



Reading OptionBench

A working guide to every screen, chart and matrix in the platform — what each number means, how it was computed, and where it stops being reliable.

1. How to read a number here	the two conventions, samples, percentiles
2. The tools, and what each one answers	where to start
3. Today — the Opportunity Map	what opens in the days ahead
4. Weeks — the calendar	two weeks at a glance
5. The scanner table	every row that cleared the filters
6. The expanded row	the edge, the position, the price
7. The Playbook	the four steps, and why they say that
8. The cycles chart	this cycle against every past one
9. The entry × exit matrix	reaching a target and keeping it
10. Best OS ETF	and why it is sometimes empty
11. Ticker-Options Ideas	start from a ticker
12. Trade Doctor	a health check before you trade
13. Forecast by Options	what the chain implies
14. Glossary	every term in one place

Everything here describes historical backtests and live market data. None of it is investment advice.
optionbench.com

How to read a number here

Four ideas govern every figure on the platform. Learn these and the rest of this guide is detail.

1.1 Every return appears under two conventions

This is the most important thing to understand, and it is where most options tools quietly flatter themselves.

A backtest with a +10% profit target has to decide what “the trade hit its target” means. Looking once a day, it records the closing price on the day the threshold was crossed. That close is the first observation *after* the threshold was passed, so it almost never sits on the target. It sits above it.

At the observed exit	The close of the first day the position was up by the target or more. The industry-standard figure, and the one most tools publish alone.
At the target	Exactly the target. What an unattended limit order books when it fills.

How large is the gap

Measured across 4,038 cycle-entries on intraday quotes: 2.8 points on average, and it widens the longer a position is held — about 1.4 points at one day out, over 3 points at eleven. On small pre-earnings returns that is often enough to flip the sign.

Plan on the second number. The first tells you the move happened; the second tells you what you could have taken home.

1.2 A percentage without a sample size is not a claim

A 75% win rate over four cycles and over forty are not the same statement. Four observations put the true rate somewhere between roughly 30% and 95%. Every win rate on the platform carries its cycle count, and you should read the two together, always.

- Under 8 cycles — greyed or flagged. Indicative only.
- 8 to 16 cycles — two to four years. Readable, but one unusual quarter moves the average by several points.
- 20 cycles and above — five years or more. Where a win rate starts to mean something.

1.3 Percentiles compare a setup to itself

When the platform says an entry cost sits at the 21st percentile, it means: across this ticker's own past cycles, measured at the same point in the cycle, only 21% were cheaper. Low is cheap.


This is deliberately not a comparison against other tickers or a broad IV range. A strangle on a quiet utility and one on a semiconductor are not comparable in absolute terms; each is only comparable to its own history.

1.4 Nothing includes commissions

Historical figures are priced from archived option chains. Positions are valued at the ask when buying and the bid when selling, which is conservative, but no brokerage fee is deducted anywhere. Add your own — at small position sizes they are not a rounding error.

The tools, and what each one answers

Each tool answers one question. Knowing which is which saves more time than any shortcut.



Today	What opens in the next few days, ranked by win rate. The screen to start a session on.
Weeks	Every entry date from every scanner, two weeks at a time. Planning rather than acting.
IV Scanner	Where implied volatility sits against a ticker's own recent range. A regime reading, not a setup.
Pre-Earnings	Backtested setups around company announcements. The core scanner.
Pre-Events ETF	The same around macro events — FOMC, CPI, NFP, ISM — on index and commodity ETFs.
Best OS ETF	Credit spreads on ETFs ranked by expected edge per dollar at risk.
Ticker-Options Ideas	Start from a ticker rather than from a scanner.
Trade Doctor	Diagnoses a position you already hold, or are about to open, at today's prices.
Forecast by Options	The probability distribution the option market is implying for an underlying.

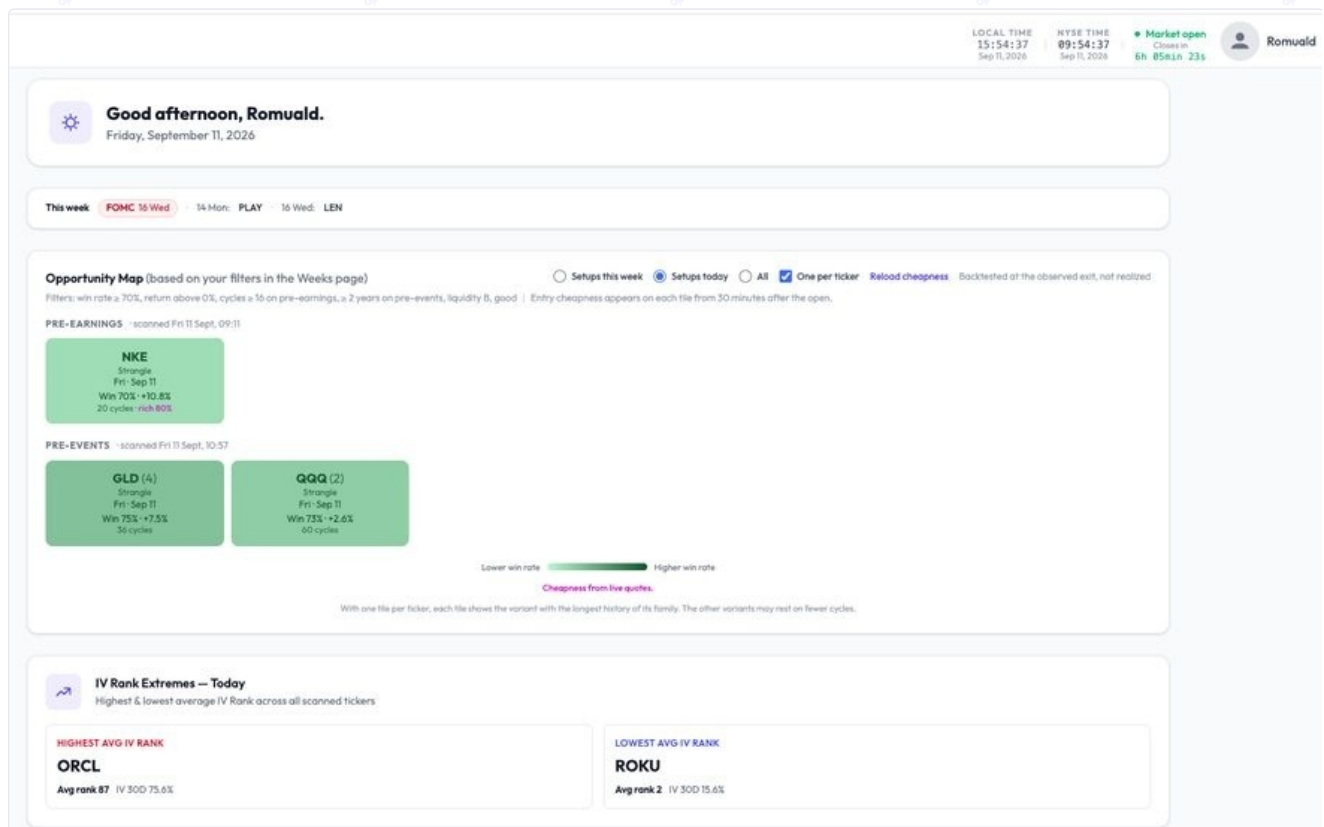
Scanners look backwards, Trade Doctor looks at now

The scanners tell you what past cycles did with this structure on this ticker. Trade Doctor tells you what *this* position is worth and what it risks at today's prices. They answer different questions and will not agree — the gap between them is usually informative rather than a fault.

Everything in this guide applies to the pre-earnings and pre-events scanners alike: they share a layout, and differ only in what triggers the cycle.

Today — the Opportunity Map


The first screen after login. It answers one question: which backtested entry windows open in the days ahead.



The Today page. Market state and both clocks sit top right; the bar under the greeting lists what the week holds; the Opportunity Map follows, then the IV rank extremes across every scanned ticker.

3.1 What the tiles are

Each tile is a setup whose entry day falls inside the selected window. This is not a list of companies reporting this week — it is a list of days to open a position. A ticker reporting three weeks out appears here if its backtested entry day is Friday.




The map on its own, with "Setups today" selected. NKE carries a live cheapness reading in fuchsia; GLD and QQQ show a bracketed count because several variants are collapsed into one tile.

Ticker	A number in brackets — GLD (4) — means four variants cleared the filters and are collapsed into one tile.
Strategy	Straddle, Strangle, Calendar or Diagonal.
Entry day	The session on which the backtested position opens.
Win rate	Share of past cycles that reached the target at some point before the exit.
Return	Average, at the observed exit. See 1.1 — the fillable figure is lower.


Weeks — the calendar

Every entry date from every scanner, laid out two weeks at a time. Where the Opportunity Map is for acting, this is for planning.



Next week, five columns. Yellow cards are pre-earnings, blue are pre-events. "+4 more" expands a day that holds more setups than fit.

Each card names the ticker, the event where there is one, the strategy, and the three figures that decide whether a setup is worth opening: win rate, average return and cycle count.



A single day expanded. Five setups here, two more behind the link. The cycle count on the right is what separates a claim from a coincidence.

4.1 The filters live here

The thresholds set on this page govern the Opportunity Map as well. Raising the minimum win rate or cycle count thins both screens at once, which is usually what you want: a setup you would not take from the calendar is not one you want surfacing on the home page either.

Why the same ticker appears several times in a day

A ticker can clear the filters on several entry days, with several structures, over several cycle depths. Each is a distinct backtest. On the Opportunity Map they collapse into one tile; here they are shown separately because the day is the unit.

4.2 Reading the week rather than the day

The value of this screen is the shape. A week with fifteen setups spread across five days is a normal week. A week where everything lands on one day means several announcements clustered, and the entry windows will compete for the same capital.

The scanner table

Every row that cleared the filters, sortable on any column. Pre-Earnings and Pre-Events share this layout.

Ticker

Strategy

Buy day

Winrate

Return

Liquidity

Earnings

Cycles

All

All

All

Any

Any

Any

Any

Any

Scanner

Calendar

Click any row to expand details

Strategy	Ticker	Buy Day	Cycles	N	Win%	Ret%	Liquidity	Days	Next Date	Time
Long Strangle	MU	T-30	24	24	75.0%	+9.0%	94 A	T-13	2026-09-30	AMC
Long Strangle	NFLX	T-15	4	4	75.0%	+11.7%	91 A	T-27	2026-10-20	AMC
Long Strangle	NFLX	T-14	4	4	75.0%	+13.3%	96 A	T-27	2026-10-20	AMC
Long Strangle	NFLX	T-15	8	8	75.0%	+11.6%	90 A	T-27	2026-10-20	AMC
Long Strangle	NFLX	T-15	20	20	65.0%	+7.2%	94 A	T-27	2026-10-20	AMC
Long Strangle	NFLX	T-14	20	20	69.0%	+9.1%	94 A	T-27	2026-10-20	AMC
Long Strangle	NKE	T-15	12	12	66.7%	+11.9%	81 B	T-14	2026-10-01	AMC
Long Strangle	NKE	T-13	12	12	75.0%	+12.5%	89 A	T-14	2026-10-01	AMC
Long Strangle	NKE	T-11	12	12	66.7%	+10.0%	91 A	T-14	2026-10-01	AMC
Long Strangle	NKE	T-10	12	12	66.7%	+5.6%	92 A	T-14	2026-10-01	AMC
Long Strangle	NKE	T-15	16	16	68.8%	+12.9%	87 A	T-14	2026-10-01	AMC
Long Strangle	NKE	T-13	16	16	75.0%	+15.1%	92 A	T-14	2026-10-01	AMC
Long Strangle	NKE	T-12	16	16	68.8%	+10.5%	91 A	T-14	2026-10-01	AMC
Long Strangle	NKE	T-11	16	16	68.8%	+11.3%	92 A	T-14	2026-10-01	AMC
Long Strangle	NKE	T-15	20	20	75.0%	+14.6%	91 A	T-14	2026-10-01	AMC
Long Strangle	NKE	T-14	20	20	70.0%	+10.8%	92 A	T-14	2026-10-01	AMC

NKE

Long Strangle

Enter T-14 - now T-14

Next earnings - 2026-10-01 (AMC)

LIVE - Intraday

Rich entry
richer than 80% of past cycles

The pre-earnings scanner. Filters across the top, the Scanner / Calendar toggle below them, and one row per backtested (ticker, strategy, entry day, cycle depth). Clicking a row expands it.

Strategy	The structure the backtest holds.
Buy Day	The entry day, in business days before the announcement.
Cycles / N	Cycles is how many past announcements were looked at. N is how many of them produced a backtested trade. They are usually identical; when they are not, a cycle lacked usable option data and dropped out. N is the number the statistics rest on — read it, not Cycles, when judging whether a win rate means anything.
Win%	Share of cycles that reached the target. The bar beside it is the same figure.
Ret%	Average return across all completed cycles, winners and losers alike, as a percentage of the entry premium. Read it with the win rate, not instead of it.
Liquidity	A score out of 100 and a grade. A is 85 and above, B is 70, C is 55. It is built from spread and volume, and a poor grade invalidates everything to its left — a backtested edge you cannot fill is not an edge.
Days	Business days from today to the next announcement. Compare it against Buy Day: if Days is smaller, the entry window has passed.
Next Date / Time	The next announcement and whether it lands before the open (BMO) or after the close (AMC). This decides which session is T-0.

The same ticker on several rows

One ticker can appear a dozen times: several entry days, several structures, several cycle depths. They are not duplicates — each is a separate backtest, and they will disagree. A 4-cycle row showing 75% and a 20-cycle row showing 65% on the same ticker are both true; the second is the one to believe.

The asterisk, and why it is not just a small-sample mark

A figure followed by * and dimmed is flagged as unusually noisy. The two columns are tested differently.

On the win rate, the mark appears when the lower bound of the confidence interval falls below 50% — that is, when the sample is small enough that the true rate could be a coin toss. This is why a 70% over 20 trades can be marked while a 75% over 24 is not: the bound, not the headline, decides.

On the return, it is the average against its own scatter, scaled by sample size. The threshold is deliberately set at 1 rather than the textbook 2: measured across a full scan the median setup sits at 1.47, so a textbook cut would dim 71% of the table and the mark would stop meaning anything. At 1 it flags the noisiest quarter — the +276% on four cycles.

Read it as *unusually noisy*, not as *not significant*. Hovering the figure gives the number behind the flag.

5.1 Sorting is not ranking

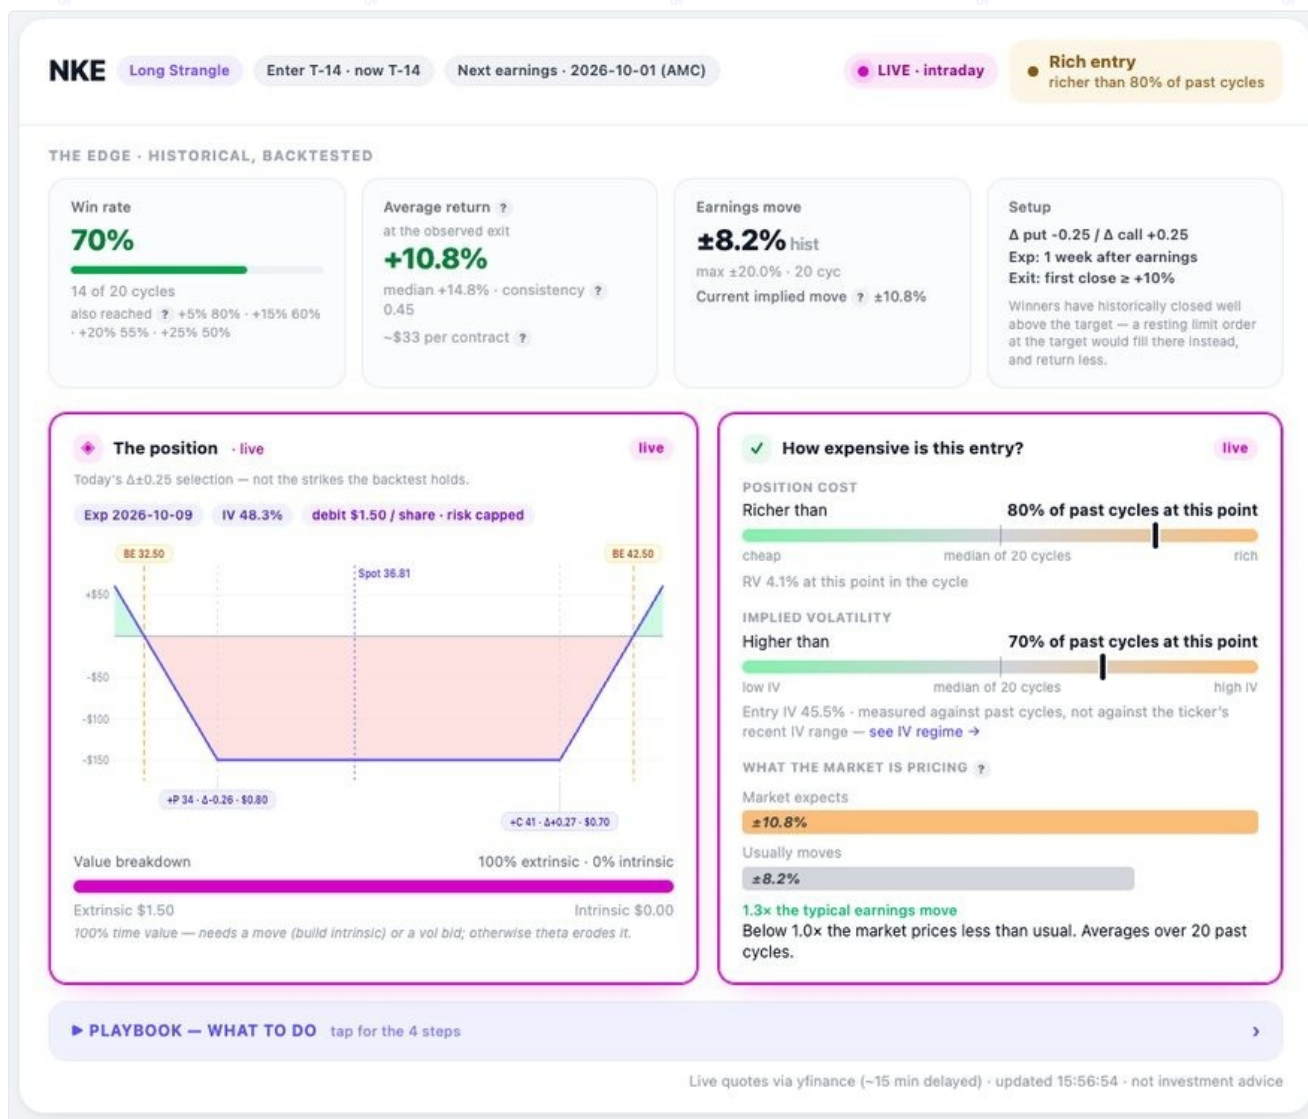
Sorting by win rate puts the thinnest samples at the top, because a small sample produces extreme values more easily. Sort by win rate, then read the cycle column before anything else.

5.2 The Calendar toggle

The same rows, arranged by entry date rather than as a list. Useful once you have narrowed the filters and want to see when the work falls.

The expanded row

Click any row. The header states what you are looking at; four tiles give the historical edge; two panels give today's position and today's price.



A pre-earnings row expanded. The fuchsia LIVE badge means the prices are from the open market; the amber verdict on the right is today's entry against this ticker's own past cycles.

6.1 The header strip

Ticker, strategy, Enter T-14 · now T-14, and the next announcement. On the right, the source badge — LIVE · intraday in fuchsia when the market is open, Last close · EOD in grey when it is shut — then the verdict: Cheap entry or Rich entry, with the percentile behind it.

6.2 The four edge tiles

Win rate	Cycles that reached the target, with the count. The smaller thresholds underneath — “also reached +5% by 80%, +25% by 50%” — show how quickly the odds thin as you ask for more.
Average return	Labelled <i>at the observed exit</i> . Median and consistency sit below. Consistency near zero means the outcomes were scattered and the average is carried by a few cycles.
Earnings move	What the underlying historically did, median and max, over the stated cycles — and, on the last line, what the market is pricing for the coming announcement.
Setup	Deltas, expiration rule, exit rule. Read this first: it defines what the other three tiles describe.

6.3 The position, live

Today's strikes at the stated deltas, the debit, the underlying, implied vol on each leg, a payoff diagram with the break-evens marked, and the value breakdown. 100% extrinsic means the position is pure time value: it needs a move, or a bid for volatility, or theta erodes it.

These are *today's* strikes at $\Delta \pm 0.25$, not the strikes the backtest held in past cycles. They will differ.

6.4 How expensive is this entry

Three readings on pre-earnings rows, and they can disagree — which is the point.

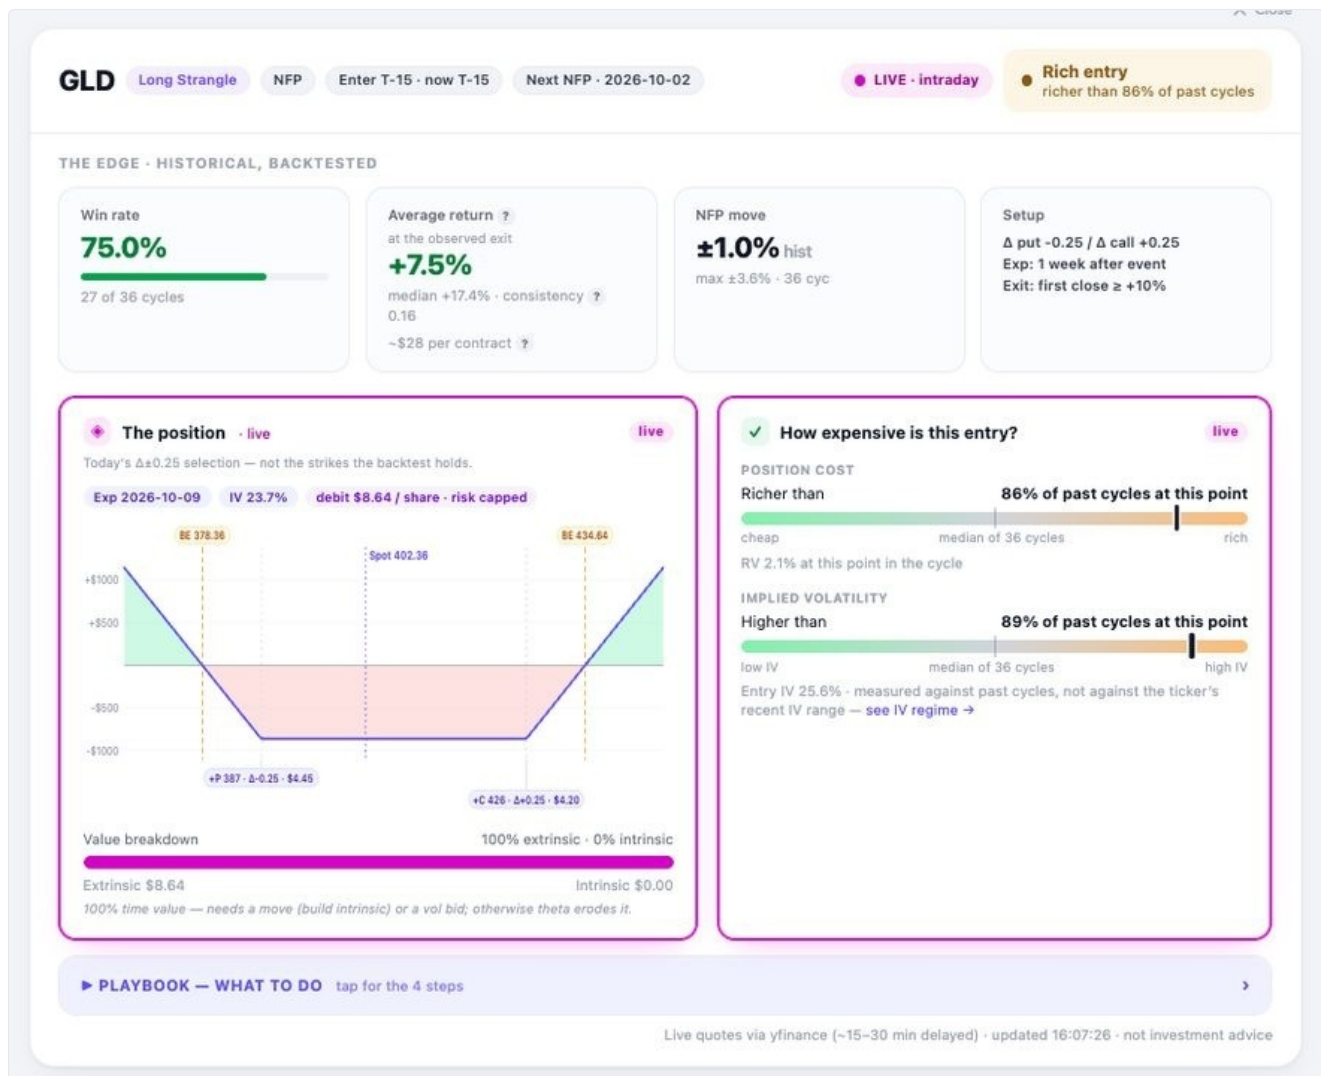
Position cost	Today's debit against the same point in past cycles.
Implied volatility	Entry IV against past cycles. Note this is measured against <i>past cycles at this point</i> , not against the ticker's recent IV range — that is what IV rank does, and it is a different question. The link takes you there.
What the market is pricing	The implied move against the typical realised move. Above 1.0× the option market expects more than this stock usually delivers, which is expensive for a long position.

When two readings disagree

A setup can be cheap against its own history and expensive against what the stock actually does. That is not a contradiction; it means the whole ticker is in a stretched regime. Size accordingly.

6.5 The same layout on pre-events

Pre-events rows read identically. Only the trigger differs: a macro release rather than a company announcement, so the third tile names the event and counts its past occurrences.



A GLD strangle around an NFP print. Same four tiles, same two panels, and the source badge in fuchsia because the market is open.

Why pre-events has only two gauges

The third is absent on macro events on purpose. Measured across a full scan, the median ratio is 4.4× on straddles and 5.1× on strangles, and all 111 rows read as demanding. An index moves 1–2% on an NFP print while the position needs 5–10%, so the ratio is structural rather than informative. A line that says the same thing on every card teaches you to stop reading it.

The Playbook

Four steps under the panels, collapsed by default. It restates the numbers above in plain language and adds what they imply.

▶ **PLAYBOOK — WHAT TO DO** tap for the 4 steps

!
Entry looks rich. The trade is pricier than most past entries here. Across 31,001 pre-earnings cycles, cheaper entries have historically returned more — but waiting shortens the window, and shorter windows have historically won less often.

!
Confirm the fill before trading. Closing quotes overstate the real cost — bid/ask spreads tighten after the open. Send a limit near mid.

- × **Liquidity / spread** — worst leg 39%
- ✓ **Defined risk** — loss capped at premium paid

♦
You are not buying the gap. The position closes before the announcement — the earnings move never enters the result. What it buys is the rise in implied volatility that anticipates the report. The risk is a weaker ramp than usual, not a disappointing move. No directional call needed.

e
Plan the exit up front. Defined risk — most you can lose is the debit paid. The backtest exits at the **close of the first day** the position is up by the target or more, otherwise at the last clean session before the announcement. Winners arrive early: median day 3, and the daily chance of reaching the target falls by about 40% over the first week.

Live quotes via yfinance (~15 min delayed) · updated 15:56:54 · not investment advice

The Playbook expanded. The icons carry the verdict: an exclamation is a caution, a tick a pass, a diamond a point of understanding rather than a check.

7.1 Step one — how the entry is priced

Cheap or rich against this ticker's own history, with the consequence stated: across 31,001 pre-earnings cycles, cheaper entries have historically returned more. It also names the cost of waiting — a later entry is a shorter window, and shorter windows have historically won less often.

7.2 Step two — confirm the fill

The execution check. Closing quotes overstate the real cost because bid/ask spreads tighten after the open, so the step tells you to work a limit near mid rather than take the market. Below it, the two things that can invalidate the trade before it starts: the worst leg's spread, and whether the loss is capped.

7.3 Step three — what you are actually buying

The step most worth reading twice. The position closes before the announcement, so the earnings move never enters the result. What the trade buys is the rise in implied volatility that anticipates the report. The risk is a weaker ramp than usual, not a disappointing move — and no directional call is needed.

7.4 Step four — plan the exit up front


What you can lose, when the backtest exits, and one figure worth acting on: winners arrive early. Median day three, and the daily chance of reaching the target falls by about 40% over the first week.

The practical consequence

A position that has not reached its target within a week or so is usually one where the move did not come, and from there theta does the rest. That argues for a time-based exit alongside the target — not because holding causes the loss, but because the information that it will not work is already in.

The cycles chart

This cycle plotted against every past one, on four measures at once. The densest screen in the platform and the most rewarding.




With only the average shown. The black line is the median across cycles, the shaded band the P25–P75 range, and the green and red dots mark which past cycles reached the target.

8.1 The four panels

Value	What the position was worth, in dollars or as a share of spot — the toggle top right.
Relative Value	The same as a percentage of the underlying: is this trade expensive? Labelled at both ends.
IV %	Implied volatility through the cycle. The ramp into the announcement is what a long-vol position is buying.
Close Δ%	The underlying itself, indexed to the entry day. The panel that explains the other three.


Close Δ% is the one people skip

When a cycle's line leaves the shaded band on this panel, the underlying did something unusual — and that, not the setup, explains the outcome. A position that paid handsomely while the stock fell 18% did not work because the entry was well chosen. Reading the other three panels without this one produces confident wrong conclusions.



The Close $\Delta\%$ panel alone. The black line is the median path of past cycles, the shaded band the middle half of them, and the dashed line the cycle running now.

Read it first, then the other three. Where the dashed line sits inside the band, this cycle is behaving like its predecessors and the value and IV panels can be taken at face value. Where it leaves the band, it cannot.



Every cycle shown. Individual series can be toggled on the right; the purple diamond on the left of each panel is the live value now.

8.2 The legend

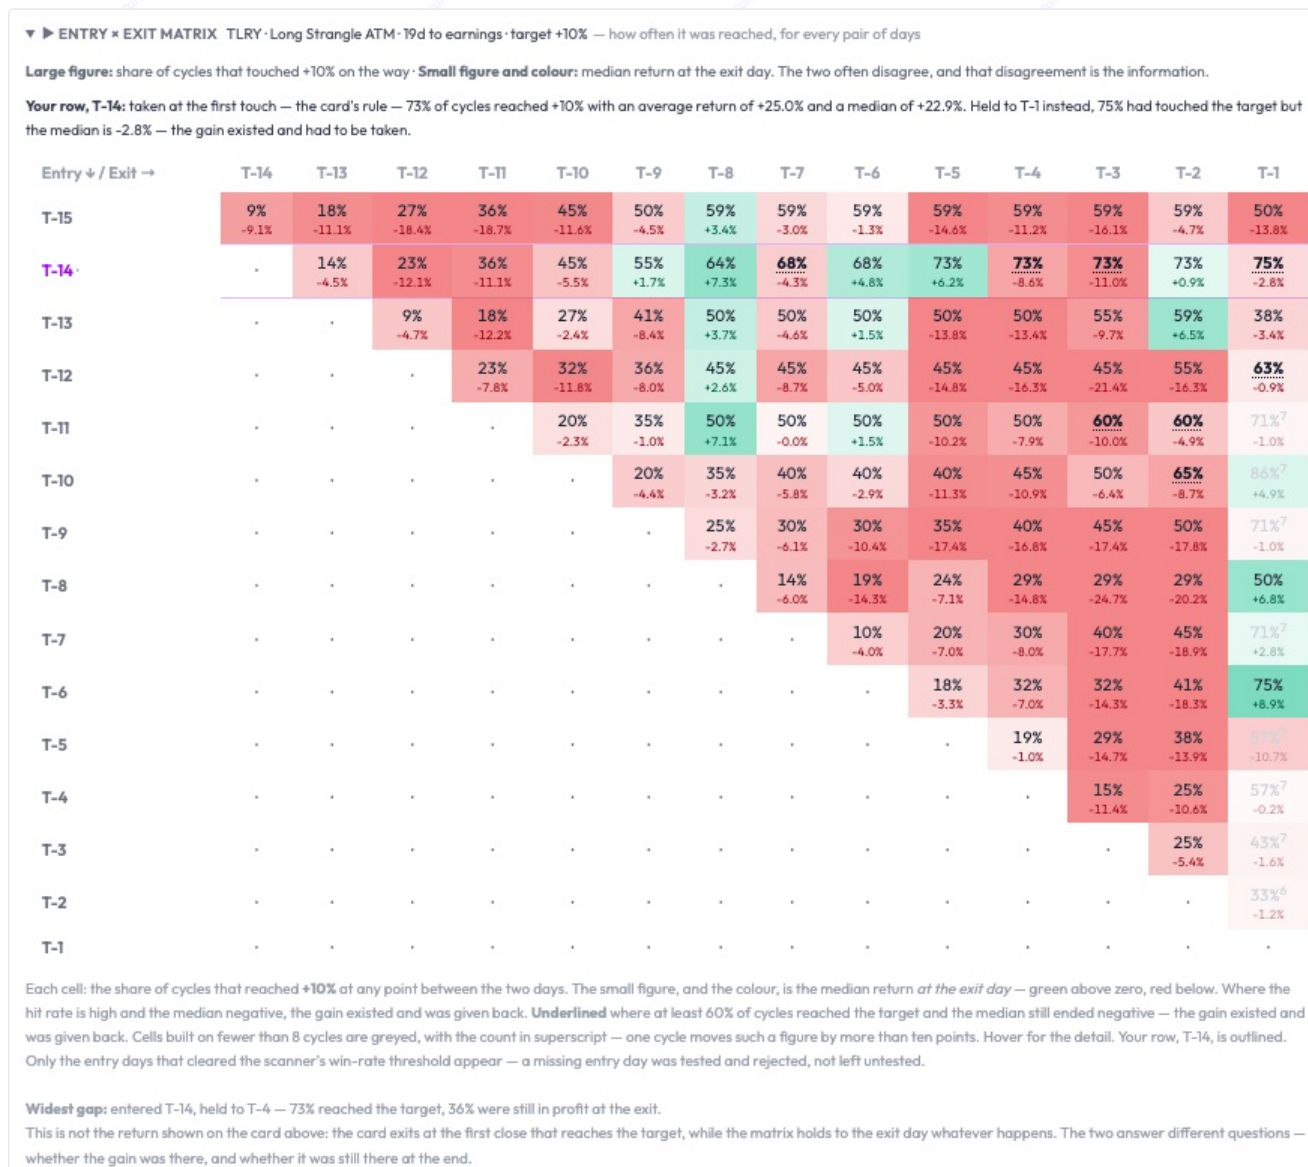
- Green dot — that cycle reached +10% at some point.
- Red dot — it did not.
- Dashed vertical — the buy day on the left, today on the right.
- Purple diamond — the live value, only while the market is open.

The Outcomes tab replaces the per-cycle lines with the distribution of results, which answers “how often and how much” rather than “what did the

path look like”.

The entry × exit matrix

The block below the chart, and the part of OptionBench that answers a question competing tools do not ask.



Rows are entry days, columns exit days. Every cell carries two figures. Hovering it gives the detail behind them: cycle count, median and average at the exit, share still in profit, share that touched the target, and the P25/P75 spread.

9.1 How to read a cell

The large figure The share of past cycles that reached the target at any point between those two days.

The small figure The median return at the exit day. The colour repeats it: green above zero, red below.

Those are two different questions, and the matrix shows them together precisely because they often disagree. A row can read 75% all the way across and be red all the way across: three cycles in four touched the target, and holding past that point gave the gain back.

Your row, in one line Above the grid, a sentence restates your own entry day under both conventions. First, taken at the first touch — the card's rule — with the win rate, the average and the median, which reproduce the card above to the decimal. Then, held to the exit day where the touch rate peaks, with the median there. The gap between the two is what the same cycles paid for not taking the target.

A high figure over a negative one

The pattern worth learning. It means the gain existed and was given back: most cycles touched the target on the way, and most were underwater by the exit day. Cells where at least 60% reached the target and the median still ended negative are underlined.

So a touch rate rising from one column to the next is not automatically an improvement. If the small figure turns red as the large one rises, more cycles poked at the target and fewer held it.

9.2 Greyed cells and missing rows

A cell built on fewer than eight cycles is greyed with its count in superscript — one cycle moves such a figure by more than ten points.

Only entry days that cleared the scanner's win-rate threshold appear. A missing entry day was tested and rejected, not left untested. If a row you expected is absent, the setup was backtested on that day and did not make the cut.

9.3 Why the matrix and the card can differ


The card exits at the first close that reaches the target. The matrix holds to the exit day whatever happens in between. They answer different questions: whether the gain was there, and whether it was still there at the end.

The line above the grid makes this explicit: the first-touch figures on that line are the card's own numbers, recomputed on the same cycles from the same archive, and they match. Everything that follows on the row is the price of holding.

The widest gap line under the matrix names the pair where those two diverge most on this setup. It is the single strongest argument for a resting limit order rather than a held position.

Best OS ETF

The full Best Options Strategy engine across 26 liquid ETFs, keeping the best structure for each — directional spreads, straddles, strangles, condors and butterflies alike.



An empty result. It happens, and it is information rather than a failure: 0 of 26 ETFs cleared the bar on this scan.

10.1 How candidates are ranked

Every candidate is priced at a realistic fill and simulated twice: once at implied volatility, once at the volatility the underlying has actually been realising. The ranking uses the second, net of commissions and exit slippage.

Why the second simulation is the one that counts

At implied volatility every position prices to roughly zero expectancy by construction. That is what risk-neutral pricing means: the premium collected equals the expected cost of buying it back. A screen ranked on that number would be ranking noise.

The gap between implied and realised is where the edge is. A premium seller is paid for that gap and for nothing else.

10.2 Why the screen is sometimes empty

A setup only appears if its expectancy at realised volatility is positive *after* commissions and exit slippage. Three things can empty the screen, and all three are worth knowing:

- Implied sits at or below realised. There is no premium to harvest, so no credit structure pays for its risk.
- The premium is real but thin. It exists and the costs eat it. A position that only works before commissions is not a position.
- Spreads are wide. The realistic fill is far enough from mid that the edge disappears at the point of execution.

An empty screen is therefore a reading of the volatility regime, not a broken scan. It says: today, across 26 liquid ETFs, nothing is paying enough to be worth the risk. That is a more useful answer than a ranked list of marginal trades.

10.3 ETFs only


Single stocks carry pre-earnings implied volatility, which would distort the ranking: their implied sits above realised for a reason that has nothing to do with a harvestable premium. Keeping the universe to ETFs makes the implied-minus-realised comparison mean the same thing on every row.

10.4 Force Scan

The scan takes several minutes across the universe, so it runs on a schedule rather than on every page load. The timestamp under the description tells you how old the result is; Force Scan reruns it now.

Ticker-Options Ideas

Start from a ticker rather than from a scanner. Strategies are ranked by expected return on capital at risk, simulated at realised volatility and net of commissions and slippage.



Two candidates on QQQ. The bell curve on each card is the outcome distribution, with the loss zone, the break-even and the profit zone marked.

11.1 The header line


The ticker, the price it was scanned at, and a badge: LIVE or EOD. The amber note matters — the live spot activates about thirty minutes after the open, and option greeks stay on end-of-day data all session by design. If you scan in the first half hour, re-scan shortly after.

Next to it, the IV rank over twelve months, with a plain reading: debit-friendly when volatility is cheap, credit-friendly when it is dear. That single word decides which half of the strategy list is worth looking at.

11.2 What the cards show

Expected return on risk	The headline figure, and the ranking key. Per share — multiply by 100 for the per-contract amount.
POP	Probability of profit, with its uncertainty band.
credit / debit	Whether the structure collects or pays premium at entry.

Under the list, a line states how many strategies were filtered out for having no positive expected P&L. Eight of ten filtered out is common and is the same reading as an empty Best OS ETF screen.



A card expanded. The trade in full, what it pays when it works, the payoff at expiration, and — in red — the typical loss and the bad case.

11.3 The expanded card

Legs and expiration, deltas, and three groups of figures: if it wins with the average P&L, the reward-to-risk ratio, the maximum gain and the expected days to close; the payoff diagram with spot and break-even marked; and if it loses, which gives the typical loss and the bad case at the

worst five per cent.

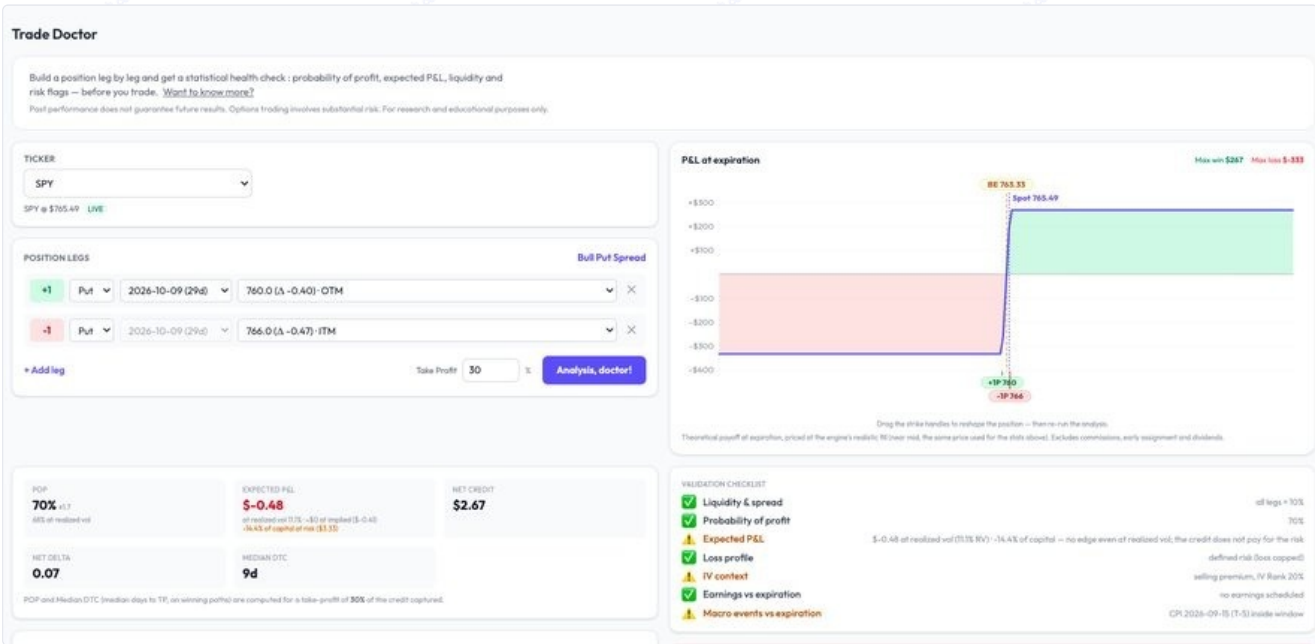
Read the loss box before the return

A high POP with a small expected return is the most common shape here, and the least useful. The typical loss against the maximum gain tells you what you are really being paid: a 1% expected return with a bad case fourteen times the average win is not a favourable trade, whatever the probability of profit says.

Open in Trade Doctor carries the position across for a live check, which is the subject of the next chapter.

Trade Doctor

Build a position leg by leg and get a statistical health check before you trade — probability of profit, expected P&L, liquidity and risk flags.



A bull put spread on SPY. Legs on the left, payoff on the right, the statistics below, and the validation checklist beside them.

12.1 The position builder

Pick a ticker, add legs, set a take-profit. The strike dropdowns carry the delta and whether the leg is in or out of the money. The payoff diagram updates with the position, and the strike handles can be dragged to reshape it — then re-run the analysis.


12.2 The five statistics

POP	Share of simulated paths reaching the take-profit, with its band. A second figure gives the same at realised volatility.
Expected P&L	The mean outcome, stated at realised volatility <i>and</i> at implied, and as a share of capital at risk. See 12.4.
Net credit / debit	What the position collects or pays.
Net delta	The directional exposure. Near zero is market-neutral.
Median DTC	Median days to close on the winning paths. How long the trade typically ties up capital.

POP and Median DTC are both computed for *your* take-profit, not a fixed one. Changing it changes both.

12.3 The validation checklist

Seven checks, each with the figure behind it on the right. A tick is a pass, an amber triangle a caution, and the colour is not decoration: liquidity and loss profile can invalidate everything else. A backtested edge on a position you cannot fill, or whose loss is uncapped, is not an edge.



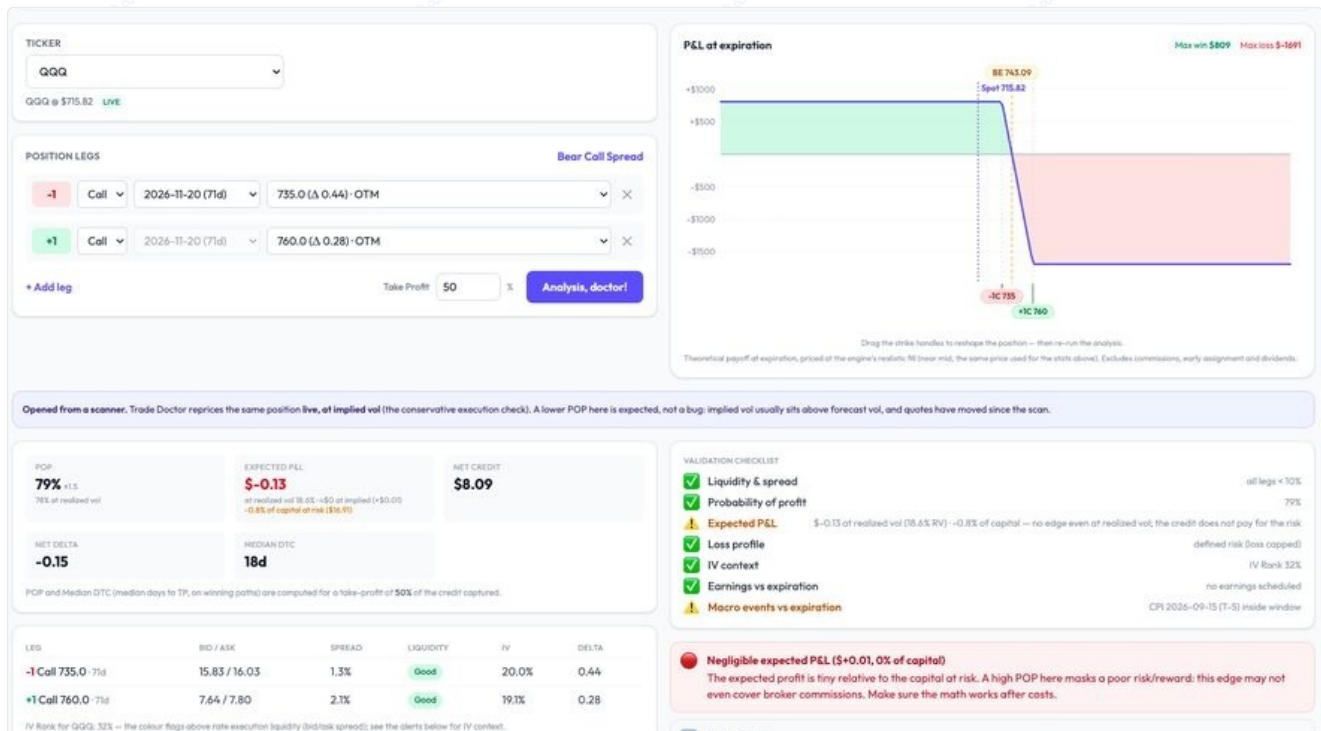
The leg table and the alerts. Red is a finding that should stop the trade; amber is a caution; blue is context.

12.4 Why a credit spread shows an expectation near zero

Simulated at implied volatility, a credit position has an expectation of approximately zero by construction — the credit collected equals the expected cost of buying it back. That is not a defect in the trade.

So the engine runs a second pass at realised volatility, the volatility the underlying has actually delivered. The difference between the two is the volatility risk premium, and it is the real question for a premium seller: is implied above realised, and by enough to pay for the risk?

When the answer is no, the checklist says so in as many words — *no edge even at realised vol; the credit does not pay for the risk* — and a high POP beside it is a warning rather than a comfort.




The same tool opened from a scanner card. The blue banner explains why the numbers differ from the ones on the card.

12.5 Opened from a scanner

When a position arrives from Ticker Ideas or a scanner, Trade Doctor reprices it live, at implied volatility — the conservative execution check. A lower POP than the card showed is expected, not a bug: implied volatility usually sits above forecast volatility, and quotes have moved since the scan.

Forecast by Options

The probability distribution the option market is implying for an underlying — read out of the chain itself, not modelled.



TLT at 21 trading days. The histogram is the implied distribution, split at the threshold; the rising curve is the cumulative probability on the right-hand axis.

13.1 The method

The price of every strike on an expiration carries information about how likely the market thinks each outcome is. Differentiating the option price across strikes recovers the implied distribution directly — the Breeden-Litzenberger result. No assumption about the shape of returns is imposed and no simulation is run: the chain is asked what it already believes.

The header states how many strikes went into it. More strikes means a smoother and more trustworthy density.

13.2 Asking it a question

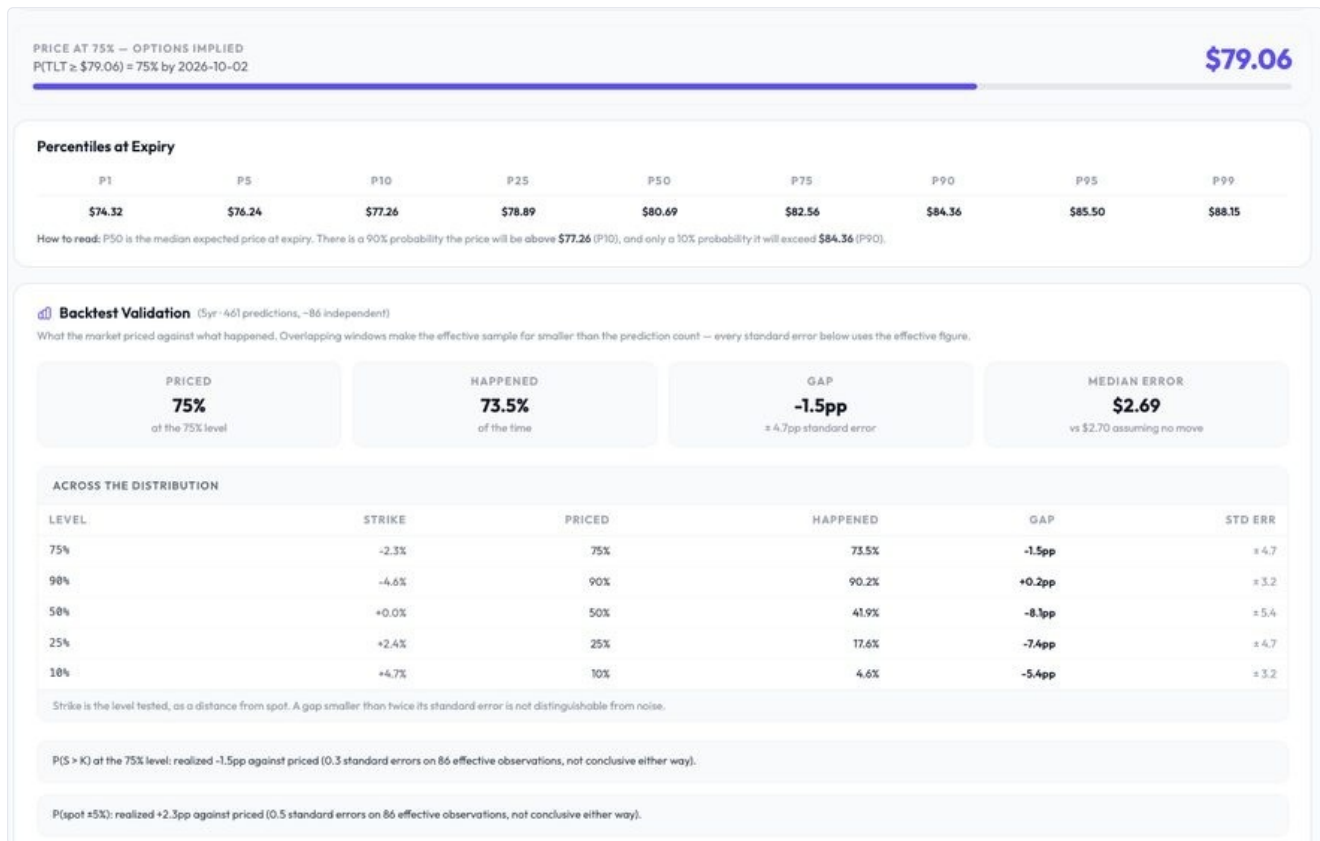
Two directions, on the Target toggle:

- Price → give a level, get the probability of finishing above it.
- Probability → give a confidence, get the level that corresponds.

The percentile table answers the same question at nine points at once. P50 is the median expected price at expiry; there is a 90% probability of finishing above P10 and only a 10% probability of exceeding P90.

Risk-neutral, not real-world

Implied probabilities embed the market's risk premium. They are what you would need to believe for today's option prices to be fair, not a forecast of what will happen. Downside outcomes in particular look more likely than history suggests, because protection carries a premium.



The backtest validation. Five levels tested across the distribution, each with what was priced, what happened, the gap, and the standard error on that gap.

13.3 The validation, and how to read a gap

Every forecast is checked against what actually happened over five years. The figure to read is not the gap alone but the gap against its standard error: a gap smaller than twice its standard error is not distinguishable from noise, and the table says so on every row.

Why the effective sample is smaller than the prediction count

461 predictions, about 86 independent. Rolling windows overlap, so successive forecasts share most of their underlying path and cannot be counted as separate observations. Every standard error shown uses the effective figure, which is why the bands are wider than 461 tests would suggest.

A tool that quoted 461 would look far more confident and would be wrong.

A persistent gap is not model error. It is the variance risk premium and the regime of the window — it is what pays the premium seller.

HOW TO READ THIS

Options-implied probabilities are risk-neutral. A persistent gap is the variance risk premium and the regime of the window, not model error – it is what pays the premium seller. Overlapping windows make the effective sample (86) far smaller than the prediction count (461); all standard errors here use the effective figure.

PARTIAL ROLLING BACKTEST RESULTS

First 10 + last 10 of 461 tests

DATE	SPOT	THRESHOLD	REALIZED	RESULT
2021-09-10	\$148.40	\$145.68	\$146.91	Exceeded ✓
2021-09-10	\$148.40	\$145.24	\$145.35	Exceeded ✓
2021-09-10	\$148.40	\$146.69	\$141.88	Did not ✗
2021-09-24	\$146.91	\$143.59	\$141.88	Did not ✗
2021-09-24	\$146.91	\$143.26	\$145.03	Exceeded ✓
2021-09-24	\$146.91	\$142.82	\$144.13	Exceeded ✓
2021-10-08	\$141.88	\$139.25	\$144.13	Exceeded ✓
2021-10-08	\$141.88	\$138.74	\$147.69	Exceeded ✓
2021-10-08	\$141.88	\$138.04	\$149.30	Exceeded ✓
2021-10-22	\$144.13	\$140.73	\$149.30	Exceeded ✓
2026-07-23	\$83.17	\$81.47	\$82.04	Exceeded ✓
2026-07-23	\$83.17	\$81.43	\$82.05	Exceeded ✓
2026-07-23	\$83.17	\$81.80	\$83.00	Exceeded ✓
2026-07-23	\$83.17	\$81.69	\$82.76	Exceeded ✓
2026-08-06	\$82.52	\$81.30	\$82.05	Exceeded ✓
2026-08-06	\$82.52	\$81.20	\$82.88	Exceeded ✓
2026-08-06	\$82.52	\$81.36	\$83.02	Exceeded ✓
2026-08-06	\$82.52	\$80.82	\$82.21	Exceeded ✓
2026-08-20	\$82.34	\$80.92	\$81.95	Exceeded ✓
2026-08-20	\$82.34	\$80.77	\$82.21	Exceeded ✓

The rolling results themselves, first and last ten of the tests, so the validation can be inspected rather than taken on trust.

optionbench.com

Reading OptionBench

optionbench.com

Glossary

At the observed exit	Return measured at the close of the first day the position was up by the target or more.
At the target	Return measured at the target itself — what a resting limit order books.
AMC / BMO	After market close / before market open. Decides which session is T-0.
Consistency	How tightly past outcomes clustered. Near zero means the average rests on a few cycles.
Cycle	One past announcement and the trade around it. Twenty cycles is roughly five years of quarterly earnings.
Defined risk	A position whose maximum loss is capped by construction.
Entry cost percentile	Today's debit against the same point in this ticker's own past cycles. Low is cheap.
Extrinsic value	The part of an option's price that is time and volatility rather than moneyness. It decays.
Implied move	The move the option market is pricing into the announcement.
IV rank	Where implied volatility sits within this ticker's own recent range. Different from the IV percentile against past cycles.
Liquidity score	0 to 100 from spread and volume. A is 85 and above, B is 70, C is 55.
P25 / P75	The range holding the middle half of past outcomes. A wide spread means the median hides a lot.
Realised move	What the underlying actually did around past announcements.
RV%	Position value as a percentage of the underlying price — how rich the structure is, independent of the share price.
T-x	Business days before the announcement. T-0 is the last session that closes before it.
Take profit	The threshold at which the backtest closes. 10% on long-volatility structures, 20% on calendars.
Vol ramp	The rise in implied volatility as an announcement approaches. What a pre-earnings long-vol position is buying.

A closing note on honesty

Several figures here look worse than they would elsewhere. The target-convention return is lower than the observed one. Thin samples are greyed rather than rounded up. Unconfirmed dates are marked. Entry costs are shown at their percentile even when that percentile is the highest on record.

That is deliberate. A tool that is more optimistic than the fill you actually get is worse than no tool, because it costs money to discover.

[See these screens with today's data](#)

Seven-day free trial at optionbench.com — no commitment, cancel anytime.