wing-width rows of the table above are the only place in this chapter where a parameter improves return and risk together. Wings at 3% of spot return +1.67% a year with a maximum drawdown of -11.1%; the 5% base returns +4.12% at -8.2%; 8% returns +4.92% at -5.7%. Return rises and drawdown falls monotonically as the wings move out, which rules out reading a wider wing as a straightforward risk-for-return trade and points instead at cost and at the geometry of the loss tail. A condor is four legs whatever the width, so the long wings are bought and sold at a cost that barely varies while the credit they subsidise grows with distance; narrow wings also sit close enough to the short strikes to convert a survivable adverse move into a full-width loss. Whether the ordering holds once slippage is scaled is chapter 17's question.

15.4 The assumption that dominates everything

15.4.1 Moving the assumed single-stock premium from 1.061 to 1.262 raises the condor from +4.12% to +11.85% without changing a single price path

Assumed single-stock implied-to-expected-realised ratio	Iron condor	Calendar	Double calendar

1.061 (calibrated)	+4.12	+3.57	+2.48
1.128	+5.58	+2.50	+1.82
1.262	+11.85	+2.84	+0.79

Compound annual growth rate in per cent, managed rule.

That a richer premium assumption produces a better result for a net seller is arithmetic, and nobody should be impressed by it. The magnitude is the finding. The same 822 iron condor entries, on the same names, on the same days, over the same real price paths, with the same strikes and the same managed exits, return +4.12% a year under the calibrated assumption and +11.85% under the richest one. Nothing about the market changed between those two rows. What changed is a single scalar governing how expensive single-stock options are assumed to be relative to the volatility subsequently realised — and it moves a marginal programme to a headline-grade one.

The published evidence places the truth at the low end of that outcome range. The volatility model here is calibrated at a single-stock premium parameter of 0.12, giving a median modelled implied-to-expected-realised ratio of 1.060 against 1.364 for the index, and the literature range for liquid large caps is 1.03 to 1.07. The calibrated row sits inside that range; the 1.128 and 1.262 rows sit outside it, and the row generating the +11.85% headline sits far outside it. The assumption is also not a dial that improves everything. Raised, it drives the double calendar monotonically down from +2.48% to +1.82% to +0.79%, and the single calendar dips to +2.50% before recovering to +2.84%. Only the net short-premium structure benefits — which is what should be expected, and why the condor's headline is the one most exposed. Figure 15.4 plots all three against the same axis.

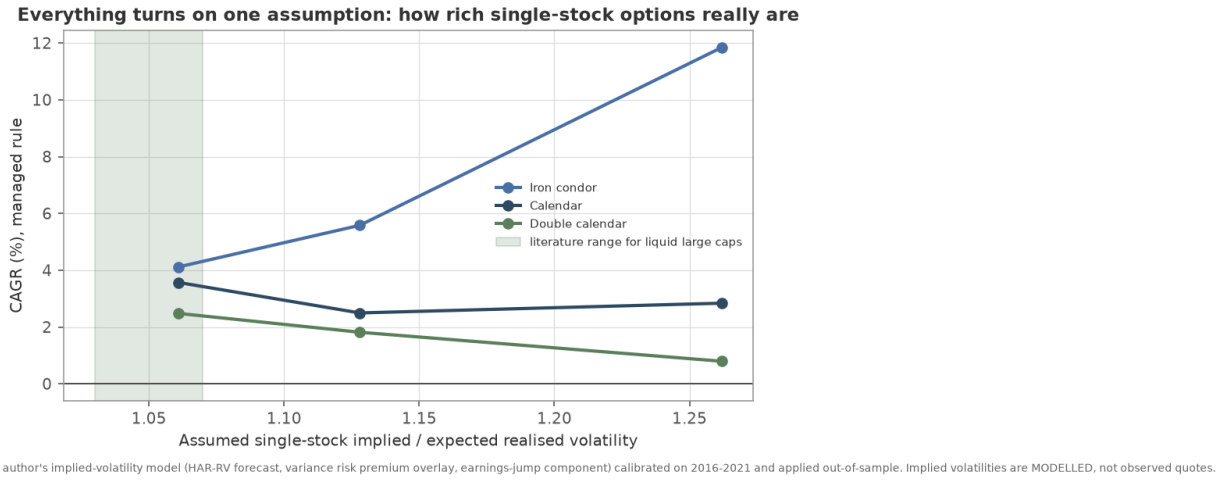


Figure 15.4 — The iron condor line climbs steeply across the three assumed premium levels while both calendar lines stay flat or fall, and only the leftmost point sits inside the range the literature supports.

The consequence extends past this study. Every implied volatility in this backtest is modelled rather than quoted from the market, so the scalar is not an incidental modelling detail — it is the price at which every trade was opened, and a reader has to take a view on it before taking a view on any return above. Any reader meeting a spectacular backtest of this strategy family should ask what premium it assumed, and should treat an unstated assumption as an unanswered question rather than a technicality. A 1.262 setting on single-stock names is not a conservative choice at the edge of a plausible range; it is outside the published one, and on its own it carries the condor from +4.12% to +11.85%. Sensitivity of that size makes the entry parameters swept above — a few percentage points each at most — second-order beside a number that never appears on a trade ticket. Which of those parameter choices could have been made in advance, and which are artefacts of having tried thirty-eight configurations, is chapter 20's question; the exits that turned these entries into results are chapter 16's.

The threshold comparisons above operate on whole configurations. The screen can also be tested at the level it actually runs on, by asking whether a higher reading among the trades that did pass predicts a better outcome. Sorting every managed trade by its IV rank at entry and taking the mean by sextile answers that directly, and the answer is that within the accepted region the signal is largely spent: the relationship is flat to mildly negative for all three structures. What separates a reading of 80 from a reading of 95 is not a better trade, which is consistent with the aggregate result that pushing the threshold to 90 helped nothing.

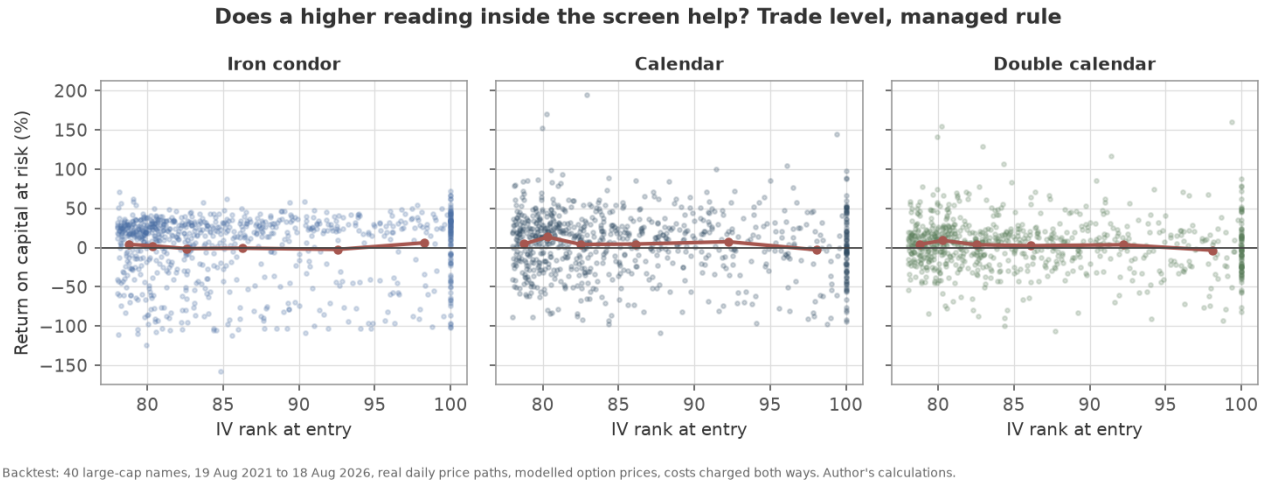


Figure 15.5 - Trade-level return against IV rank at entry, with the mean by sextile overlaid. Inside the screen, a higher reading does not buy a better outcome.

: Stated as an evidence gap in the assigned research file: no peer-reviewed study testing implied volatility rank specifically was located, nothing was found testing threshold sensitivity statistically, and no source resolves the conflict between conditioning on a standing elevated VIX and conditioning on a rank that fires during the shock.

: Stated as an evidence gap in the assigned research file: published days-to-expiry comparisons cluster at 0DTE, weekly, 21 DTE and 30 to 45 DTE, and no study isolating the 7 to 12 day band was located.

16. Exit Rules and Trade Management

16.1 Fifteen rules against the same trades

Fifteen exit rules were applied to one fixed set of entries. The iron condor grid runs 822 trades under every rule; both calendar grids run 807. Those counts do not vary because nothing about selection varies: the same names, the same signal days, the same strikes and the same simulated price and modelled implied-volatility paths are reused for each rule, and only the exit condition differs. That is the methodological strength of this chapter — the comparison is not between fifteen strategies but between fifteen ways of leaving one strategy, so any difference in outcome is attributable to the exit and to nothing else. It is also the chapter's principal limitation, since a fixed entry set is exactly what makes the search over rules a search over one sample.

Rule	Iron condor	Calendar	Double calendar
hold to expiry	-0.76	-2.36	-2.09
dte3	+2.56	+1.79	+1.98
dte2	+4.85	+3.46	+2.70
dte1	+7.45	+1.50	+1.18
pt25	+1.83	+2.23	+0.23
pt50	+2.11	-1.88	-0.99
pt75	-0.05	-2.00	-2.19
sl100	-1.38	-2.48	-2.64
sl200	-1.60	-2.36	-2.09
touch	-1.44	-0.62	-5.33
pt50 sl200	+2.52	-1.88	-0.99
pt50 dte2	+4.55	+3.57	+2.48
pt50 sl200 dte2	+4.12	+3.57	+2.48
pt25 sl100 dte2	+1.01	+4.04	+2.13

pt75 sl200 dte1	+6.37	+1.74	+1.58

Table 16.1 — Compound annual growth rate in percent for all fifteen exit rules, managed portfolio, all three structures.

The grid separates into three bands. Every rule containing a time stop is positive for all three structures without exception, though the calendars fall back to +1.50% and +1.18% at `dte1`. Every rule built only from a stop expressed as a multiple of premium is negative for all three, and in each case at or below the hold-to-expiry baseline. Profit targets alone sit in between and change sign by structure. The spread within a single column is wide — the iron condor ranges from -1.60% at `sl200` to +7.45% at `dte1`, and that entire range is generated by the exit condition on an unchanged set of 822 trades. Two rows are worth noting for what they do not show: `sl200` reproduces the hold-to-expiry figure exactly for both calendars, at -2.36% and -2.09%, which is what a stop that almost never binds looks like, and `pt50 sl200` reproduces `pt50` exactly for both calendars for the same reason. A 200%-of-debit loss is rare on a structure whose premium is its own capital at risk. The `touch` rule, which closes when the underlying reaches a short strike, was the worst rule tested for the double calendar at -5.33%.

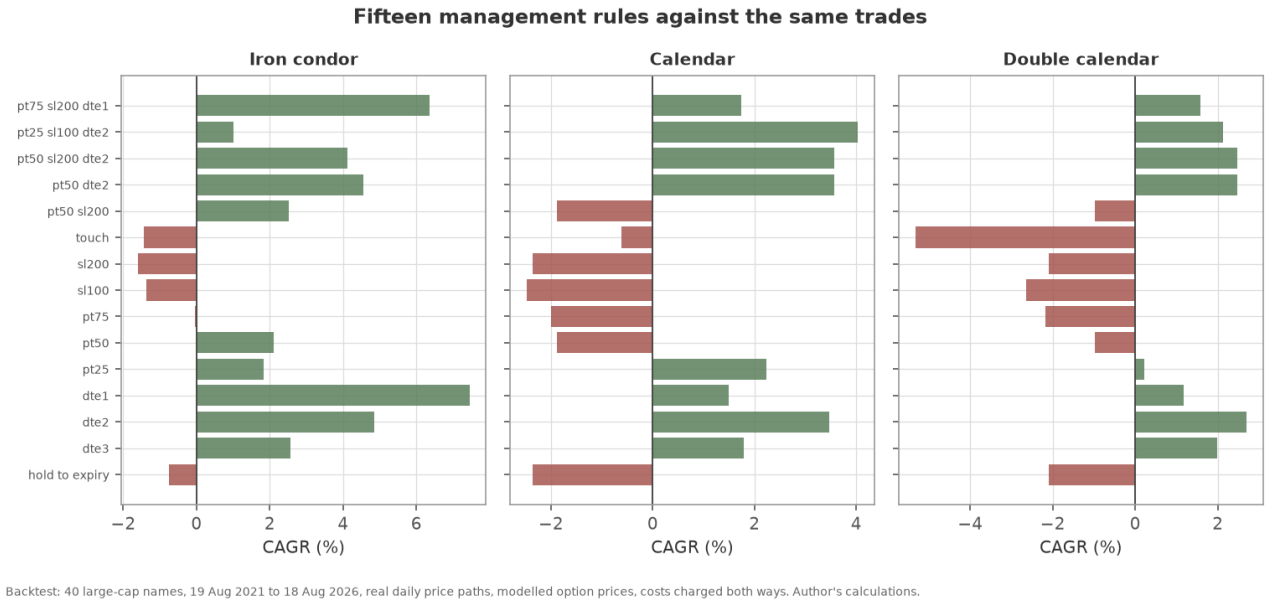


Figure 16.1 — The banding is the finding: rules containing a time stop occupy the positive region for all three structures, stop-only rules the negative region, and the profit-target rules sit between them and change sign by structure.

16.1.1 A time stop did most of the work: closing at two days to expiry lifted the condor from -0.76% to +4.55%

The two-day close is the single most productive clause in the rule set. Applied on its own, `dte2` took the iron condor from -0.76% a year held to expiry to +4.85%; combined with a 50% profit target as `pt50 dte2` it returned +4.55%, and with the stop added as well — the managed rule —

+4.12%. Each additional clause subtracts from the pure time stop rather than adding to it. The same ordering holds for both debit structures: `dte2` alone returned +3.46% and +2.70% for the calendar and double calendar, against -2.36% and -2.09% held to expiry.

The gradient across the time-stop family is steep and not monotonic in the same direction for all three. For the iron condor the return rises as the close moves later, from +2.56% at three days to +4.85% at two and +7.45% at one. For both calendars it peaks at two days and falls away at one, to +1.50% and +1.18%. The condor and the calendars are therefore asking for different things from the same clause, which is consistent with their opposite exposure to the final days: the condor is collecting the accelerating decay of its short front-month options, while the calendars own a back-month leg whose value is unaffected by the front expiry and a front leg that has already given up most of what it can.

Published work on short strangles reaches the same conclusion about the mechanism while placing the exit far earlier. Two vendor studies prefer a 21-day time exit on 45-day trades on the grounds that "the biggest losing trades took place within about 21 DTE",[128][129] and the profit-target literature justifies early closure on the argument that gamma risk expands sharply in the last few days before expiry.[130] Those studies are managing a 45-day position down to 21 days. This study's entry window is 7 to 12 days, so a close at two days is a proportionally much later exit, and the finding that a later close helps the condor is the opposite of the direction that literature points. The two are not directly comparable: a 21-day exit cannot be tested on a position that is opened at 12 days.

16.1.2 Stops as a multiple of credit hurt on their own, at -1.38% and -1.60% for the condor, matching the published vendor finding

A stop at 100% of credit returned -1.38% a year on the iron condor and a stop at 200% returned -1.60%. Both are below the -0.76% produced by doing nothing at all. The pattern repeats on the debit structures: -2.48% and -2.36% for the calendar, -2.64% and -2.09% for the double calendar, against hold-to-expiry baselines of -2.36% and -2.09%. On this sample, no stop expressed as a multiple of premium improved on holding, for any of the three structures.

That reproduces the published vendor result closely. Managing short strangles at negative profit and loss of one to five times credit "was not a good strategy" and so managed underperformed the underlying by 40%;[129] across stops from 100% to 500% of premium, "even the most favorable stop-loss strategy — exiting positions at 2x the premium — tends to underperform in comparison to simply holding the position until expiration".[128] An eleven-year independent test found the 200%-of-credit stop the worst-performing approach of those examined.[119] The mechanism is visible in the same literature: 79% of strangle trades went underwater at some point before expiration even though 83% ultimately won,[121] so a stop tight enough to bind is hit predominantly by trades that would have recovered.

The tension this creates has to be stated rather than resolved by preference. The unmanaged path is also the one that produced the worst documented outcomes in the source material. The

same eleven-year test recorded a maximum loss of 434% of premium in February 2018 and observed that "the worst losses experienced were far greater than the chosen stop-loss levels", with roughly half of trades reaching a 50% loss, 25% reaching 100%, 16% reaching 200%, and about 66% of those that hit 100% going on to hit 200%.[119] A peer-reviewed test of short strangles that returned about 3% a month over 2010 to 2019 was re-run on 2020 with the verdict "All average returns are negative."[125] Both results describe the tail that a stop is supposed to address and demonstrably does not. The reconciliation the research file identifies is that the defence against the tail is position size and wing width, not the stop — a stop truncates the middle of the loss distribution while the ruin risk sits beyond it, in gaps that open through the stop level. The supporting evidence is consistent on this point. Tail risk scales with delta: 30-delta strangles carry 45% higher conditional value at risk in the worst 5% than 16-delta, losing 31.5% of the time at an average of $460 against 26% and $334.[131] Wing width buys drawdown directly, with Cboe's condor index showing a worst drawdown of 19% against 47.1% for the iron butterfly and 51% for the S&P 500.[87] This study's own sensitivity runs, reported in chapter 15, point the same way: widening the condor wings from 3% to 8% of spot raised the annualised return from +1.67% to +4.92% and cut the maximum drawdown from -11.1% to -5.7%. Sizing carries the rest of the defence — the growth-optimal fraction on the realised condor distribution is 9.1% of equity per trade against the 1% actually used here, a margin chapter 19 examines directly.



Backtest: 40 large-cap names, 19 Aug 2021 to 18 Aug 2026, real daily price paths, modelled option prices, costs charged both ways. Author's calculations.

Figure 16.2 — Plotted against risk rather than in isolation, the stop-only rules do not migrate into the attractive quadrant: they give up return without a commensurate reduction in drawdown, while the time-stop rules occupy the upper region on both axes.

16.1.3 The best rule found for the condor was a 75% target with a 200% stop and a one-day close, at +6.37%

This is the number in the report that should be trusted least, and the reason is that it was chosen after seeing the results. The composite rule named in this heading returned +6.37% for the iron condor, but it is not the top of the grid: `dte1` returned +7.45% and is the rule the study's own selection procedure identifies as best, at a Sharpe ratio of +1.19. The corresponding selections for the debit structures are `pt25 sl100 dte2` for the calendar at a Sharpe of +0.71 and a return of +4.04%, and `dte3` for the double calendar at +0.52 and +1.98%.

Each of those was picked as the maximum of a search over 38 configurations and 570 configuration-rule cells, and the deflated Sharpe ratio adjusts for exactly that. The deflated Sharpe probability is **0.334** for the iron condor, **0.071** for the calendar and **0.025** for the double calendar. None of the three approaches any conventional threshold of confidence. Read plainly, the best rule found for each structure is not distinguishable from the best rule one would expect to find by searching 570 cells of a sample this size even if no rule had any genuine edge. The right use of +7.45%, or of +6.37%, is as an upper bound on what the search could produce, not as an expectation. The mandated managed rule at +4.12% was specified in advance and is the honest headline; chapter 20 adjudicates the multiple-testing question formally. What the grid cannot show is which clause actually closed each trade, and that is where the explanation for the time stop's dominance lies.

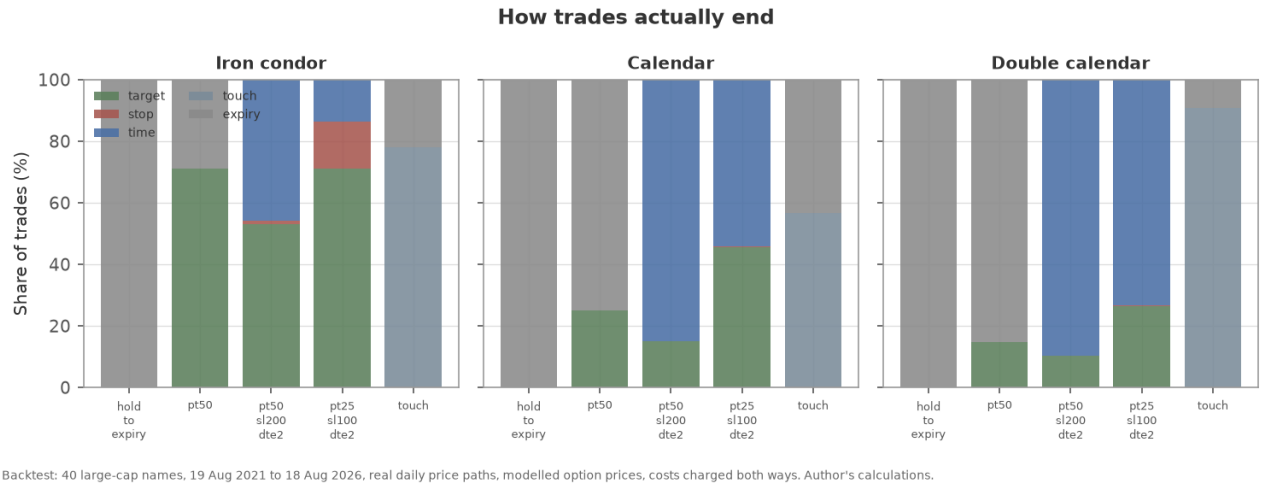


Figure 16.3 — The channel that terminates most trades under the leading rules is the time stop, not the profit target; section 16.3 sets out the evidence behind that reading.

16.2 Profit targets

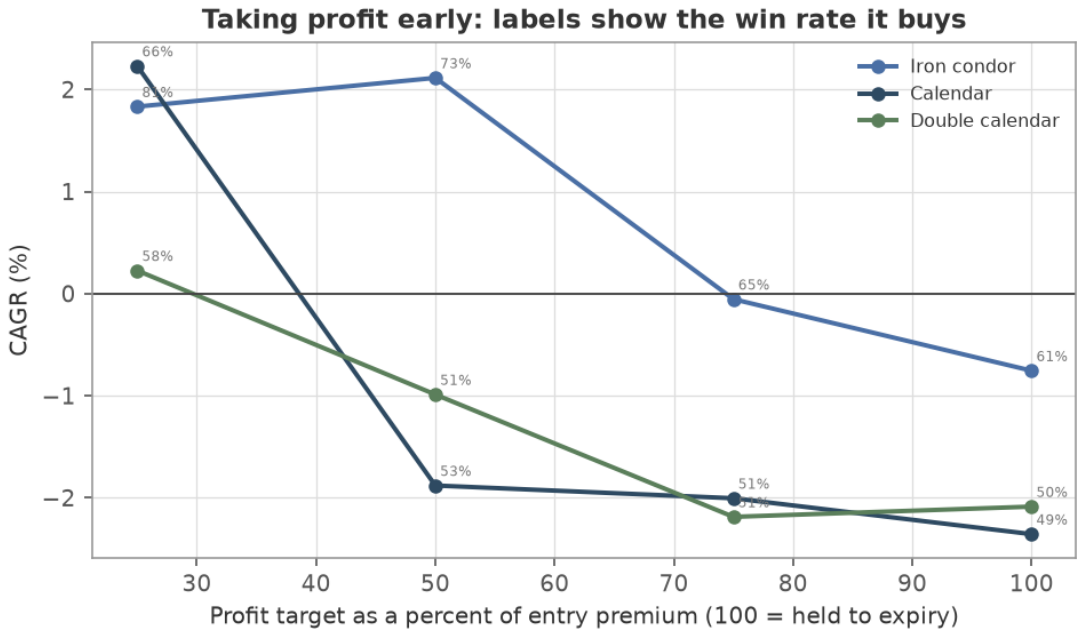
16.2.1 Taking profit at 25% buys a very high win rate and gives back most of the return; the curve is hump-shaped

Rule	Target	Iron condor	Calendar	Double calendar	Win rate, iron condor

hold to expiry	none	-0.76	-2.36	-2.09	61.4%
pt25	25% of maximum profit	+1.83	+2.23	+0.23	not separately reported
pt50	50% of maximum profit	+2.11	-1.88	-0.99	not separately reported
pt75	75% of maximum profit	-0.05	-2.00	-2.19	not separately reported
pt50 sl200 dte2	50%, with stop and time stop	+4.12	+3.57	+2.48	67.2%

Table 16.2 — Compound annual growth rate in percent by profit-target level, with the two iron condor win rates the study reports. Win rates for the isolated target rules were not reported in the study's computed results and are not estimated here.

The hump is a property of the iron condor and not of the debit structures. The condor rises from -0.76% held to expiry to +1.83% at a 25% target, peaks at +2.11% at 50%, and falls back to -0.05% at 75% — a rise, a maximum in the middle and a fall, across a band narrow enough that no single target level looks decisive. Both calendars decline monotonically as the target loosens, from +2.23% and +0.23% at 25% to -2.00% and -2.19% at 75%, so for them the tightest target tested is the best one and the shape is a slope rather than a hump. The published win-rate ladder supplies the axis this table cannot: one-standard-deviation strangles at 45 days recorded success rates of 95% managing at 25% of maximum profit, 90% at 50% and 82% held to expiration, with managing at 50% delivering about 77% greater daily profit and loss at 23.5 average days in trade against 13.5.[120] Win rate rises monotonically as the target tightens while return does not, which is the whole of the argument: a 25% target converts a large number of positions into small certain gains and forfeits the tail of the win distribution that pays for the losses. On this study's condor that forfeit shows up as the shortfall of +1.83% at the 25% target against +2.11% at 50%, while at 75% the position is held long enough that the target stops binding before the loss distribution catches up.



Backtest: 40 large-cap names, 19 Aug 2021 to 18 Aug 2026, real daily price paths, modelled option prices, costs charged both ways. Author's calculations.

Figure 16.4 — The condor's return traces a shallow hump peaking at the 50% target while the calendars fall away monotonically, and no target level on its own approaches what the same structures earn once a time stop is attached.

16.3 How trades actually end

16.3.1 Under the managed rule the majority close on the time stop rather than the target, which is the opposite of how these rules are usually described

The managed rule is conventionally described as a 50% profit target with a stop as a backstop and a time exit as a formality. The decomposition inverts that ordering.

Exit channel	Rule in isolation	Iron condor	Calendar	Double calendar
Profit target only	pt50	+2.11	-1.88	-0.99
Stop only	sl200	-1.60	-2.36	-2.09
Time stop only	dte2	+4.85	+3.46	+2.70
All three (the managed rule)	pt50 sl200 dte2	+4.12	+3.57	+2.48

Average hold under the managed rule, days	—	4.8	5.2	5.3

Table 16.3 — Each clause of the managed rule run alone, in compound annual growth rate percent, against the combination, with the realised average hold beneath.

The time stop alone beats the full managed rule for the iron condor, +4.85% against +4.12%, and for the two calendars it sits alongside it — marginally below for the calendar at +3.46% against +3.57%, marginally above for the double calendar at +2.70% against +2.48%. The profit target alone is worth +2.11% on the condor and is negative on both calendars. The stop alone is negative everywhere. Whatever the managed rule is earning, it is earning it through the two-day close; the other two clauses are, on this sample, a small tax on the condor and approximately neutral on the calendars. The average hold corroborates the reading directly: against an entry window of 7 to 12 days, positions lasted 4.8, 5.2 and 5.3 days on average, so the typical trade was closed well before its front expiry rather than carried to it. The per-rule count of trades terminating on each channel was not reported in the study's computed results, so the claim in this heading rests on that decomposition, on the average hold and on Figure 16.3 rather than on a tabulated frequency.

Two qualifications limit how far this travels. Every option price behind these fifteen columns is a model price rather than an observed market quote, and the implied volatilities driving the exits are modelled; a rule that closes at two days to expiry is being priced by the same surface that priced the entry, and the surface's behaviour in the final days is the part of it least disciplined by data. And the entire grid is gross of the difference between a rule that trades often and one that does not — a two-day close on a 7 to 12 day entry is a shorter hold and a higher turnover than holding to expiry, and the round-trip cost of $18.88 on the condor already consumes 7.0% of a $247 median credit at the base slippage assumption. What that assumption is worth, and what happens to the ranking in Table 16.1 when it is doubled, is the subject of the next chapter.

The distribution of realised holding periods shows how hard each clause actually binds. Held to expiry, the median condor trade runs eight trading days and the mean 7.4. Adding a 50% profit target alone pulls the median to five days and the mean to 5.65. The full managed rule brings the mean to 4.83 against an entry window of 7 to 12 calendar days, and the tighter variant to 3.68. The position is therefore closed roughly halfway through its nominal life, which recasts what the mandate is: the specification names a 7-to-12-day option, but the strategy that works holds it for about a week of trading, and the last sessions before expiry, where the decay is fastest and the gamma worst, are mostly never experienced.

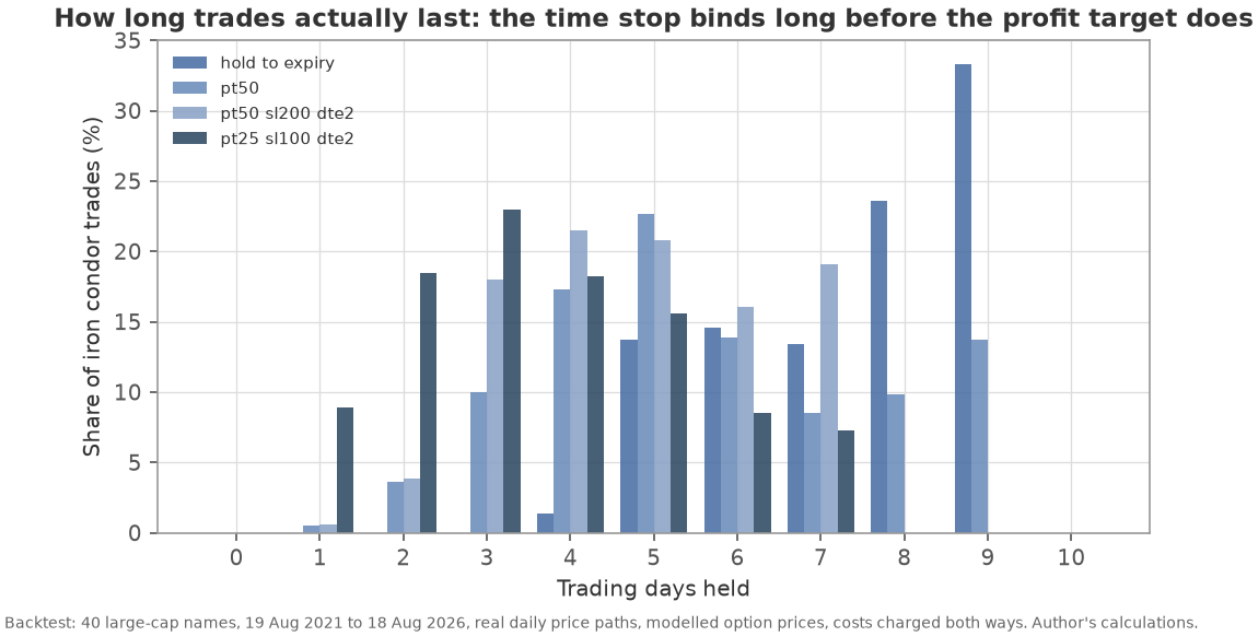


Figure 16.5 - Realised holding period for the iron condor under four rules. The mass moves left as clauses are added, and almost none of it reaches expiry.

---

17. Monte Carlo, Path Risk and Costs

17.1 Block bootstrap

17.1.1 Four thousand resampled paths put the condor's median terminal wealth at 1.099 with a 24.7% chance of a loss and a median worst drawdown of 11.1%

Chapter 14 quoted the terminal-wealth distribution beside the headline Sharpe ratios and left the construction to this chapter. The resampled unit is the per-trade return on capital at risk under the managed rule — a 50% profit target, a 200% stop and a close at two days to expiry, whichever comes first — taken from the realised trade record net of the base cost assumption. Nothing is modelled forward: no return enters a path that the backtest did not produce.

Paths are assembled by block bootstrap rather than by independent draws. The procedure samples blocks of ten consecutive trades from the realised sequence and concatenates them until the path reaches the realised trade count, 822 trades for the iron condor and 807 for each calendar. Four thousand paths are generated per strategy, each compounding at 1% of equity risked per trade, the sizing used throughout the backtest. The block preserves local clustering, which is the property an independent draw destroys. Entries are not spread evenly: a median of ten trades enter per month and as many as 39 positions would be open simultaneously against a funded cap of twelve, so consecutive trades share a market state and lose together. A block of ten is roughly a month of entries at the realised rate, and keeping those runs intact is what allows the drawdown column below to mean anything.

Strategy	5th percentile	Median	95th percentile	Probability of a loss	Median maximum drawdown
Iron condor	0.882	1.099	1.363	24.7%	-11.1%
Calendar	1.257	1.499	1.782	0.0%	-5.1%
Double calendar	1.102	1.277	1.470	0.3%	-4.9%

Table 17.1 — Terminal wealth across 4,000 block-bootstrapped paths per strategy, path length equal to the realised trade count, 1% of equity risked per trade.

The condor's median path ends at 1.099 and its fifth-percentile path at 0.882, with the ninety-fifth at 1.363; nine tenths of outcomes fall inside a band running from below starting equity to well above it. A quarter of paths — 24.7% — finish below where they started, on a strategy whose point estimate is a Sharpe ratio of 0.84 and a compound annual growth rate of 4.12%. That gap between a respectable ratio and a one-in-four chance of ending down is the reason the exercise is here: a single realised path reports one draw from Table 17.1, and the backtest's headline is that draw rather than the distribution behind it. The median worst drawdown of 11.1% is also the more honest planning figure than the -8.2% the realised condor path actually suffered, since the realised path was not the median one for drawdown.

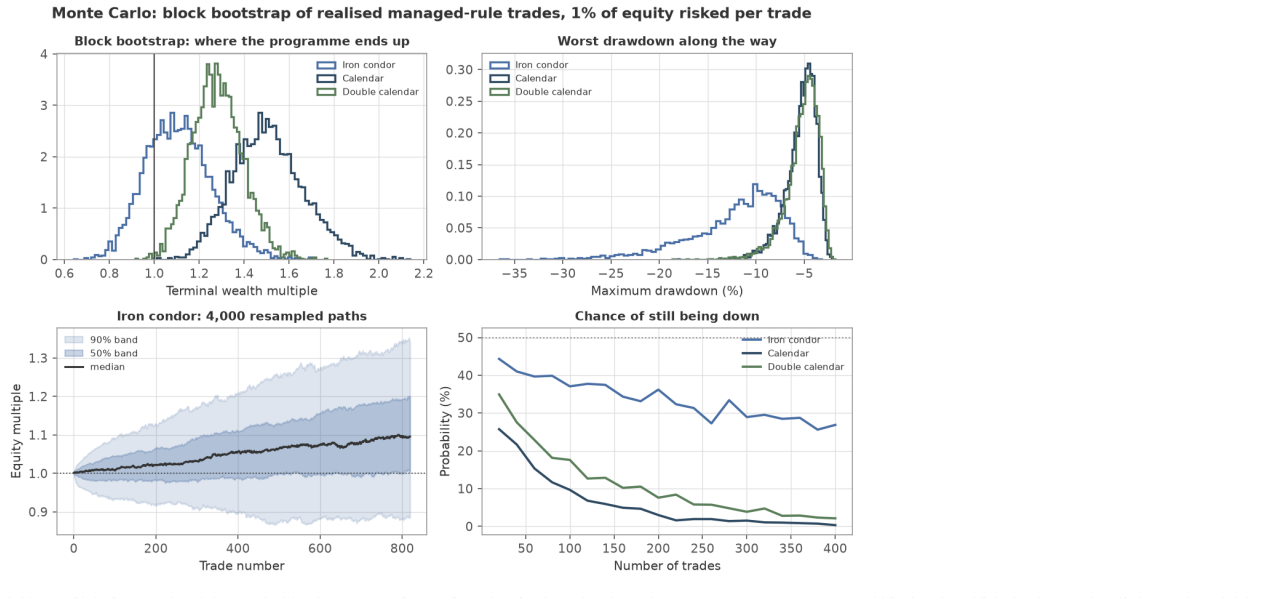


Figure 17.1 — The three terminal-wealth distributions from Table 17.1, showing the condor's wider spread and the mass of its left tail below 1.0.

17.1.2 The calendars show tighter distributions because their per-trade loss is capped by construction

Figure 17.1 separates the three distributions clearly. The single calendar spans 1.257 to 1.782 between the fifth and ninety-fifth percentiles with no losing path at all in 4,000 draws, and the double calendar spans 1.102 to 1.470 with 0.3% of paths losing. Both sit entirely to the right of the condor's fifth percentile, and both carry roughly half its median drawdown, -5.1% and -4.9% against -11.1%.

The conventional explanation is that a debit structure's loss is bounded by its premium, and the trade record only partly supports it. Median capital at risk equals the median debit for both calendars, $371 and $737, so the denominator of every per-trade return is the amount paid — but the bound is not clean. Chapter 9 measured a far downside of -$661 on a $623 double-calendar debit, about 6% beyond the premium, because the deep in-the-money put-side calendar tends to a negative value at a non-zero rate. Nor is the realised tail milder: the mean of the worst 5% of single-calendar trades is -208% of capital at risk against -191% for the condor, and the worst single trades are -468% and -469% respectively.

What actually separates the distributions is the sign of the skew. The condor's per-trade returns are skewed -0.74 with excess kurtosis +5.01; the calendars are skewed +0.35 and +0.58 with excess kurtosis +9.34 and +10.23. Both calendars are more fat-tailed than the condor, but their fat tail points the right way, and compounding at a fixed 1% of equity is far more forgiving of a positively skewed input than a negatively skewed one. A block bootstrap propagates that asymmetry directly: strings of the condor's small wins are interrupted by losses large enough to undo many of them, while the calendars' rarer large outcomes are mostly gains. The cap on the debit is a secondary effect at best; the skew is the mechanism.

17.2 What a bootstrap cannot do

17.2.1 A resample cannot manufacture an event absent from the sample, so a cluster of maximum losses is injected by hand

Table 17.1 is a lower bound on risk, not an estimate of it, and the reason is structural rather than a matter of sample size. Broadie, Chernov and Johannes state the limitation exactly: "the nonparametric bootstrap... has difficulties dealing with rare events. In fact, if an event has not occurred in the observed sample, it will never appear in the simulated finite sample distribution."[116] Every path in Figure 17.1 is a rearrangement of trades that happened. If the backtest window contains no genuine short-volatility catastrophe — no session in which every short strike across 40 correlated names is breached at once — then no resampling of it can produce one, however many paths are drawn. Four thousand is a statement about sampling error, not about the support of the distribution.

The same source shows how far this reaches into conventional inference. Broadie, Chernov and Johannes argue that "central limit approximations are problematic for option return statistics because of small samples sizes... and the irregular nature of option return distributions", and on replacing the asymptotic test with a finite-sample one, p-values rise "by roughly a factor of 10 for

the ATM options and by a factor of more than 10000 for deep OTM put options"; an in-sample average return of -60% on a 6% out-of-the-money put carries a p-value just over 8%.[116] They also report that excluding the single largest positive return cuts average returns at that strike by about 15% — one crash carries the result. A short-premium book sells precisely this payoff, so the standard t-test on its mean return is not merely imprecise here, it is the wrong test by orders of magnitude. Two further caveats compound it: every option price in the resampled pool comes from the study's modelled volatility surface rather than from historical quotes, and a model surface fitted to ordinary conditions is the least likely place for a genuine tail to originate.

The response is to stop asking the resampler for something it cannot supply and to impose the event instead. A cluster of maximum losses is injected by hand into the condor paths and the distribution re-run.

17.2.2 Three injected clusters move the condor's median terminal wealth from 1.042 to 0.821 and the fifth percentile from 0.920 to 0.726

Iron condor, 300-trade paths	Median terminal wealth	Fifth percentile
As observed	1.042	0.920
One injected cluster of eight maximum losses	0.959	—
Three injected clusters	0.821	0.726

Table 17.2 — Stress test on the iron condor over 300-trade paths, with clusters of eight maximum losses injected into the resampled sequence. The fifth percentile is reported for the observed and three-cluster cases only.

The baseline differs from Table 17.1 because the path length differs: 300 trades rather than the realised 822, which is why the unstressed median is 1.042 here and 1.099 there. Against that baseline, one cluster of eight maximum losses is enough to put the median path below its starting value, from 1.042 to 0.959. Three clusters take it to 0.821, a loss at the median rather than in the tail. The fifth percentile moves from 0.920 to 0.726 over the same three injections, so the adverse case deteriorates by more than the median does. Nothing about the strategy changed between those rows; only the assumption that the sample contained the worst thing that can happen. Eight consecutive maximum losses is not an extreme premise for a book that would hold up to 39 concurrent positions in 40 names whose mean pairwise correlation has ranged to 0.694 — one gap through every short strike on one morning delivers a cluster of exactly this shape, and the backtest window supplied none. Figure 17.2 shows the consequence as a shift in the whole distribution rather than a thickening of its left tail alone, which is the signature of a repeated event rather than a single outlier.

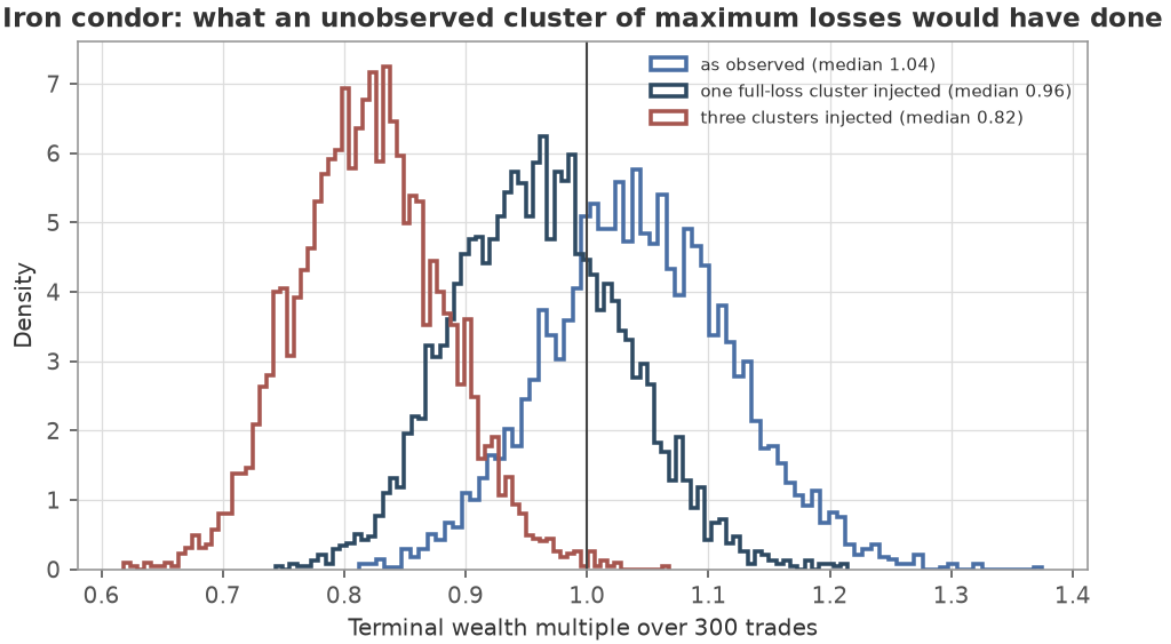


Figure 17.2 — The condor's stressed distributions, showing the whole distribution translating left as clusters are injected rather than only the tail thickening.

17.3 Costs decide the outcome

17.3.1 At zero slippage all three earn between 7.4% and 7.8%; at twice the base assumption the condor earns 0.85% and both calendars lose money

Slippage multiple	Iron condor	Calendar	Double calendar
0.0x	+7.42	+7.75	+7.44
0.5x	+5.82	+5.59	+4.90
1.0x (base, about $3.25 per contract per leg)	+4.12	+3.57	+2.48
1.5x	+2.54	+1.46	+0.20
2.0x	+0.85	-0.78	-2.23
3.0x	-2.42	-4.87	-6.58

Table 17.3 — Compound annual growth rate (%) under the managed rule against the assumed execution cost, all other parameters held at the mandated specification: an implied volatility rank above 78, 7 to 12 days to expiry, and short-leg deltas of 25 to 35 for the condor and 30 to 40 for

the calendars.

At zero cost the three structures are barely distinguishable: +7.42% for the condor, +7.75% for the calendar and +7.44% for the double calendar. Every difference between them reported anywhere in this study is therefore a difference in what they pay to trade, not in what they earn from the volatility surface. The ordering inverts along the ladder — the single calendar leads at 0.0x and the condor leads from 0.5x onward — which is the cross-sectional inversion Goyal and Saretto document when a portfolio is priced at the midpoint instead of at achievable fills.[107] Their magnitudes are of the same order as Table 17.3: a straddle portfolio returning 22.5% per month at the midpoint falls to 7.5% when the effective spread equals the quoted spread, and delta-hedged call returns fall from 2.1% to statistically insignificant returns of around 0.6%.[107] Round-trip costs measured in this backtest are $18.88 for the condor, $22.46 for the calendar and $43.95 for the double calendar — 7.0%, 4.4% and 4.4% of premium. The condor pays the largest share of its premium and survives the ladder best, because its capital at risk of $524 exceeds its $247 premium while each calendar's capital at risk equals its debit.

The base assumption is calibrated rather than chosen. NYSE's Q1 2020 options review, covering the top 15% most active symbols and 95% of multiply-listed volume, reports an average January 2020 quoted spread of $0.09, or 374 basis points of option price, with 752 contracts posted per side.[109] Muravyev and Pearson show that orders fill inside those quotes — effective half-spread 76% of quoted, an adjusted measure at 60%, and 42% for traders who time executions — so charging the full quoted spread is over-conservative and charging zero is indefensible.[108] Taking 65% of a $0.045 half-spread gives about $3.00 of spread cost on a 100-multiplier contract, and adding roughly $0.20 of institutional commission[111] and $0.04 of clearing and regulatory fees[112] gives approximately $3.25 per contract per leg. Two caveats travel with it: the anchor is a January 2020 measurement with no 2026-vintage replacement located, and Christoffersen, Goyenko, Jacobs and Karoui find short-dated contracts strictly less liquid than long-dated ones at every moneyness, with relative quoted spreads of 34.02% on 20-to-70-day out-of-the-money calls against 22.67% at 71 to 180 days.[113] This study trades the illiquid end of that comparison, at 7 to 12 days to expiry with an average hold of 4.8 to 5.3 days, on four legs for the condor and the double calendar. On a cycle measured in days rather than months, cost is not a rounding item in the result — it is most of the result.

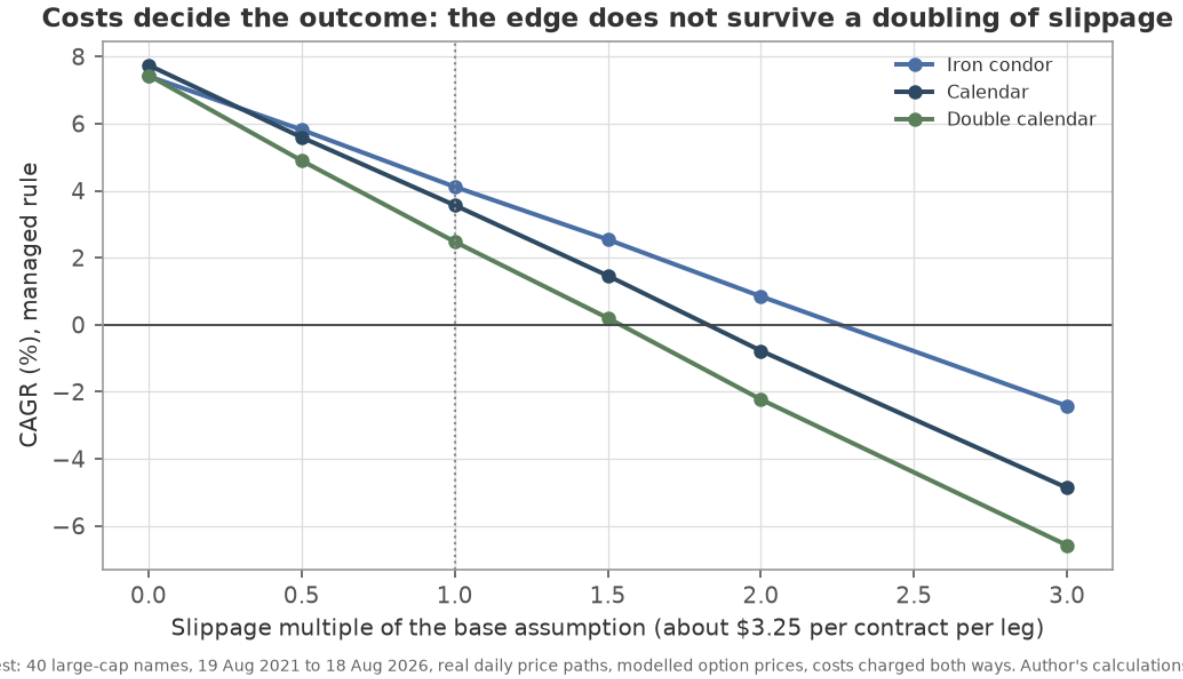


Figure 17.3 — The cost ladder of Table 17.3 drawn as three lines, the calendars crossing zero before the 2.0x column and the condor between 2.0x and 3.0x.

The right-hand end of Figure 17.3 is not a hypothetical. From January to March 2020 the average quoted spread rose from $0.09 to $0.32, a 256% increase in dollars and 123% in percentage terms, from 374 to 835 basis points, while posted size fell 53%, from 752 to 354 contracts per side.[109] Single names dispersed far beyond that market-wide average: on 28 February 2020 a vendor measurement put Citigroup's width at roughly 25 times its level a week earlier, with Exxon at 12 times and the major index products at 10 to 11 times.[132] A market-wide multiplier of 2.5 to 3.5 times therefore spans the 2.0x and 3.0x columns; at 3.0x the condor loses 2.42% a year, the calendar 4.87% and the double calendar 6.58% — and that regime arrives in exactly the week a short-premium book most needs to close. That evidence rests on a single episode, the only one for which options-specific measurements were retrieved, so the multiplier is indicative rather than a distribution.

The two halves of this chapter fail together and neither table shows it. Table 17.1 holds costs fixed at the base level in every path, and Table 17.3 re-runs the backtest along the realised price path with no injected cluster. The stress event and the cost multiplication are the same event, so the honest reading is that the 0.821 median of Table 17.2 and the negative rows of Table 17.3 are two views of one week, not two independent risks that can be planned for separately. Neither is a forecast; both are lower bounds on how a modelled edge of a few percent a year behaves when the assumptions holding it up move at once. That leaves one lever able to respond to all of it, which is how much is committed per trade and how many trades are permitted to be open together — the subject of chapter 19, where the growth-optimal fraction on the realised condor distribution is 9.1% of equity against the 1% used throughout this study.

18. Earnings: The Hidden Driver of High IV Rank

A scheduled report falls inside the next twelve days on **32.6%** of the days on which IV rank exceeds 78, against **15.2%** of all name-days in the backtest window — an over-representation of 2.1 times. The selection statistic established in chapter 6 fires on 8.85% of name-days, and on close to a third of those the trade will straddle an announcement. Screening on IV rank is therefore partly an earnings screen, and one the trader did not consciously choose. Nothing in the statistic makes that visible: IV rank reads position within a 52-week minimum-to-maximum range[65], and for a name whose largest implied-volatility prints of the year are its four scheduled reports, the top of that range is itself an earnings artefact.

18.1 The mechanism

18.1.1 Pre-event implied variance carries a jump term that inflates the front expiry and vanishes at the first post-announcement quote

Annualised implied variance ahead of a scheduled announcement separates into a diffusive component and an event component:

$$\sigma_t^2 = \sigma^2 + \frac{\left(\sigma_j^Q\right)^2}{T-t}$$

where σ^2 is diffusive variance, σ_j^Q the risk-neutral earnings jump volatility, and $T-t$ the time remaining to the announcement[60]. The event term is a fixed quantity of variance divided by the horizon over which it is annualised, and both effects that matter follow from that division.

The first is inversion. As $T-t$ shrinks the jump term grows at rate $(T-t)^{-1}$, so an expiry landing just after the print carries a far larger annualised implied volatility than one landing weeks later. The source expects the near-dated implied volatility to "blow-up" when maturity sits just after the announcement, so that "longer dated options will have lower Black-Scholes implied volatilities than shorter dated options"[60]. Chapter 5 set out what an inverted curve does to a position spanning two expiries; what matters here is that earnings are a recurring, calendar-known generator of it. The entry sample agrees: 43.3% of high-IV-rank days are inverted against 18.5% of all days.

The second is that the collapse is a discontinuity, not a decay. Once the announcement is public σ_j^Q has been realised and the event term is absent, so implied variance reverts to σ^2 at the first post-announcement quote[60]. No exit placed in the minutes after the release avoids it. The implied volatilities here are modelled rather than quoted, and the event is inserted explicitly as a variance ratio of 0.23 — an implied jump volatility of 1.108 times each name's realised announcement-day volatility.

18.1.2 Documented magnitudes: a worked large-cap case with 26 points of inversion and a 28-point overnight drop

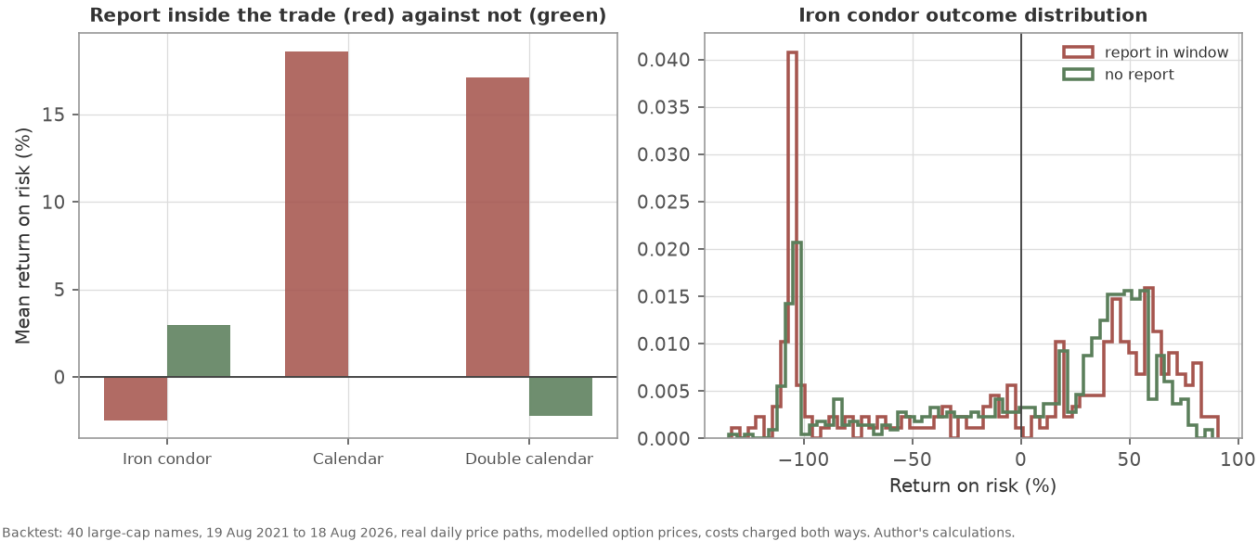
Intel, reporting after the close on 15 July 1997, is the worked large capitalisation case. Its July at-the-money call carried a Black-Scholes implied volatility of 71.15% against 45.19% for the August at-the-money call, a front-over-back inversion of 25.96 volatility points. The day after the announcement the July call "fell drastically to 42.96 percent" — a one-session drop of 28.19 points, 60% of the pre-event level — while the August expiry fell only to 43.92%, a move of 1.27 points[60]. That gap is the mechanism in one observation, and the authors state it "holds more generally for other earnings announcement dates and for other firms". The jump component itself is smaller than those headline points imply: two estimators put the Intel event at 8.6% and 7.2%, and across the sample the estimates "are on the order of 10 percent", with a 75th percentile near 15%[60].

18.2 What it did to each strategy

18.2.1 The condor lost on earnings trades and made money without them; both calendars did the reverse, and by a wide margin

Structure	Trades spanning a report	Mean return, report inside	Mean return, no report	Worst trade, report inside	Worst trade, no report
Iron condor	28.8%	-2.47%	+2.97%	-157%	-124%
Calendar	28.3%	+18.59%	-0.00%	-109%	-97%
Double calendar	28.3%	+17.13%	-2.22%	-107%	-100%

Returns in the table and in Figure 18.1 are means on capital at risk, and close to 28% of each book spanned a report. The condor's sign is negative with a report inside and positive without, so its positive mean came entirely from trades that avoided an announcement. Both calendars invert that: the calendar's non-earnings trades round to -0.00% and the double calendar's are negative at -2.22%, so the calendars' returns on this sample were earnings returns. The tails move the same way but far less sharply — the condor's worst trade goes from -124% to -157% of capital at risk, the calendars only from -97% to -109% and from -100% to -107%. A wide gap in the mean beside a narrow one in the tail points to an effect on the body of the distribution, not a handful of disasters.



Backtest: 40 large-cap names, 19 Aug 2021 to 18 Aug 2026, real daily price paths, modelled option prices, costs charged both ways. Author's calculations.

Figure 18.1 — The sign of the earnings effect reverses between the condor and the two calendars, large relative to each unconditional mean.

18.2.2 The same event is a liability for a short-vega structure and an asset for a long-forward-variance one

The iron condor is short both wings in a single expiry, and it is the front expiry — the one carrying the jump term. Selling into an inflated front looks like an advantage and is priced as one, the premium being larger precisely because $(\sigma_j^Q)^2/(T-t)$ sits inside it. What it then owns is the realised jump: reaction days here carry a mean absolute move of 5.48% against 1.38% on ordinary days, a ratio of 3.96, while the mandated specification places the short strikes at 25 to 35 delta, inside one standard deviation of the forward. The premium compensates for a risk the condor is fully exposed to, and the measurement says it was insufficient.

The calendars sell the front expiry and buy the back, so their result is the spread between two expiries rather than the level of either, which makes an inflated front leg the thing they are long. The jump term enters the front leg divided by a small $T-t$ and the back leg divided by a large one, so when it vanishes at the first post-announcement quote it vanishes almost entirely from the leg the structure is short — Intel's front expiry lost 28.19 volatility points overnight, the back 1.27[60]. The same event is therefore a cost to a structure short the front and a benefit to one that sells the front and buys the back.

The earnings split and the term-structure split are consequently one finding measured twice. Held to expiry, the calendar returns +10.7% on capital at risk when the curve is inverted at entry against -7.3% in steep contango, and the double calendar +6.8% against -4.5%. The two conditioning variables are not independent evidence.

Three qualifications belong with the result. The earnings jump modelled here is built from each name's own past squared earnings-day returns, so it reproduces the average size of a reaction but not the dispersion of surprises: the fat right tail of surprise magnitudes is absent by construction.

Second, the published direction on the pre-event window runs against short premium — at-the-money straddles held from three days before an announcement to the announcement date earn a "highly significant 3.34% return"[62], on the reading that "the average option trader underestimates future volatility before forthcoming earnings announcements"[72]. Third, no source located in this study's research documents finds unconditional premium selling into earnings profitable after costs in large capitalisation names: the strongest documented short-volatility premium around announcements is conditional on an ex-ante signal, concavity in the short-dated implied volatility curve[133], and its abstract gives no net-of-cost return.

18.3 The practical rule

18.3.1 Earnings exposure should be a deliberate choice per structure rather than a by-product of the IV rank screen

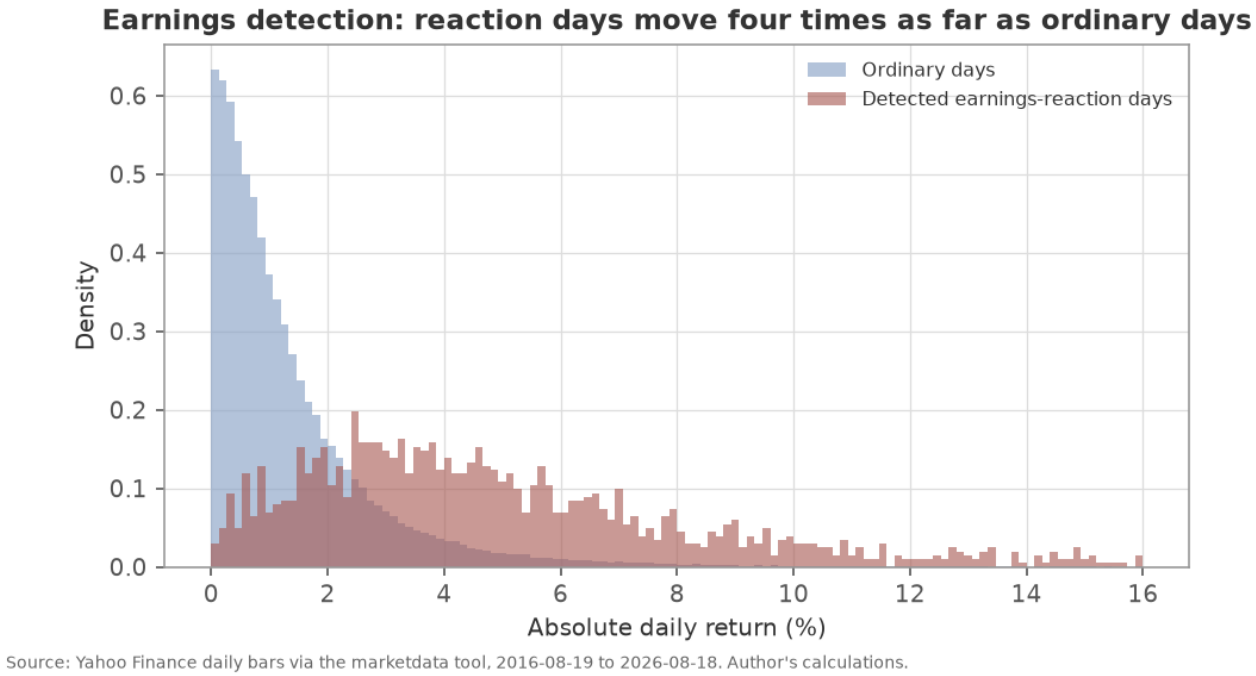
The rule is one line per structure, and it is a filter, not a forecast. For the iron condor, require that no scheduled report falls inside the window: the screen supplies one on 32.6% of qualifying days, and the measured mean return on capital at risk there is -2.47% against +2.97% without. For both calendars, treat a report inside the window as a condition for entry rather than a hazard, because on this sample that is where the return was. Chapter 21 carries the checklist.

Since the chapter rests on knowing which days are earnings days, detection was audited against Securities and Exchange Commission (SEC) filings.

Diagnostic	Value
Reaction days identified	1,577
Names covered	40
Mean absolute move, reaction days	5.48%
Mean absolute move, ordinary days	1.38%
Ratio	3.96
Median offset to nearest SEC filing date	1 day
Modelled implied jump volatility, x realised announcement-day	1.108
Earnings jump premium, variance ratio	0.23

Detection is behavioural rather than calendrical: days are flagged by the size of the move, then checked against the filing record. The separation is wide — 5.48% against 1.38%, a ratio of 3.96

— so these are a distinct population, not the upper tail of the ordinary one, and the median offset to the nearest SEC filing date is one day, which is what a correctly identified announcement should look like given that a release after the close moves the following session. Both are consistency checks only: no match against a commercial earnings calendar was performed, and names whose largest moves come from elsewhere will be misclassified. The last two rows are model inputs, recording how much extra variance the pricing model places in an expiry containing a flagged day. Figure 18.2 shows both checks.



Source: Yahoo Finance daily bars via the marketdata tool, 2016-08-19 to 2026-08-18. Author's calculations.

Figure 18.2 — Detected reaction days separate cleanly from ordinary days by move size and cluster one day after the corresponding filing.

The screen's earnings content is a timing problem as well as a selection one. Reports arrive in season, and the backtest records up to 39 positions open at once against a funded cap of twelve; how much of that clustering is calendar-driven was not measured, but a filter admitting earnings to the calendar book and excluding them from the condor book cannot spread the two evenly through the year. Chapter 19 takes the concurrency constraint on its own terms.

---

19. Portfolio Construction and Sizing

19.1 Sizing a negatively skewed book

19.1.1 The growth-optimal fraction on the realised condor distribution is 9.1% of equity per trade; this study uses 1%

Every equity curve in this report was produced under one portfolio rule, applied identically to all three structures: one percent of realised equity risked per trade, sized from the equity standing at

the entry date so the curve compounds without look-ahead. Capital at risk, the denominator of per-trade returns, has a median of $524, $371 and $737 by structure.

Rule	Setting	Consequence in the backtest
Capital risked per trade	1% of realised equity	Also the Monte Carlo fraction
Sizing basis	Equity at the entry date	Compounding without look-ahead
Risk denominator	Capital at risk per spread	Median $524, $371 and $737 by structure
Concurrent funded positions	12	Peak demand reached 39
Contracts per trade	40	Binds as equity grows
Signals above the cap	Skipped, not funded	822 and 807 trades are net of skipping

Table 19.1 — The portfolio rule behind every result here.

Two rows do more work than they appear to. Sizing from terminal rather than entry-date equity would fund early positions with capital the strategy had not yet earned, an error compounding across 822 trades. The skip rule matters more: a signal arriving when twelve positions are already open is discarded rather than queued or funded smaller, so the strategy declines its own signal exactly where the screen fires hardest. The headline figures — a 4.12% compound annual growth rate for the managed iron condor against an 8.2% maximum drawdown — are the returns of a book that turned trades away, not of one that took everything the mandated specification offered.

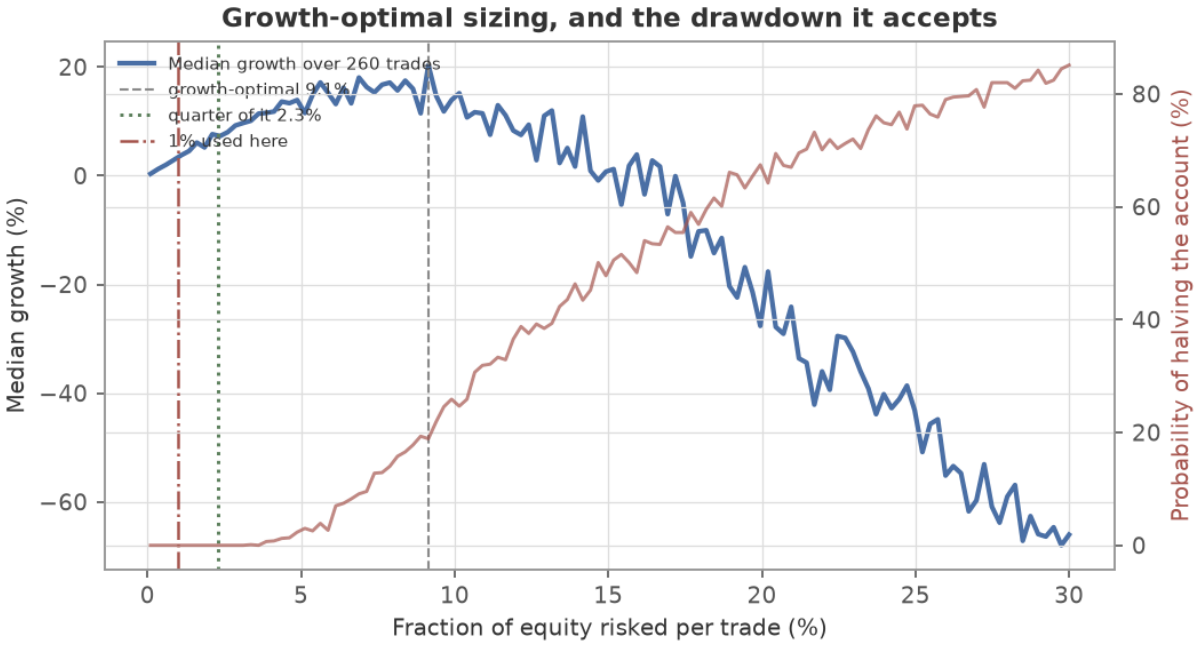


Figure 19.1 — Growth against the fraction of equity risked per trade on the realised iron condor distribution: the optimum sits at 9.1%, and the study funds 1%, roughly a ninth of it.

19.1.2 Why full Kelly is the wrong target when the win rate and the tail are estimated from a window that may contain no tail event

The Kelly fraction maximises long-run compound growth, and its principal authors are direct about the cost. MacLean, Thorp and Ziemba note that full Kelly "maximizes the limiting exponential growth rate of wealth," but that its "suggested wagers may be very large," so that "the Kelly criterion is relatively risky in the short term... it is possible, through bad scenarios, to lose most of one's wealth"[134]. Fractional Kelly is thus a risk-preference choice, not the repair of an error.

Here the estimation problem is worse than the risk-preference one. The 9.1% figure derives from 822 managed iron condor trades with a 67.2% win rate, a skew of -0.74 and excess kurtosis of +5.01, over a five-year window. Both inputs that drive it — the win probability and the size of the tail — are estimates drawn from a sample that may contain no tail event; Goetzmann, Ingersoll, Spiegel and Welch call this "peso problem" one that "may be ubiquitous" where high Sharpe ratios are rewarded[122]. The bias runs one way: a window that missed its tail returns a fraction that is too large, precisely when funding it would be most damaging. Chapter 14's point that Sharpe understates this payoff's risk applies directly, the ratio and the fraction being functions of the same unobserved tail. Sinclair and Brooks add that "the higher order moments of option positions have an economically significant effect on optimal trade sizing"[135] — a correction two-moment Kelly ignores, and one that only cuts the fraction further[136].

Structure	5th pct	Median	95th pct	Probability of a loss	Median max drawdown	Realised max drawdown
Iron condor	0.882	1.099	1.363	24.7%	-11.1%	-8.2%
Calendar	1.257	1.499	1.782	0.0%	-5.1%	-5.6%
Double calendar	1.102	1.277	1.470	0.3%	-4.9%	-7.4%

Table 19.2 — The sizing ladder at the funded fraction: terminal wealth from a block bootstrap of 4,000 paths at 1% of equity per trade, beside the managed rule's realised maximum drawdown.

At a ninth of the growth-optimal fraction the iron condor still loses money on 24.7% of bootstrap paths, its fifth-percentile path ending at 0.882 of starting equity. That is the payoff shape, not the sizing. The calendars sit differently, at 0.0% and 0.3% probability of a loss. The instructive disagreement is between the last two columns: for the iron condor the simulated -11.1% is worse than the -8.2% realised, so history was kinder than a typical resample of its own trades, while for the double calendar the realised -7.4% is worse than the simulated -4.9%. Neither settles the question alone, and at 9.1% per trade both would be multiples of what is shown.

19.2 Concurrency and correlated exposure

19.2.1 A median of ten trades enter per month and signals cluster: up to 39 would be open at once against a funded cap of twelve

A median of ten trades enter per month, and that demand is not spread evenly: at peak, 39 positions would have been open simultaneously against the funded cap of twelve, the excess skipped. An IV rank above 78 fires on 8.85% of name-days, and elevated implied volatility arrives across names together. Mean pairwise correlation across the forty names stood at 0.094 at the 18 August 2026 close against a ten-year mean of 0.317, but the ten-year range runs to 0.694 — and the top of that range is where a volatility-keyed screen fires everywhere at once.

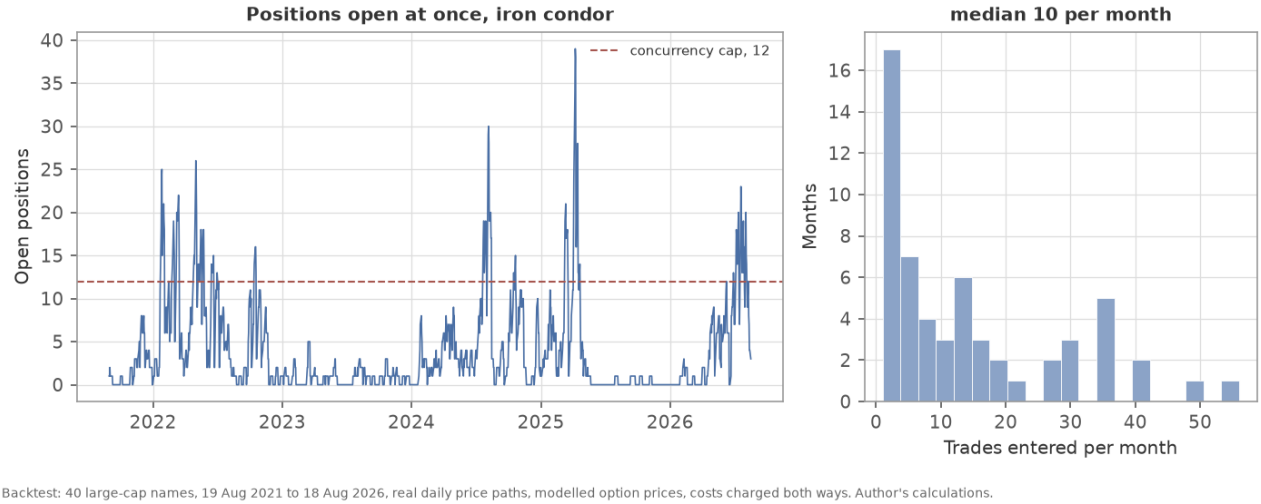


Figure 19.2 — Concurrent open positions through the backtest against the funded cap of twelve; peak demand reaches 39.

The Monte Carlo behind Table 19.2 preserves this: resampling in blocks of ten consecutive trades keeps bad clusters intact instead of averaging them away. Over 300 iron condor trades, median terminal wealth is 1.042 as observed, 0.959 with one injected cluster of eight maximum losses and 0.821 with three, while the fifth percentile falls from 0.920 to 0.726.

19.2.2 Clustering is the binding constraint, not per-trade risk, and it is what destroyed the named option-selling failures

Programme	Date	Loss	Proximate cause
LJM Preservation & Growth Fund	2–5 February 2018	More than $1 billion, about 80% of value, in two days[126]	S&P 500 down more than 6%, VIX up 177%; broker-ordered liquidation to meet margin[126]
OptionSellers.com	November 2018	~300 accounts liquidated; $31.3m receivable as of 10 December 2018[137]	Naked selling; all accounts answered to one natural-gas move, clients liable for all losses[138]
Allianz Structured Alpha	February–March 2020	More than $7 billion, from above $11 billion of assets[139]	Promised crash hedges "repeatedly failed to buy", making a defined-risk book undefined[139]

Malachite Capital	March 2020	~$600 million fund dissolved[140]	Margin calls from counterparties including BNP Paribas[140]

Table 19.3 — Four named option-selling failures. None was caused by a bad signal.

Every row failed on leverage and correlated exposure rather than on a wrong forecast. LJM's signal had run since 2013; what changed was size, and the SEC alleges the manager "increased the level of risk in the portfolios in order to chase return targets" during "a period of historically low volatility"[141]. LJM's own scenario reporting flagged the exposure months ahead: for a -5% S&P 500 and +50% implied-volatility shock, estimated loss averaged 17.6% from January 2016 to September 2017, then rose "from 18.6% at the beginning of October to 61.4% by January 31st" — while the disclosed 20% "worst case" had been produced by "doubling the P&G Fund's worst single-day performance and rounding up"[126]. A single joint-shock number, trending, said what was coming; the win rate said nothing. Allianz Global Investors US pleaded guilty, stipulating $3,238,748,199.69 in restitution[142].

Inferred from these records rather than stated in them, the mechanism is that nominally separate positions were one bet: LJM's whole book answered to a single -5%/+50%-volatility vector[126], and roughly 300 OptionSellers accounts to one natural-gas move[137]. What they support is a portfolio joint-stress budget — size so the aggregate loss under a specified simultaneous spot-and-volatility shock, and the margin call it triggers, is survivable. The twelve-position cap is a crude version, on count rather than shocked loss. Two of the four were institutional multi-strategy books, not short-dated premium selling; the mechanism transfers, the parameters do not.

19.3 Assignment and pin risk

19.3.1 Legs do not net: a short leg assigned without the long produces a loss that can exceed the spread width

Clearing provides no offset between the legs of a spread. An assigned writer must perform "regardless of the overall risk of their position when taking into account other options that may be owned as part of the overall multi-leg strategy"[143]. FINRA's worked example: short 50-strike puts against long 55-strike puts, the stock closes at $56 with both legs out of the money, then news drops it to $40 after the bell. The short puts are assigned; the long puts, out of the money at the official close, are not. The writer buys stock at $50 against $40 — a loss exceeding the width the structure was designed to cap, unhedgeable until the next session, and over an expiration weekend roughly 65 hours of single-name gap exposure plus a settlement obligation of strike times 100 per contract. On a calendar, assignment of the near-dated short leaves stock to be paid for until the longer-dated leg is exercised[143].

Exercise by exception is a one-cent threshold measured against the official OCC closing price, and is not strictly automatic: a clearing member may decline to exercise past it, or exercise short of it[144]. Early exercise of a call is a dividend-capture event, rational when the dividend exceeds remaining time value and typically taken the day before the ex-dividend date; the writer learns of it only the next business day, and exercising the long leg to flatten the short stock does not avoid the dividend[98]. Deep in-the-money puts are exercised early for interest on the strike proceeds[145]. Under pin risk the writer "will not know how many of their short options will be assigned"[111], yet the hedge must be chosen before the close, without the assignment count.

The frequently quoted statistic that about 6% to 7% of contracts are exercised must not be read as an assignment probability: the denominator is all contracts, of which more than 72% are closed before expiry and 22% expire worthless[144]. OIC disclaims the inference directly — "That does not mean that you can only be assigned on 7% of your short option"[146]. The usable claim is conditional: a writer short an option expiring in the money "should expect assignment... though assignment is not guaranteed"[145]. This study's defence is the managed rule's close at two days to expiry; the held-to-expiry variants that carry the exposure in full are also the ones that lose money, at -0.76%, -2.36% and -2.09% compound annual growth. No assignment outcome appears among the computed results, and the option prices behind them are modelled, not observed quotes.

Figure 19.1 fixes what a single trade may cost; Figure 19.2 shows how many such trades the signal wanted open at once, which is what turns a per-trade fraction into a portfolio exposure. Neither can show whether the five-year window that produced the 9.1% optimum contained its own tail — the question Chapter 20 takes up, and on which the rule above rests.

---

20. Limitations and What Would Falsify This

20.1 The modelled surface

20.1.1 No historical option quotes were used; every option price in the backtest is a model price, and the report should be read as conditional on that model

Not one option price in this report was observed. Every entry credit, exit debit, intermediate mark and Greek in chapters 7 to 19 came out of the pricing pipeline; the underlying paths, the Cboe indices, rates and filing dates are real, the options are not. There is therefore no way, from inside this study, to separate the proposition that these structures earned what is reported from the proposition that the surface was built in a way that pays them.

Bias	Direction on reported returns	Controlled here	How, and what is left

Modelled rather than quoted prices	Ambiguous; overstating near events	Partly	Earnings jump priced explicitly; nothing controls unmodelled demand
Midpoint fills	Overstates	Yes	65% of the quoted half-spread charged both ways; slippage swept to 3.0x
Effective/quoted spread uncertainty	Overstates	Partly	Published ratio contested between under 0.5 and 0.8–1.0; the sweep spans it
Liquidity interaction inverting the cross-section	Overstates, and distorts ranking	Not applicable	Universe is 40 large caps, the liquid side of the sort
Measurement error and conditioning on endpoints	Overstates	Not applicable	No noisy quotes to condition on; model error replaces it
Stale and non-synchronous quotes	Overstates	Not applicable	Modelled prices are synchronous with the close by construction
Look-ahead through filters and full-sample statistics	Overstates	Yes	Volatility model fitted 2016–2021 and frozen; IV rank uses a trailing window
Selection bias from the specification search	Overstates	Yes, and quantified	Deflated Sharpe probability, section 20.2.1
Sample shorter than the search justifies	Overstates	Partly	38 configurations against guidance of about 45; 570 cells is far beyond it

Rare events absent from a five-year sample	Overstates	No	Block bootstrap cannot repair it; injected-cluster stress is the only answer
Survivorship in the underlying universe	Overstates	No	Point-in-time index membership was not used
Concurrency beyond the funded book	Overstates	No	Up to 39 positions open at once against a cap of 12

Eleven of the twelve rows point the same way and the twelfth is ambiguous rather than favourable, which is the honest summary of this study's error structure: none of it is symmetric noise a larger sample would average away. The twelve split evenly four ways — three controlled, three partly, three that do not arise, three not controlled at all. Two of the three that do not arise are excused only because a modelled surface cannot suffer a defect requiring noisy quotes to condition on,[105] which substitutes one error for another rather than removing it. Survivorship needs a further caveat: no primary study quantifying it for a US large-cap *options* backtest was located, so the point-in-time remedy is asserted in the literature without a citation this report can stand behind.

The strongest defence of the surface is the out-of-sample reconstruction of the VIX from S&P 500 prices alone over 2021 to 2026: correlation 0.818, R-squared +0.642, root mean squared error 3.15 points, bias +0.12 points, 1,254 observations. That validates one thing — the at-the-money level of a broad index — and 3.15 volatility points of error is of the same order as the gross index premium the condor is designed to harvest. The tenor biases are not neutral either: VIX9D reconstructs with a bias of +0.60 points and VIX3M with -1.70. If that pattern carries to the single-stock surface, an inference rather than a measurement, the model runs rich at the front and cheap at the back, which flatters a seller of front premium and a buyer of back premium at once — the two sides of every trade here. One reproducibility snag belongs with it: the calibrated premium is kappa = 0.12, but `volmodel_meta.json` carries a stale 0.5 from an earlier pass, and a replicator trusting the metadata rebuilds a different surface.

20.1.2 What the model cannot produce: idiosyncratic volatility demand, supply shocks, and the true shape of the far wings

The criticism aimed at this class of proxy is that a realised-volatility forecast scaled by a variance risk premium is defensible for a broad index at horizons of a week or more, and not for short-dated single-stock options near earnings, where true implied volatility carries an event term scaling as one over time to the announcement and the market charges for a jump the proxy prices at the diffusive level.[101][60] The partial answer here is that the jump is modelled explicitly from each

name's announcement history and priced with a separate premium, at a variance ratio of 0.23 and an implied jump volatility of 1.108 times realised announcement-day volatility, validated against 1,577 reaction days that move 5.48% on average against 1.38% on ordinary days. That covers the scheduled, quantifiable part of event risk and leaves four things uncovered.

First, idiosyncratic volatility demand. Implied volatility rises when someone wants optionality, and the causes leave no trace in a price range: takeover speculation, index-inclusion flow ahead of a rebalance, an activist position being assembled. High modelled IV rank here can only ever be a high diffusive forecast, a scheduled report inside the window, or a high VIX. Real high-IV-rank episodes are a strictly larger set, and the surplus members are disproportionately those where the market's price is high for a reason a backward-looking model cannot see.

Second, supply shocks in the option market itself. Prices move when supply moves — dealer inventory turning over, a large systematic overwriter withdrawing, a structured-product hedge recycled into listed strikes. The pipeline has no supply side: implied volatility here is a forecast times a constant, so it cannot go bid because nobody will sell.

Third, the true shape of the far wings. The smile is one quadratic slope scaled off the index SKEW reading, currently 143.6, not a per-name measured skew, so cross-sectional variation in single-stock skew is absent by construction. A scalar premium ratio reproduces the at-the-money level but cannot generate skew and is biased whenever risk-neutral skewness moves.[106] The parametrisation is not globally admissible either: an unbounded quadratic smile prices a deep in-the-money option above its own intrinsic value in the far wings, which is why the structure diagrams cap standardised moneyness at plus or minus 2.5. That cap touches only the diagrams, never the backtest, where every strike sits inside one standard deviation of the forward — but the reason it is needed is a real defect in the wing.

Fourth, and most consequential, the model reproduces the average size of an earnings surprise and not its dispersion. The jump premium is a single ratio on a per-name mean absolute reaction, and compressing the distribution around that mean removes precisely the observations that decide whether trading through a report pays. The calendar's entire measured edge sits there: mean return on risk of +18.59% on trades spanning a scheduled report against -0.00% on those that do not, and +17.13% against -2.22% for the double calendar. The most load-bearing result in this study is the one most exposed to the invented part of the surface, and exposed in the favourable direction. A standing disagreement sharpens it further: later work finds that implied volatility subsumes the information in past volatility rather than merely echoing it,[147] and on that view the proxy is impoverished, not just noisy.

20.2 Statistical honesty

20.2.1 Thirty-eight configurations and 570 configuration-rule cells were evaluated; the deflated Sharpe ratio adjusts the best result for that search

Thirty-eight configurations were run, each against fifteen management rules, for 570 configuration-rule cells. Published guidance holds that with five years of data no more than about forty-five independent configurations should be tried before a strategy with no skill is expected to show an in-sample Sharpe of 1 and an out-of-sample Sharpe of zero, with a rule of thumb in years of

$$\mathrm{MinBTL} < \frac{2 \ln N}{E[\max SR]^2}$$

described by its authors as necessary and not sufficient.[114] Thirty-eight sits inside forty-five; 570 does not, by any reading. The deflated Sharpe ratio corrects an observed Sharpe for selection under multiple testing and for non-normal returns, using the trial count, the dispersion of Sharpe ratios across trials, the sample length, and the skew and kurtosis of the selected strategy.[115]

Input	Iron condor	Calendar	Double calendar
Best rule found	`dte1`	`pt25_sl100_dte2`	`dte3`
Sharpe of that rule	+1.19	+0.71	+0.52
CAGR of that rule	+7.45%	+4.04%	+1.98%
Sharpe under the managed rule	+0.84	+0.61	+0.52
Trades in the sample	822	807	807
Skew of trade returns	-0.74	+0.35	+0.58
Excess kurtosis of trade returns	+5.01	+9.34	+10.23
Configurations tried	38	38	38
Configuration-rule cells tried	570	570	570
Deflated Sharpe probability	**0.334**	**0.071**	**0.025**

The convention is easy to read backwards, so it needs stating plainly. The deflated Sharpe probability is the probability that the observed Sharpe exceeds what the best of that many random trials would have produced. Higher is better. A value below 0.95 means the result is not distinguishable, at conventional confidence, from what the search itself would generate on data with no edge in it. All three fall well below 0.95, and not marginally: 0.334, 0.071 and 0.025. None

of the three best-performing rules in this report survives its own search. The double calendar scores worst at 0.025 even though its winning rule improved nothing — Sharpe of +0.52 for `dte3` against +0.52 under the mandated management, on a lower CAGR of +1.98% against +2.48% — which is the point of the statistic: the winner is judged against the dispersion of the field it beat and the shape of its own tail, and excess kurtosis of +10.23 is punished. For scale, the framework's worked example takes a Sharpe of 2.5 over 1,250 observations found after 1,000 trials, with negative skew and excess kurtosis, and returns only a 90% probability that the true Sharpe exceeds zero; the same result from 46 trials would have scored 0.9505 and passed.[115] The Sharpe ratios here are near 1, not 2.5.

The competing treatment converts the Sharpe to a t-ratio, adjusts the p-value for prior tests and converts back. Its authors call the usual flat 50% haircut a serious mistake, because the haircut is nonlinear: almost always more than 50%, sometimes much more, when the annualised Sharpe is below 0.4, but at most 25% above 1.0.[148] Only the condor's best rule, at +1.19, sits in the lenient band; both calendars fall short of it, into a range the source does not quantify. The two estimators are complements and must not be averaged.

20.2.2 Rare events break the standard t-test for exactly this payoff, and a non-parametric bootstrap cannot fix it

Central limit approximations are unreliable for option return statistics because the samples are small and the distributions irregular. A parametric bootstrap simulating underlying paths raises p-values by roughly a factor of ten for at-the-money options and by more than ten thousand for deep out-of-the-money puts, to the point where an in-sample average return of -60% on a 6% out-of-the-money put carries a p-value just over 8%, and dropping the single largest positive return cuts that strike's average return by about 15%.[116] The payoff that description fits is the one transacted here: 822 condor trades with skew of -0.74, excess kurtosis of +5.01, a worst single trade of -469% of capital at risk and a mean of the worst 5% at -191%.

The Monte Carlo in chapter 17 is a block bootstrap over blocks of ten consecutive trades across 4,000 paths, and it does preserve serial dependence. It cannot repair this. The same source is explicit that a non-parametric bootstrap has difficulty with rare events, and that an event which has not occurred in the observed sample will never appear in the simulated distribution.[116] The proof that the objection bites is in this study's own output: the block bootstrap returns a probability of loss of 0.0% for the calendar and 0.3% for the double calendar. Those are properties of the resample, not of the strategies, because a five-year window containing no short-volatility catastrophe cannot produce one by rearrangement. The injected-cluster stress test is the only analysis that addresses this, by adding events the data does not contain: on the condor over 300 trades, median terminal wealth falls from 1.042 as observed to 0.959 with one cluster of eight maximum losses and 0.821 with three, and the fifth percentile from 0.920 to 0.726. That is a deliberate assumption, and the right kind — but no result in this report has been tested against a volatility crisis.

20.3 What would change the conclusion

20.3.1 A single-stock premium at the index level, materially tighter execution, or a longer sample containing a genuine short-volatility catastrophe

Three tests would move these conclusions, each specified tightly enough to run.

The first is the premium. Re-measure the implied-to-expected-realised ratio for the same 40 names against actual quoted surfaces over the same window. This study calibrates to a median modelled ratio of 1.060 against 1.364 for the index, with a published range for liquid large caps of 1.03 to 1.07. A higher measured ratio does not simply improve the book: at an assumed 1.128 the condor rises to +5.58% while the calendar falls to +2.50% and the double calendar to +1.82%, and at 1.262 the condor reaches +11.85% while the double calendar drops to +0.79%. A richer premium confirms the condor and refutes the calendars; a ratio below 1.03 removes the condor's case entirely. One measurement on purchasable data decides more than any other number here.

The second is execution. Reprice the trade log against a live quote feed for a quarter and compare realised fills to the 65% of quoted half-spread charged here. The published effective-to-quoted ratio is itself contested, between under 0.5 and 0.8 to 1.0 for large stocks, and a short-dated programme pays the spread every cycle so the haircut scales with turnover — though no primary measurement at 7-to-12-day tenors was located.[107] At 1.5 times the base slippage the three strategies return +2.54%, +1.46% and +0.20%; at 2.0 times, +0.85%, -0.78% and -2.23%. At roughly twice the charge assumed here, both calendars are loss-making and the condor is indistinguishable from cash.

The third is the sample. Extend the backtest through a period containing a genuine short-volatility catastrophe, using the same frozen 2016-2021 model and holding selection at the rule already chosen, so the extension is fully out of sample. Two cheap checks belong beside it: test that rule on data not used for selection, the only procedure that raises a deflated Sharpe probability rather than explaining it away; and impose the funded concurrency cap of twelve against the 39 positions the signal generated at peak, since a book that cannot take every trade does not earn the backtest's return.

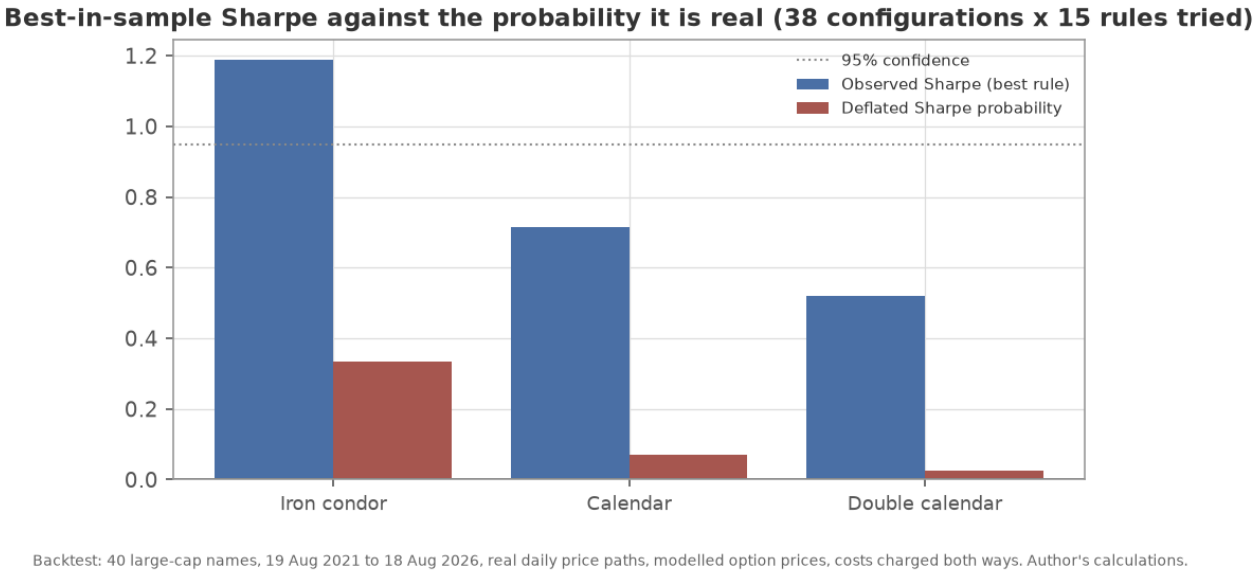


Figure 20.1 — All three deflated Sharpe probabilities sit far below the 0.95 line, the double calendar furthest at 0.025, so the best result found for each structure lies inside the range the search alone could have produced.

Figure 20.1 condemns quoting the best cell of a 570-cell grid as a finding. It does not establish that these strategies have no edge: a probability below 0.95 is a failure to demonstrate, not a demonstration of absence. That distinction governs how the next chapter reads. Chapter 21 sets out what a practitioner would do with this evidence, and every element of it inherits the three unresolved dependencies above — the premium ratio, the fill assumption and the absent crisis — as conditions to act under rather than settled questions.

---

21. Conclusion and Implementation Playbook

21.1 What a practitioner should take from this

21.1.1 The mandated specification is workable but thin, and its edge is smaller than its win rate suggests

The managed iron condor won 67.2% of its 822 trades and compounded at +4.12% a year. Those two figures sit at very different orders of impressiveness, and the distance between them is the most useful single thing this report measured. Held to expiry the same entries won 61.4% of the time and compounded at -0.76%: a win rate comfortably above three in five, attached to a strategy that lost money. Win rate on this payoff describes the frequency of the small outcomes, not the sign of the sum.

The size of what is being won is the reason. The condor's median credit was $247 against median capital at risk of $524, its distribution is left-skewed at -0.74 with excess kurtosis of +5.01, and the mean of its worst 5% of trades was -191% of median capital at risk. The round-trip cost of $18.88

consumes 7.0% of that median credit before the trade has done anything: gross of slippage the managed condor returns +7.42% a year, and at about $3.25 per contract per leg, +4.12%.

The bootstrap says the same in the language of outcomes. Resampling in blocks of ten consecutive trades at 1% of equity per trade, 24.7% of 4,000 iron condor paths ended below their starting equity, with a median terminal wealth of 1.099 and a fifth percentile of 0.882. Workable, then, in that three specifications compound positively under a rule fixed in advance; thin, in that a quarter of the condor's resampled histories are losing ones.

21.1.2 Three changes with the largest measured effect: a time stop, wider condor wings, and a longer expiry than 7 to 12 days

The time stop is the most robust finding in the study and the least like a discovery. Closing at two days to expiry, applied alone, took the iron condor from -0.76% to +4.85%, the calendar from -2.36% to +3.46% and the double calendar from -2.09% to +2.70%. Every rule in the fifteen-rule grid containing a time stop is positive for all three structures; every rule built only from a stop expressed as a multiple of premium is negative for all three, at or below the hold-to-expiry baseline. The clause is not a refinement of the managed rule — it is the managed rule, and the 50% target and 200% stop subtract from it on the condor, taking +4.85% down to +4.12%.

Widening the condor's wings from 3% to 8% of spot raised the compound rate from +1.67% to +4.92% and cut the maximum drawdown from -11.1% to -5.7%. It is the only parameter swept anywhere in this study that improved return and risk together, monotonically across all three settings, which makes it a shape rather than a point estimate. Length of expiry moves more: at 25 to 35 days the condor returns +7.45% and the double calendar +7.71%, against +4.12% and +2.48% in the mandated window, while the single calendar peaks earlier at +3.96% over 13 to 20 days. Below the mandate the surface collapses, to -8.80%, -5.99% and -7.40% at 3 to 6 days.

The three do not carry equal weight. The time stop holds across fifteen rules and three structures; the wing result is monotone in two dimensions at once; the expiry result rests on four windows, contradicts the short-dated premise the mandate was built on, and points at exactly the region — long tenor, late close — from which the study's post hoc maxima were drawn.

21.2 The screening rule that follows from the evidence

21.2.1 Condors want high IV rank, no earnings inside the window and a curve in contango; calendars want the opposite on two of the three

Read across the chapters, the report's three conditioning variables are not three findings. An implied volatility rank above 78 fires on 8.85% of name-days; on 32.6% of those a scheduled report falls inside the next twelve days against 15.2% of all days, and 43.3% of high-IV-rank days are inverted against 18.5% of all days. The screen selects events, and the events invert the curve. The condor is short the front expiry carrying the jump term and owns the realised move; the calendars sell that front and buy a back leg the collapse barely touches. One exposure, opposite

signs, entered on the same names on the same days.

Rule	Iron condor	Calendar	Double calendar
IV rank at entry	Above 78: +4.12%, best of five, but unordered (+3.22% at 50, +1.82% at 65, +1.26% at 90, -1.84% unscreened)	Above 65: +5.37% (+4.55% at 50, +3.57% at 78, -0.39% at 90)	Above 50: +8.44% (+4.76% at 65, +2.48% at 78, -0.26% at 90)
Term structure at entry	Not measured. Curve state was conditioned on for the calendars only; contango is inferred from their result and the earnings channel	Inverted: +10.7% mean return on capital at risk against -7.3% in steep contango above 1.10, unmanaged	Inverted: +6.8% against -4.5% in steep contango above 1.10, unmanaged
Scheduled report inside the window	Avoid: -2.47% with a report against +2.97% without; worst trade -157% against -124%	Accept, and prefer: +18.59% against -0.00%	Accept, and prefer: +17.13% against -2.22%
Short-leg delta	30: +4.12%, but nearly flat (+2.87% at 25, +3.91% at 35)	30: +7.71%, monotone (+3.57% at 35, +3.02% at 40)	30: +5.96%, monotone (+2.48% at 35, +2.38% at 40)
Expiry window	25 to 35 days: +7.45% (+4.12% mandated, -8.80% at 3 to 6)	13 to 20 days: +3.96% (+3.57% mandated, +3.40% at 25 to 35, -5.99% at 3 to 6)	25 to 35 days: +7.71% (+2.48% mandated, -7.40% at 3 to 6)

Wing or strike width	Wings at 8% of spot: +4.92% at -5.7% drawdown (5%: +4.12% at -8.2%; 3%: +1.67% at -11.1%)	Not applicable; a single strike	Not measured as a return. Geometry only: base strikes 5.8% of spot apart give peaks of +$311 and +$434 either side of a +$305 trough, the middle falling to +$180 and +$72 at 1.6 and 2.3 times that width
Management	Two-day close: +4.85% alone against +4.12% managed. Later is better, +7.45% at one day	Two-day close: +3.46% alone against +3.57% managed; one day falls to +1.50%	Two-day close: +2.70% alone against +2.48% managed; one day falls to +1.18%

Table 21.1 — Per-structure entry and construction rules, each cell traceable to a measured result on this sample, with unmeasured quantities marked as such. Growth rates are on the managed rule; the term-structure and earnings figures are unmanaged mean returns on capital at risk, on a different denominator from the headline tail statistics.

Three cells matter more than their row length suggests. The term-structure row is empty for the condor because the study never conditioned condor returns on curve state; the heading's claim that condors want contango is an inference from the calendars' measured preference and from the earnings split, and it should be held at that weight. The delta row disagrees with the mandate for both debit structures, which were assigned a 30 to 40 band and paid for the wrong half of it. And every column's management cell says the same thing: the time stop alone is within a few tenths of a percentage point of the full managed rule, above it for two structures of the three. A practitioner adopting one line from this table should adopt that one, because it is the only entry that survives being tested against fifteen alternatives rather than four or five.

The screening rule is necessary and not sufficient: none of the four option-selling failures Chapter 19 examined was caused by a bad signal.

Hazard	What the report established	Check before the trade

Assignment and non-netting	Clearing gives no offset between legs: a short assigned without the long produces a loss exceeding the width the spread was built to cap, unhedgeable until the next session[143]	Is the short strike inside a plausible overnight gap? The only defence measured here is the two-day close; the held-to-expiry variants carrying the exposure in full return -0.76%, -2.36% and -2.09%
Pin risk near the short strike	The writer cannot know how many shorts will be assigned, yet must choose the hedge before the close[111]. The 6% to 7% exercise statistic is not an assignment probability — over 72% of contracts are closed before expiry[144]	Flatten rather than carry a position pinned at a short strike; expect assignment on anything expiring in the money, though it is not guaranteed[146]
Dividend-driven early exercise	A call is exercised for a dividend exceeding remaining time value, typically the day before the ex-dividend date; the writer learns the next business day, and exercising the long leg does not avoid the dividend[98]	No ex-dividend date inside the window on the call side of a condor or a call calendar
Concurrency and correlated exposure	A median of 10 trades enter per month and up to 39 would be open at once against a funded cap of 12; the headline returns are net of skipped signals	Count open positions against the cap, and size against a simultaneous spot-and-volatility shock rather than per-trade risk. Correlation reached 0.694 within the last ten years
Cost sensitivity	Round trips of $18.88, $22.46 and $43.95 — 7.0%, 4.4% and 4.4% of median premium. At twice the base slippage: +0.85%, -0.78% and -2.23%; at three times: -2.42%, -4.87% and -6.58%	Measure realised round-trip cost against $3.25 per contract per leg before sizing. The condor is the only structure still positive at 2.0x

Table 21.2 — Pre-trade checklist for the operational hazards this report established, none of which appears in any return figure above.

The last two rows are the ones that scale with success. Assignment, pin and dividend exposure are per-trade hazards with per-trade consequences; concurrency and cost are portfolio properties that worsen as the book grows into its signal. The concurrency row carries a further implication the study did not measure directly: because the condor and calendar books are opposite sides of one earnings-driven front-expiry exposure, selected by the same screen, running both is closer to netting than to diversifying, and the residual exposure is smaller than the gross position count implies. No combined-book correlation was computed, so that is an inference from the measured signs rather than a result. What is measured is that no assignment outcome appears anywhere in the computed results, and every option price behind them is a model price.

21.3 Conditions today

21.3.1 With the VIX 1.27 points below fair value, correlation at 0.094 and the curve in contango at 1.094, the current environment favours the calendar side of the book over the condor side

At the 18 August 2026 close the VIX stood at 15.83, the 40th percentile of the last ten years, against an economic fair value of 17.10 from a model explaining 79.8% of in-sample variation in log VIX — a gap of -1.27 points, a residual of -0.52 standard deviations. Mean pairwise correlation across the forty names was 0.094 against a ten-year mean of 0.317 and a range whose floor is 0.090. The VIX3M/VIX ratio was 1.094, contango, the ordinary state on 92.3% of the last ten years of days.

Two of those three point away from the short-premium side. Implied volatility below what macro conditions support means the condor is being paid less for the same risk, and the condor is the structure most sensitive to how richly premium is priced: the same 822 entries return +4.12% at the calibrated assumption and +11.85% at the richest tested. Correlation at 0.094 against a ten-year floor of 0.090 is a position rather than an estimate, and the distance back to the 0.317 mean is travelled by index volatility rising — the state in which a volatility-keyed screen fires across names at once, bringing the concurrency hazard precisely when the short-premium book is largest.

The third reading cuts the other way and should not be smoothed over. Contango is the calendars' adverse state: +10.7% and +6.8% on inverted entries against -7.3% and -4.5% in steep contango. But 1.094 sits below the 1.10 cut defining that bucket, and the split conditions on each name's modelled 30-day over 7-day ratio at entry, not on the index curve. Single-name inversion remains available whatever the index does — 43.3% of high-IV-rank days were inverted — so the reading constrains the tilt without removing it.

The tilt the evidence supports is therefore toward the calendar side of the book — the structures paid for owning the front-to-back spread and for the events that invert it — and away from the structure whose return is a function of premium being rich at a moment when the fair-value model says it is not. That is a conditional reading of one day's state against relationships estimated on a

five-year sample, not a forecast. The fair-value gap is well inside the model's own root mean squared error of 3.53 points and carries little alone; the correlation reading is the more informative of the three; and none of the three says anything about the direction of the underlying market, which is not what any of these structures is positioned on.

That is as far as the evidence reaches, and the honest summary of its reach is unflattering. The mandated specification works, in that all three structures compound positively under a rule fixed in advance. Its edge does not survive a doubling of transaction costs: at 2.0x slippage the condor returns +0.85% and both calendars are negative. Its best variants cannot be distinguished from the product of the search that found them, at deflated Sharpe probabilities of 0.334, 0.071 and 0.025 against a threshold of 0.95, higher being better. And the largest single determinant of every return in this report is not a finding at all but a scalar chosen by the modeller — the assumed ratio of single-stock implied to expected realised volatility, calibrated here to give a median of 1.060 against 1.364 for the index, and worth more to the condor's headline than every entry parameter in Table 21.1 combined. A reader who accepts that number accepts the study; a reader who does not has no result to evaluate.

References

1. Reuters. S&P 500 notches record-high close as rate-hike worries ease. 13 August 2026. https://www.reuters.com/business/retail-consumer/wall-st-futures-tick-higher-oil-retreats-ahead-inflation-data-2026-08-13/
2. FactSet. S&P 500 Earnings Season Update. 7 August 2026. https://insight.factset.com/sp-500-earnings-season-update-august-7-2026
3. FactSet. Amazon Drives S&P 500 Earnings Growth for Q2 Above 45% on Valuation Gains. 3 August 2026. https://insight.factset.com/amazon-drives-sp-500-earnings-growth-for-q2-above-45-on-valuation-gains
4. CNBC (Michael Santoli). Earnings bonanza that lifted market to record may not be all that it appears to be. 17 August 2026. https://www.cnbc.com/2026/08/17/santoli-earnings-bonanza-that-lifted-market-to-record-may-not-be-all-that-it-appears-to-be.
5. DollarLiquidity. Forward P/E Ratios. 18 August 2026. https://dollarliquidity.com/en/valuation
6. International Monetary Fund. Global Financial Stability Report. 14 April 2026. https://www.imf.org/en/publications/gfsr/issues/2026/04/14/global-financial-stability-report-april-2026
7. VanEck. Understanding Concentration Risk in the S&P 500. 3 August 2026. https://www.vaneck.com/us/en/blogs/emerging-markets-equity/understanding-concentration-risk-in-the-sp-500/
8. Columbia Threadneedle. The rise of the Magnificent 7: concentration risk versus earnings power. 15 September 2025. https://www.columbiathreadneedle.com/en/insights/the-rise-of-the-magnificent-7-concentration-risk-versus-earnings-power/
9. Reuters. Magnificent 7 results set to test broadening US stock market. 29 July 2026. https://www.reuters.com/business/magnificent-7-results-set-test-broadening-us-stock-market-2026-07-29/
10. OptionMetrics. The Worst Crisis? 20 May 2025. https://optionmetrics.com/blog/the-worst-crisis/
11. SIFMA Insights. April 2025 Market Metrics. April 2025. https://www.sifma.org/wp-content/uploads/2025/04/SIFMA-Insights-Monthly-Metrics_April_2025.pdf
12. Bank for International Settlements. Market whiplash after the 2025 tariff shock. BIS Working Paper 1282. 2025. https://www.bis.org/publ/work1282.pdf

13. Associated Press. A brutal tech-led sell-off slams Wall Street in its worst day since October. 5 June 2026. https://apnews.com/article/stocks-markets-iran-oil-trump-b5e10863b81cb1d6399f688ad8885c46

14. Cboe Global Markets. Index Insights: March 2026. 1 April 2026. https://www.cboe.com/insights/posts/index-insights-march-2026

15. Cboe Global Markets. 'Spot Down, Vol Down' as Investors Monetized Hedges. 16 March 2026. https://www.cboe.com/insights/posts/spot-down-vol-down-as-investors-monetized-hedges/

16. Cboe Global Markets. Index Insights: April 2026. April 2026. https://www.cboe.com/insights/posts/index-insights-april-2026/

17. CNN Business. Nasdaq, S&P 500 suffer worst day of year as AI stocks tumble. 5 June 2026. https://www.cnn.com/2026/06/05/markets/stock-market-sell-off-fed

18. Reuters. Chip slump erases $1.3 trillion in stock market value. 5 June 2026. https://www.reuters.com/business/media-telecom/chip-selloff-erases-over-1-trillion-stock-market-value-2026-06-05/

19. Allianz Global Investors. When calm no longer means safe: navigating a new volatility regime. 6 May 2026. https://www.allianzgi.com/en/insights/navigating-a-new-volatility-regime

20. IncomeShares. The Rise of Options Income ETPs. February 2026. https://incomeshares.com/en-eu/rise-of-options-income-etps-europe-report

21. Cboe Global Markets. Evaluating the Market Impact of SPX 0DTE Options. 8 September 2023. https://www.cboe.com/insights/posts/volatility-insights-evaluating-the-market-impact-of-spx-0-dte-options/

22. Cboe Global Markets. Inside Volatility Trading: Is VIX Backwardation Necessarily a Sign of a Future Down Market. 26 July 2022. https://www.cboe.com/insights/posts/inside-volatility-trading-is-vix-backwardation-necessarily-a-sign-of-a-future-down-market

23. OptionMetrics. VIX Futures Term Structure: A Warning Sign? 20 March 2025. https://optionmetrics.com/blog/vix-futures-term-structure-a-warning-sign/

24. Brenner, Shu and Zhang. The New Market for Volatility Trading. NYU Stern. https://pages.stern.nyu.edu/~mbrenner/research/vix_futures.pdf

25. Cboe Global Markets. VIX Term Structure. https://www.cboe.com/tradable-products/vix/term-structure/

26. Cboe Global Markets. Cboe Volatility Index Mathematics Methodology. https://cdn.cboe.com/api/global/us_indices/governance/Cboe_Volatility_Index_Mathematics_Methodology.pdf

27. Johnson, Travis L. Risk Premia and the VIX Term Structure. Journal of Financial and Quantitative Analysis 52(6), 2017, 2461–2490. https://www.travislakejohnson.com/pdfs/Johnson%20VIXTS%202017%20(JFQA).pdf

28. Simon, David P. and Campasano, Jim. The VIX Futures Basis: Evidence and Trading Strategies. Journal of Derivatives 21(3), Spring 2014, 54–69. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2094510

29. CNBC. Wall Street's 'fear gauge' is doing something unusual as stocks hit record highs. 5 August 2026. https://www.cnbc.com/2026/08/05/wall-streets-fear-gauge-is-doing-something-unusual-as-stocks-hit-record-highs.html

30. Benzinga. S&P 500 Call Option Volume Hits Record 4 Million Contracts. 7 August 2026. https://www.benzinga.com/markets/options/26/08/61029397/sp-500-call-option-volume-hits-record-4-million-contracts-as-bull

31. Bloomberg. FOMO-Driven Buying of S&P Calls Fuels 'One-Way Flow' of Orders. 6 August 2026. https://finance.yahoo.com/markets/options/articles/fomo-driven-buying-p-calls-121443645.html

32. Federal Reserve. Federal Reserve issues FOMC statement. 29 July 2026. https://www.federalreserve.gov/newsevents/pressreleases/monetary20260729a.htm

33. CNBC. Divided Fed holds interest rates steady. 29 July 2026. https://www.cnbc.com/2026/07/29/fed-rate-decision-july-2026.html

34. Reuters. Dollar treads water, Fed hike bets pared on benign US inflation. 13 August 2026. https://www.reuters.com/world/asia-pacific/dollar-treads-water-fed-hike-bets-pared-benign-us-inflation-2026-08-13/

35. CNBC. VIX: Wall Street's 'fear gauge' hits 2026 low. 17 August 2026. https://www.cnbc.com/2026/08/17/stock-market-volatility-vix-wall-street.html

36. Federal Reserve. FOMC meeting calendars. Updated 29 July 2026. https://www.federalreserve.gov/monetarypolicy/fomccalendars.htm

37. Cboe. White Paper Shows Volatility Risk Premium Facilitated Higher Risk-Adjusted Returns for PUT Index (Bondarenko). 29 May 2019. https://www.cboe.com/insights/posts/white-paper-shows-volatility-risk-premium-facilitated-higher-risk-adjusted-returns-for-put

38. Horstmeyer, Handley and Manuel. How Well Does the Market Predict Volatility? CFA Institute. 31 July 2024. https://rpc.cfainstitute.org/blogs/enterprising-investor/2024/how-well-does-the-market-predict-volatility

39. Maidel. Strategy Spotlight: Considerations in Volatility Trading. Cboe. Undated. https://cdn.cboe.com/resources/education/research_publications/strategy-spotlight-considerations-in-volatility-trading.pdf

40. Bollerslev, Tauchen and Zhou. Expected Stock Returns and Variance Risk Premia. Review of Financial Studies 22(11). November 2009. https://public.econ.duke.edu/~boller/Published_Papers/rfs_09.pdf

41. Bonaparte, Chatrath and Christie-David. S&P volatility, VIX, and asymptotic volatility estimates. Finance Research Letters 51. 2023. https://ideas.repec.org/a/eee/finlet/v51y2023ics1544612322005694.html

42. Almeida et al. 0DTE Asset Pricing. FMA Derivatives. 2025. Conference working paper, not peer-reviewed. https://www.fma.org/assets/docs/Derivatives2025/Almeida.pdf

43. Le Grand et al. The Variance Risk Premium in Crude Oil Futures. Undated working paper; its S&P 500 crisis figures are background rather than the paper's audited result. https://francois-le-grand.com/docs/research/LS_vrp.pdf

44. DeSimone. Extraordinary Times for the Variance Risk Premium. OptionMetrics. 9 June 2020. https://optionmetrics.com/blog/20200608_extraordinary_times/

45. Bekaert and Hoerova. The VIX, the Variance Premium and Stock Market Volatility. Netspar working paper / Journal of Econometrics 183(2). 2014. https://www.netspar.nl/wp-content/uploads/035_Bekaert.pdf

46. Driessen, Maenhout and Vilkov. The Price of Correlation Risk: Evidence from Equity Options. Journal of Finance 64(3), 1377–1406. June 2009. https://onlinelibrary.wiley.com/doi/10.1111/j.1540-6261.2009.01467.x

47. Driessen, Maenhout and Vilkov. Option-Implied Correlations and the Price of Correlation Risk. Working paper, September 2005 version. https://haas.berkeley.edu/wp-content/uploads/dmv_sept2005.pdf

48. Carr and Wu. Variance Risk Premiums. Review of Financial Studies 22(3). March 2009. https://engineering.nyu.edu/sites/default/files/2019-01/CarrReviewofFinStudiesMarch2009-a.pdf

49. Cao and Han. Cross section of option returns and idiosyncratic stock volatility. Journal of Financial Economics 108(1), 231–249. 2013. https://www-2.rotman.utoronto.ca/facbios/file/Han_JFE_published.pdf

50. González-Urteaga and Rubio. The cross-sectional variation of volatility risk premia. Journal of Financial Economics. 2016. https://www.sciencedirect.com/science/article/abs/pii/S0304405X15001762

51. Driessen, Maenhout and Vilkov. Option-Implied Correlations and the Price of Correlation Risk. Working paper, March 2005 version. https://pages.stern.nyu.edu/~rengle/Maenhout-S2005.pdf

52. Duarte, Jones and Wang. Very Noisy Option Prices and Inference Regarding the Volatility Risk Premium. Working paper, 2022; unpublished. https://faculty.marshall.usc.edu/Christopher-Jones/pdf/duarte_jones_wang_2022.pdf

53. Lin, Chang and Paxson. Smiling less at LIFFE. Journal of Futures Markets 28(1):57-81. 2008. https://ideas.repec.org/a/wly/jfutmk/v28y2008i1p57-81.html

54. Dennis and Mayhew. Risk-Neutral Skewness: Evidence from Stock Options. Journal of Financial and Quantitative Analysis. 2002. https://patrickjohndennis.com/2002-risk.neutral.skewness.pdf

55. Bakshi, Kapadia and Madan. Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options. Review of Financial Studies 16(1):101-143. 2003. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=239528

56. Branger and Schlag. Why is the Index Smile So Steep? Review of Finance 8(1):109-127. 2004. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=316140

57. Cboe. SKEW Index White Paper. January 2011. https://cdn.cboe.com/resources/indices/documents/SKEWwhitepaperjan2011.pdf

58. Bergomi. Smile Dynamics IV. Fields Institute. 2009-10. https://www.fields.utoronto.ca/programs/scientific/09-10/finance/derivatives/bergomi.pdf

59. Delemotte, De Marco and Segonne. Yet another analysis of the SP500 at-the-money skew. 2023. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4428407

60. Dubinsky and Johannes. Option Pricing of Earnings Announcement Risks. Review of Financial Studies 32(2):646-687. 2019. https://business.columbia.edu/sites/default/files-efs/pubfiles/6051/DJ_2006.pdf

61. Goldman Sachs. How Earnings, Volatility, and AI Capex Are Affecting US Markets. 17 July 2026. https://www.goldmansachs.com/insights/the-markets/how-earnings-volatility-and-ai-capex-are-affecting-us-markets

62. Gao, Xing and Zhang. Anticipating Uncertainty: Straddles around Earnings Announcements. Journal of Financial and Quantitative Analysis 53(6):2587-2617. December 2018. https://www.cambridge.org/core/journals/journal-of-financial-and-quantitative-analysis/article/abs/anticipating-uncertainty-stra

63. Implied Volatility (IV) Rank & Percentile Explained. tastylive. Undated. https://www.tastylive.com/concepts-strategies/implied-volatility-rank-percentile

64. IV Rank vs IV Percentile: A Complete Guide. Barchart. Undated. https://www.barchart.com/education/iv_rank_vs_iv_percentile

65. How to Trade Implied Volatility Rank (IVR). tastylive. 28 June 2024. https://www.tastylive.com/news-insights/how-to-trade-implied-volatility-rank-ivr-

66. IVR and IVP: The Numbers. tastylive Market Measures. 30 August 2016. https://www.tastylive.com/shows/market-measures/episodes/ivr-and-ivp-the-numbers-08-30-2016

67. IV Rank vs SKEW Rank. tastylive Market Measures. 31 July 2018. https://www.tastylive.com/shows/market-measures/episodes/iv-rank-vs-skew-rank-07-31-2018

68. Malkiel, Rinaudo and Saha. Option Writing: Using VIX to Improve Returns. SSRN. 2018. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3228589

69. Moreira and Muir. Volatility-Managed Portfolios. Journal of Finance. 6 April 2017. https://doi.org/10.1111/jofi.12513

70. Van Tassel. Equity Volatility Term Premia. Federal Reserve Bank of New York Staff Report 867. Undated. https://www.newyorkfed.org/medialibrary/media/research/staff_reports/sr867.pdf?sc_lang=en

71. Sizing the Risk: Kelly, VIX, and Hybrid Approaches in Put-Writing. Preprint. 2025. https://arxiv.org/html/2508.16598v1

72. Chung and Louis. Earnings announcements and option returns. Journal of Empirical Finance 40:220-235. 2017. https://ideas.repec.org/a/eee/empfin/v40y2017icp220-235.html

73. Haugh. The Black-Scholes Model (Columbia lecture notes). n.d. https://www.columbia.edu/~mh2078/FoundationsFE/BlackScholes.pdf

74. The Black-Scholes Greeks. Stochiki. n.d. https://www.stochiki.com/wiki/guide:63b5b5a4fa

75. Cboe. Volatility Management with Cboe's BFLY and CNDR Indices. 2021-07-26. https://www.cboe.com/insights/posts/benchmark-indices-series-volatility-management-with-cboes-bfly-and-cndr-indices/

76. Black and Szado for Cboe. Performance of Cboe SPX Options-Selling Indexes. 2016-02-23. https://s202.q4cdn.com/174824971/files/doc_news/2016/02/study-analyzes-performance-of-cboe-spx-options-selling-indexes

77. Cboe Insights. New Research Shows Options-Based Strategies Can Generate Higher Gross Premiums with Less Volatility. n.d. https://www.cboe.com/insights/posts/new-research-shows-options-based-strategies-can-generate-higher-gross-premiums-wi

78. Bondarenko. An Analysis of Index Option Writing with Monthly and Weekly Rollover. 2016 (data to 2015). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2750188

79. The Timing of Option Returns. n.d. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2909163

80. 0DTE Trading Rules. Data 09/2016-01/2024. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4641356

81. Bondarenko. Historical Performance of Put-Writing Strategies. 2019 (data to 2018). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3393940

82. Cboe. The State of the Options Industry: 2025. 2026-01-22. https://www.cboe.com/insights/posts/the-state-of-the-options-industry-2025/

83. CME Group. Option Calendar Spreads (lesson). Undated. https://www.cmegroup.com/education/lessons/option-calendar-spreads

84. Brenner, M. and Subrahmanyam, M. A Simple Approach to Option Valuation and Hedging in the Black-Scholes Model. Financial Analysts Journal 50(2). March 1994. https://rpc.cfainstitute.org/research/financial-analysts-journal/1994/a-simple-approach-to-option-valuation-and-hedging-in-the-

85. Gundersen, G. Estimating ATM Option Prices. 19 August 2023. https://gregorygundersen.com/blog/2023/08/19/atm-approximation/

86. Options Industry Council. Long Call Calendar Spread (Call Horizontal). Undated. https://www.optionseducation.org/strategies/all-strategies/long-call-calendar-spread-call-horizontal

87. Cboe Global Markets. Advanced Options Trading Strategies. 2025. https://cdn.cboe.com/resources/options/Advanced_Opts_Trading_Strategies.pdf

88. Options Analysis Suite. Calendar Spread Options Strategy: Term-Structure Trade. 1 January 2026. https://www.optionsanalysissuite.com/strategies/calendar-spread

89. Dobrowolski, T. How to Read Volatility Term Structure. FlashAlpha. 17 March 2026. https://flashalpha.com/articles/volatility-term-structure-contango-backwardation-events

90. CME Group. Option Calendar Spreads (course). Undated. https://www.cmegroup.com/education/courses/option-strategies/option-calendar-spreads

91. Calendar and Double Calendar Spreads - A Detailed Guide. m.Stock (Mirae Asset). Undated. https://www.mstock.com/mlearn/stock-market-courses/option-strategies/calendar-and-double-calendar-spreads

92. Gavin McMaster. Double Calendar Spreads: Ultimate Guide. Options Trading IQ. 12 June 2020. https://optionstradingiq.com/double-calendar-spreads/

93. Iron Condor. The Options Playbook. Undated. https://www.optionsplaybook.com/option-strategies/iron-condor/

94. Iron Condor Strike Selection. Options Trading IQ. 17 August 2026. https://optionstradingiq.com/iron-condor-strike-selection/

95. Double Calendar Spread: Overview, Example. Strike. 19 June 2025. https://www.strike.money/options/double-calendar-spread

96. Dan Passarelli. Double Calendars vs Double Diagonals. Market Taker Mentoring. 28 February 2013. https://markettaker.com/2013/02/double-calendars-vs-double-diagonals/

97. Double Diagonal Spread. The Options Playbook (Brian Overby). Undated. https://www.optionsplaybook.com/option-strategies/double-diagonal-spread

98. What Is Double Diagonal Spread? Fidelity. Undated. https://www.fidelity.com/learning-center/investment-products/options/options-strategy-guide/double-diagonal-spread

99. Calendar Spread Vs Iron Condor. Options Trading IQ. 22 June 2026. https://optionstradingiq.com/calendar-spread-vs-iron-condor/

100. Yang and Zhang. Drift-Independent Volatility Estimation Based on High, Low, Open, and Close Prices. Journal of Business 73(3):477-491. 2000. https://www.atmif.com/papers/range.pdf

101. Corsi. A Simple Approximate Long-Memory Model of Realized Volatility. Journal of Financial Econometrics 7(2). 2009. https://statmath.wu.ac.at/~hauser/LVs/FinEtricsQF/References/Corsi2009JFinEtrics_LMmodelRealizedVola.pdf

102. Craioveanu and Hillebrand. Why It Is Ok To Use The HAR-RV(1,5,21) Model. University of Minnesota Working Paper 1201. https://ideas.repec.org/p/umn/wpaper/1201.html

103. Blair, Poon and Taylor. Forecasting S&P 100 volatility: the incremental information content of implied volatilities and high-frequency index returns. Journal of Econometrics. 2001. https://www.sciencedirect.com/science/article/abs/pii/S0304407601000689

104. The Importance of the Volatility Risk Premium for Volatility Forecasting. Working paper, September 2013, University of Liverpool repository. https://livrepository.liverpool.ac.uk/3052241/1/Final_Paper_Sep_2013.pdf

105. Duarte, Jones and Wang. Very Noisy Option Prices and Inference Regarding the Volatility Risk Premium. January 2022. http://faculty.marshall.usc.edu/Christopher-Jones/pdf/duarte_jones_wang_2022.pdf

106. Rompolis and Tzavalis. Risk Premium Effects On Implied Volatility Regressions. Journal of Financial Research 33(2):125-151. 2010. https://ideas.repec.org/a/bla/jfnres/v33y2010i2p125-151.html

107. Goyal and Saretto. Cross-Section of Option Returns and Volatility. Working paper dated 18 April 2008; published Journal of Financial Economics 2009. http://www.ruf.rice.edu/~jgsfss/goyal_041808.pdf

108. Muravyev and Pearson. Review of Financial Studies 33(11):4973-5014, with Internet Appendix Table IA.2. 2020; appendix dated 18 January 2020. https://www.dmurav.com/MuravyevPearson.OptionsTradingCosts.InternetAppendix.pdf

109. NYSE Data Insights. Q1 2020 Options Review. 8 April 2020. https://www.nyse.com/data-insights/q1-2020-options-review

110. SEC SR-CBOE-2025-075 Exhibit 5, Cboe Rule 5.4 minimum quoting increments. 2026. https://www.sec.gov/files/rules/sro/cboe/2026/34-105131-ex5.pdf

111. Interactive Brokers. Options commissions. Accessed August 2026. https://www.interactivebrokers.com/en/pricing/commissions-options.php

112. OCC Schedule of Fees. As of January 2026. https://www.theocc.com/company-information/schedule-of-fees

113. Christoffersen, Goyenko, Jacobs and Karoui. Review of Financial Studies 31(3):811-851. 2018. https://academic.oup.com/rfs/article/31/3/811/4371415

114. Bailey, Borwein, Lopez de Prado and Zhu. Pseudo-Mathematics and Financial Charlatanism. Notices of the AMS. May 2014. https://www.ams.org/notices/201405/rnoti-p458.pdf

115. Bailey and Lopez de Prado. The Deflated Sharpe Ratio. Journal of Portfolio Management 40(5):94-107. 31 July 2014. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2460551

116. Broadie, Chernov and Johannes. Understanding Index Option Returns. Working paper. https://business.columbia.edu/sites/default/files-efs/pubfiles/2612/broadie.pdf

117. Pawe■ Perz. Profitability of selected 0DTE index options strategies. IBIMA Publishing. 2024. https://ibimapublishing.com/p-articles/44ACC/2024/4452224/4452224-1.pdf

118. Jesse Blom (Lorintine Capital) / SteadyOptions. Does "Managing Winners" Add Value to Short Strangles? 27 June 2022.

https://steadyoptions.com/articles/does-%E2%80%9Cmanaging-winners%E2%80%9D-add-value-to-short-strangles-r618/

119. projectfinance. Short Strangle Management. 1 February 2022. https://www.projectfinance.com/short-strangle-management/

120. tastylive. Why Manage at 50%, Not 25%, for Short Strangles. 23 September 2016. https://www.tastylive.com/news-insights/why-manage-at-50-not-25-for-short-strangles

121. tastylive. Our 3 Favorite Strangle Studies. 23 March 2015. https://www.tastylive.com/news-insights/our-3-favorite-strangle-studies

122. Goetzmann, Ingersoll, Spiegel and Welch. Sharpening Sharpe Ratios. NBER Working Paper 9116. August 2002. https://www.nber.org/system/files/working_papers/w9116/w9116.pdf

123. Red Rock Capital / CME Group. Sortino: A 'Sharper' Ratio. Undated. https://www.cmegroup.com/content/dam/cmegroup/education/files/sortino-a-sharper-ratio.pdf

124. Amundi Research Center. Tail Risk Adjusted Sharpe Ratio. 4 October 2018. https://research-center.amundi.com/article/tail-risk-adjusted-sharpe-ratio

125. Option Sellers Could Capture Time Decay with Short-Dated Index Options. SSRN working paper 5104300. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5104300

126. SEC v. Caine, Parvataneni, LJM Funds Management and LJM Partners, complaint. 2021. https://www.sec.gov/files/litigation/complaints/2021/comp-pr2021-89.pdf

127. Barchart. $CNDR Performance. Accessed 19 August 2026. https://www.barchart.com/stocks/quotes/%24CNDR/performance

128. tastylive. Do Smaller Losses Lead to Better Performance? 6 May 2024. https://www.tastylive.com/news-insights/Do-Smaller-Losses-Lead-to-Better-Performance

129. tastylive. Alternative to Managing Losers. 20 October 2016. https://www.tastylive.com/news-insights/alternative-to-managing-losers

130. tastylive. Straddle & Strangle Theory and Practice. 3 December 2015. https://www.tastylive.com/shows/the-skinny-on-options-data-science/episodes/straddle-strangle-theory-and-practice-12-03-20

131. tastylive. What Makes Tail Risk Risky in Short Premium Trades. 29 April 2025. https://www.tastylive.com/shows/market-measures/episodes/what-makes-tail-risk-risky-in-short-premium-trades-04-29-2025

132. ORATS. *Options market width during the drawdown*. 3 March 2020. Vendor measurement, not peer-reviewed. https://orats.com/blog/traders-complain-about-options-market-width-here-are-the-worst-offenders

133. Alexiou, Goyal, Kostakis and Rompolis. Pricing Event Risk: Evidence from Concave Implied Volatility Curves. Review of Finance 29(4):963–1007. 2025. https://revfin.org/pricing-event-risk-evidence-from-concave-implied-volatility-curves/

134. MacLean, Thorp and Ziemba. Long-term capital growth: good and bad properties of Kelly and fractional Kelly. August 2010. https://escholarship.org/uc/item/5mr5k8qj

135. Sinclair and Brooks. Skewness and Kurtosis of European Options and Implications for Trade Sizing. SSRN 2956161. c. 2017. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2956161

136. Talton CM. The Kelly Criterion and Option Trading. October 2017. http://taltoncm.blogspot.com/2017/10/the-kelly-criterion-and-option-trading.html

137. INTL FCStone subsequent-events note, SEC EDGAR. December 2018. https://www.sec.gov/Archives/edgar/data/913760/000091376018000157/R31.htm

138. Reuters. OptionSellers.com investors hit by natgas swings. 19 November 2018. https://www.reuters.com/article/business/optionsellerscom-investors-hit-by-natgas-swings-intl-fcstone-idUSL2N1XU1BP/

139. Reuters. Allianz to pay $6 bln over Structured Alpha fraud, fund manager charged. 17 May 2022. https://www.reuters.com/business/finance/allianz-pay-6-bln-over-structured-alpha-fraud-fund-manager-charged-2022-05-17/

140. Bloomberg via WealthAdvisor. $600 Million Volatility Fund Hits Margin Wall. 18 March 2020. https://www.thewealthadvisor.com/article/600-million-volatility-fund-hits-margin-wall

141. SEC press release 2021-89. 2021. https://www.sec.gov/newsroom/press-releases/2021-89

142. DOJ / Allianz Global Investors US plea agreement. 2022. https://www.lit-wc.aoshearman.com/siteFiles/39722/EN1b_AGI%20DOJ%20Plea%20Agreement.pdf

143. FINRA. Trading Options: Understanding Assignment. 14 December 2020. https://www.finra.org/investors/insights/trading-options-understanding-assignment

144. Options Industry Council. Options Exercise FAQ. Undated. https://www.optionseducation.org/referencelibrary/faq/options-exercise

145. Options Industry Council. Exercising Options. Undated. https://www.optionseducation.org/optionsoverview/exercising-options

146. Options Industry Council. Options Assignment FAQ. Undated. https://www.optionseducation.org/referencelibrary/faq/options-assignment

147. Christensen and Prabhala. The relation between implied and realized volatility. Journal of Financial Economics. November 1998. https://doi.org/10.1016/s0304-405x(98)00034-8

148. Harvey and Liu. Backtesting. Journal of Portfolio Management 42(1):13-28. 28 July 2015. https://people.duke.edu/~charvey/Research/Published_Papers/P120_Backtesting.PDF