



TRADERSUITE · PROFESSIONAL INDICATOR SERIES

AMD Po3

Pro

The complete user guide — Accumulation, Manipulation and Distribution detected by a strict state machine, every cycle measured against a reference model in R, and a scoreboard that survives restarts so the sample finally grows instead of resetting.

v1.0.0

tradersuite.online

Contents

01	What this indicator does	03
02	Installing it — step by step	04
03	Anatomy of the chart	05
04	The cycle — a strict state machine	06
05	Accumulation — how a range is found	07
06	Manipulation — the sweep and the return test	08
07	Distribution — the reference model	09
08	The R:R problem — and the gate that solves it	10
09	How outcomes are scored	11
10	The persistent scoreboard	12
11	The panel	13
12	Settings — license & range	14
13	Settings — manipulation & reference model	15
14	Settings — filters & statistics	16
15	Settings — appearance, alerts & NT8 tabs	17
16	Choosing your timeframe	18
17	Presets by market	19
18	Recommended use — the operator's workflow	20
19	Reading the numbers	21
20	Licensing & activation	22
21	What this indicator is not	23
22	Troubleshooting & FAQ	24
23	Changelog	25

How to read this guide

Sections 1–9 explain the engine: what a cycle is, how each stage is proven, and how a trade is measured. Section 10 covers the feature that makes this build different — a scoreboard that persists. Sections 12–15 are the full reference for every setting. **Sections 16–19 are the ones to read before you trade it:** they cover timeframe selection, per-market presets, the daily workflow and how to tell a working configuration from a broken one.

What this indicator does

AMD Po3 Pro finds the three-act structure that precedes many intraday reversals, and then keeps score of how often it actually pays. Act one: the market goes quiet and builds a range — *accumulation*. Act two: price pokes out of that range far enough to trigger the stop orders resting beyond it, then snaps back inside — *manipulation*. Act three: the real move runs the other way, through the range and out the far side — *distribution*. The indicator detects each act with a strict state machine, attaches a full reference trade to every confirmed cycle, and records the outcome in R. Unlike almost every pattern tool, it then tells you the truth about its own results — and remembers them after you close the chart.

The one detail that matters

What separates a manipulation from a genuine breakout is the **return**. Manipulated price comes back inside the range quickly; a real breakout does not. The whole engine is built around that single test — no return within the deadline and the cycle is written off as a breakout, with no trade modelled at all.

Every cycle is one experiment

A cycle cannot skip a stage, cannot exist twice, and always ends in exactly one outcome: target, stop, timeout, expired, breakout or filtered. That discipline is what makes the statistics countable rather than decorative.

Measured, not asserted

Entry, stop and target are fixed on the bar the manipulation confirms, using only data already printed. Outcomes are scored in R with a deliberately pessimistic rule: if one bar touches both the target and the stop, the stop wins.

The sample survives

Results are written to disk and keyed by close timestamp, so reloading the chart keeps your sample and grows it. Change a rule that alters which cycles exist and you automatically get a clean file — the two never mix.

Who it's for

Intraday futures traders who trade stop-runs and reversals, and want the pattern's real hit-rate on *their* instrument rather than a marketing claim. Defaults are tuned for **NQ / MNQ on 5-minute charts** (Section 16 explains why), and the engine is pure price maths — it runs on ES, YM, RTY, CL, GC, crypto and FX with the rescaling described in Section 17.

NQ / MNQ · 5-minute — reference tuning

ES · YM · RTY · CL · GC

Intraday: 1–15 min

Needs 200 bars of warm-up

Non-repainting · closed-bar decisions

A note on honesty

This indicator does not predict anything, and it will not tell you the pattern works. It tells you what happened. If the scoreboard on your instrument reads a negative average R after a hundred cycles, that is the tool doing its job — it has just saved you the tuition. Read Section 19 before you decide what the numbers mean.

Installing it — step by step

Installation is the same as every TraderSuite indicator and takes about two minutes. You need NinjaTrader 8 (64-bit) and the file from your purchase email or account download page.

- 1 **Download the file — leave a zip zipped.** If your download is a .zip, NinjaTrader imports the whole archive; do not extract it first.
- 2 **Import into NinjaTrader.** Control Center → **Tools** → **Import** → **NinjaScript Add-On...**, browse to the file, confirm. Remove any older copy first so NinjaTrader compiles a single clean version.
- 3 **Restart NinjaTrader.** A restart after import guarantees the new assembly is loaded.
- 4 **Add it to a chart.** Right-click → **Indicators...** → open the **TraderSuite** folder → double-click **TS_AMD_Po3_Pro**. Use a **5-minute NQ chart** to start (Section 16).
- 5 **Enter your license key.** In the settings find **0A. License** and paste the key from your purchase email. Leave it blank to start a 7-day trial.
- 6 **Click OK and wait for warm-up.** Nothing is drawn until 200 bars have loaded — the compression percentile needs its distribution first. The panel's **Phase** row reads **IDLE** and **Compression** shows a value once the engine is live.

Requirements

NinjaTrader 8 (64-bit), any recent build · internet for license validation (24-hour cache, 7-day offline grace) · a data series with **at least 250 bars** loaded. On a 5-minute chart that is under two days of futures data · write access to your NinjaTrader user folder, for the statistics file.

Where the statistics live

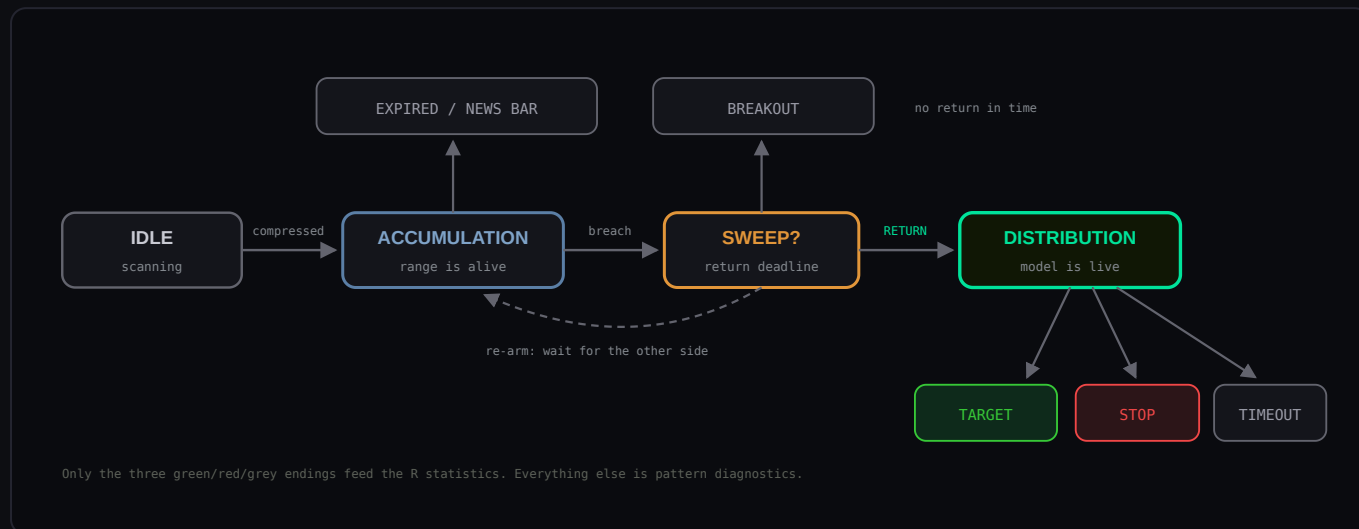
Documents\NinjaTrader 8\bin\Custom\TraderSuite\stats\
One CSV per instrument, per timeframe, per rule set. Delete a file to reset that sample deliberately; nothing else clears it.

Tip — settle your settings before you start collecting

Changing a rule that alters which cycles exist starts a brand-new statistics file. That is correct behaviour — you should not mix results from different rules — but it does mean a week of tuning produces a week of tiny samples. Decide your configuration first, then let it run untouched. Section 18 has the routine.

The cycle — a strict state machine

Every cycle walks the same path, and cannot skip a step. That is not a stylistic choice — it is what makes the statistics meaningful. If a "cycle" could mean three different things depending on how you squint at the chart, counting them would be worthless. Here each detected cycle is one countable experiment with exactly one ending.



Ending	Counts as	Meaning
TARGET	+R (statistics)	The fib-extension target was reached first.
STOP	-1R (statistics)	The stop was hit first — or the same bar touched both, which is booked as a stop.
TIMEOUT	actual R (statistics)	Neither level in 64 bars; closed at the real R from the close, not silently dropped.
BREAKOUT	diagnostic only	The sweep never returned. No trade was modelled, so there is nothing to score.
EXPIRED	diagnostic only	The range aged out before any sweep arrived.
NEWS BAR	diagnostic only	One bar breached both boundaries — proves nothing, cycle reset.
R:R / HTF REJECT	diagnostic only	A confirmed manipulation refused by a filter (Sections 08 and 14).

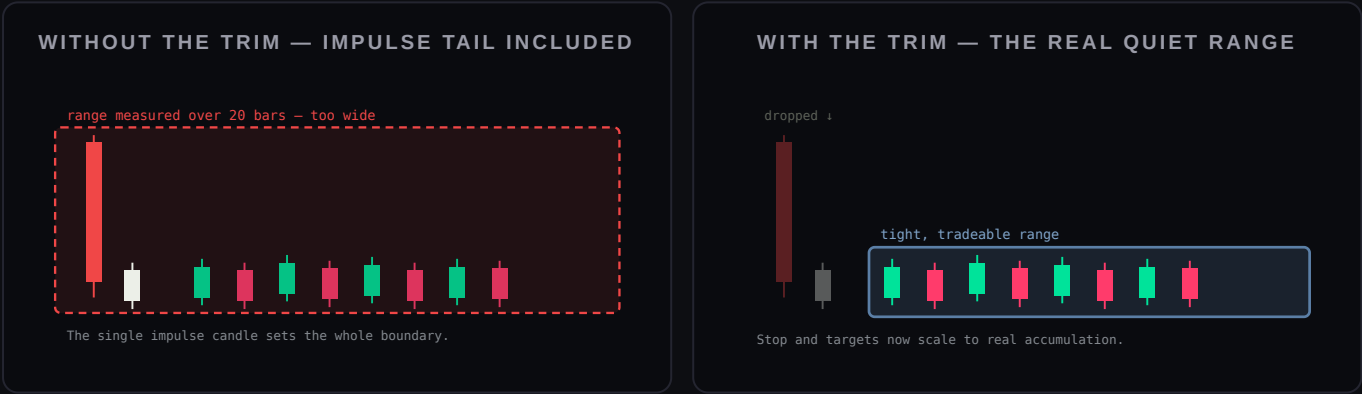
The re-arm — a second chance, once

If a confirmed manipulation is rejected by the R:R gate or the HTF filter, the range is not thrown away. It **re-arms**: the same boundaries stay live, the clock resets, and it waits for a sweep of the *opposite* side. Repeating the same side is blocked. A cycle can re-arm at most once, and cycles that traded after a re-arm are counted separately on the panel as ds — double sweeps — so you can see whether the second attempt behaves differently from the first.

Accumulation — how a range is found

A range is not drawn because price looks sideways. It is drawn because the market is measurably quiet. On every bar the engine measures the 20-bar high-low span and ranks it against its own last 200 bars. Only when that span sits in the bottom quartile of its own history does a range become eligible. "Quiet" therefore means quiet *for this instrument, right now* — the same rule works on a calm Tuesday and inside a volatile week without retuning.

Gate	Default	Why it exists
Compression percentile	≤ 25	The 20-bar width must be in the bottom 25% of its own 200-bar distribution. This is the master sensitivity control.
Min range width	0.15% of price	Rejects microscopic ranges where the stop and target would be indistinguishable from the spread. ~30 points on NQ at 20,000.
Cooldown	10 bars	A new range cannot be anchored immediately after the previous cycle ended — prevents the engine chattering through one messy area.
Impulse-tail trim	15%	Drops the oldest window bars while each still dominates the width, so the tail of the preceding impulse leg is not measured as part of the quiet range.
Boundary method	Pivot	Boundaries come from confirmed swing pivots, clamped so they can never sit inside price that has already printed (Section 10).



Why the trim matters more than it sounds


Every level downstream is derived from the range: the stop sits beyond it, R is measured from it, the target extends from it. A boundary inflated by one leftover impulse candle produces an oversized stop, a distant target and a distorted R for the whole cycle. The trim is a one-line rule that keeps the entire measurement honest.

The range can breathe

Once anchored, boundaries may drift within the tolerance (10% of width by default) without resetting the cycle — real ranges are rarely perfectly flat. Anything beyond the tolerance is a breach, and the sweep test in Section 06 takes over.

Manipulation — the sweep and the return test

This is the page that defines the whole indicator. Below and above every established range sit resting stop orders. Price poking beyond a boundary triggers them, which supplies exactly the liquidity a larger participant needs to fill the opposite way. The tell is not the poke — it is what happens next. A manipulation **comes back inside quickly**. A genuine breakout does not.



Every test a sweep must pass

Test	Default	Purpose
Range is mature	≥ 12 bars	A breach before this is an early break, not a sweep — the range never really formed.
Beyond the tolerance	10% of width	Ordinary breathing at the edge is not a sweep.
Returns inside	≤ 8 bars	The defining test. No return in time and the cycle is a breakout.
Not the re-armed side	—	After a re-arm, only the opposite boundary may be swept.
Depth cap	off (100)	Optional: reject excursions deeper than a chosen True Range percentile — a real breakout wearing a sweep's clothes.
Liquidity beyond	off	Optional: require equal highs / lows out there (2+ confirmed pivots within $0.1\times$ width) — proof the stops were actually resting.
Session window	off	Optional: only count sweeps inside your chosen hours.

Same-bar sweeps count too

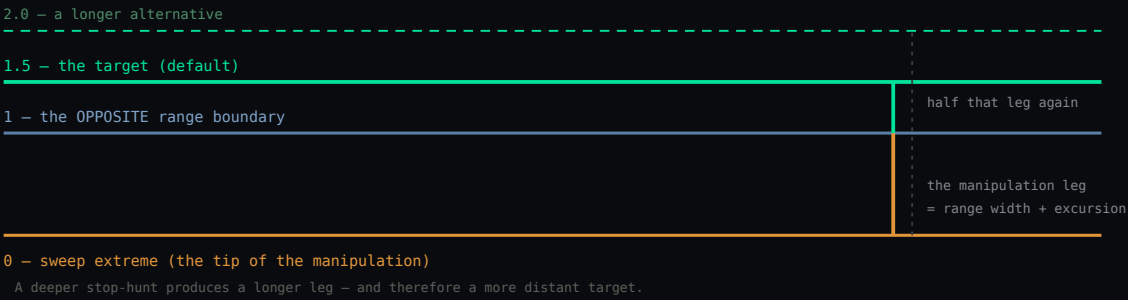
A long wick that spikes beyond the boundary and closes back inside on the *same* bar is the purest form of the pattern, and it is handled on that bar rather than being missed. The excursion tracked for the stop always includes that wick, so the stop sits beyond the full move — not beyond the close.

Distribution — the reference model

The moment a manipulation confirms, a complete trade is fixed — using only data already printed. No waiting for another candle, no lookahead, nothing that could be revised later. That is what lets the statistics mean something: every recorded outcome came from a trade that was fully specified before the next bar opened.

Level	How it is derived	Reasoning
Entry	close of the confirming bar	The first honest price at which you could have acted on the signal.
Stop	sweep extreme $\mp 0.4 \times \text{ATR}$	Beyond the <i>full</i> excursion including its wick, plus a buffer — a retest of the swept tip does not knock the model out.
R	entry – stop	The risk unit. Floored at 8 ticks so a degenerate setup cannot fake a huge R.
Target	fib extension of the manipulation leg	Scales with the move that created the setup, not with a fixed multiple.

THE FIB GRID — ANCHORED TO THE MANIPULATION, NOT THE RANGE



The grid runs from the sweep tip (0) to the far side of the range (1). The target sits at the extension level — 1.5 by default, meaning half the leg again beyond the range.

The two-unit shape, and why this build has none

This is an **indicator**, not an autotrader. The reference model is a single unit that runs to the target, the stop, or the timeout — no scale-out, no breakeven move, no trailing. That is deliberate: partial exits improve a live trader's equity curve but make the *measurement* ambiguous, because you can no longer say what one cycle was worth in R. A clean single-unit model gives an unambiguous number. What you do with real orders around that number is your business, and Section 18 has suggestions.

Entry at the confirming close is pessimistic — on purpose

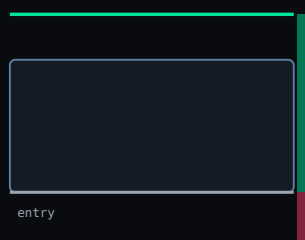
That bar is by definition the reversal candle, so entering at its close means buying after the snap-back has already happened. It is the worst realistic fill. Real execution with a resting limit nearer the boundary should beat the modelled number, which means your live results have room to be *better* than the scoreboard — never the other way round because of an optimistic assumption.

The R:R problem — and the gate that solves it

Two cycles can look identical on the chart and be worth completely different amounts. The reason is one uncontrolled variable: *where inside the range the confirming bar closes*. Entry is that close. The stop is fixed beyond the sweep. So a bar that closes just inside the boundary produces a small risk and a long run to target, while a violent reversal bar that closes near the far side produces a large risk and almost no room left. Same pattern, opposite economics.

ONE SWEEP, THREE CLOSES, THREE COMPLETELY DIFFERENT TRADES

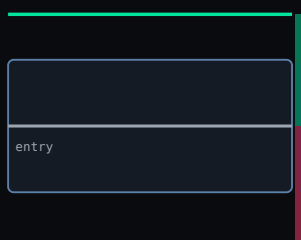
closes just inside



R:R ≈ 3.7

excellent – take it

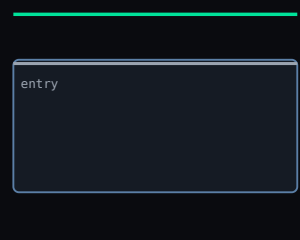
closes mid-range



R:R ≈ 2.0

normal – the default case

closes near the far side



R:R ≈ 0.4

risks more than it can win

Identical range, identical sweep, identical target — only the confirming close differs. The red bar is risk, the green bar is reward.

The Min reward:risk gate — default 1.5

Before a distribution opens, the engine computes R:R and refuses anything below your floor. The rejection is counted separately on the panel as R:R rejected, so you can see how much the filter is doing, and the range is given a chance to re-arm rather than being discarded. **This is the single highest-impact setting in the indicator.** Raising it from 1.5 to 2.0 typically removes the worst quarter of setups; lowering it to 0 reproduces the unfiltered behaviour of the original design, which took the 0.4 case above without comment.

A second-order effect worth knowing

Risk grows roughly one-for-one with how deep the sweep runs; the target only moves half as fast, because it is half a leg beyond the range. So **deeper, more dramatic-looking stop-hunts carry systematically worse payoff** than modest ones. If you find yourself drawn to the violent sweeps, check the R:R field before agreeing with yourself.

How outcomes are scored

From the bar *after* the entry, every closed bar is checked against the target and the stop until one resolves or the clock runs out. Three rules make the resulting number trustworthy.

1 · Ambiguity always loses

If one bar's range touches **both** the target and the stop, OHLC data cannot tell you which came first. The engine refuses to guess in its own favour and books a full $-1R$ stop, incrementing a separate ambiguous counter so you know how often it happened. Your live results can only be equal or better than this assumption.

2 · Timeouts are counted, not dropped

A cycle that reaches neither level within 64 bars closes at its **actual** R from the close — which might be $+0.4R$ or $-0.6R$. Silently discarding these, as many tools do, quietly removes every mediocre outcome and flatters the average.

3 · Percentages wait for a sample

Until the minimum sample (10 cycles by default) is collected, the panel shows raw counts and the word *collecting* instead of percentages and average R . A 78% win rate built on four trades is noise wearing a suit.

4 · Diagnostics stay out of the R stats

Breakouts, expiries, news bars and filter rejections never had a trade, so they are counted on a separate diagnostics row. They answer a different question — *how often does this market play the game at all?* — and mixing them into the win rate would be meaningless.

The arithmetic, in one line

```
// per closed cycle
R = target hit → |target - entry| / |entry - stop|    (typically +1.5 to +3.0)
R = stop hit   → -1.00    (always exactly one R, by definition)
R = timeout    → (close - entry) / |entry - stop|    (anything in between)

Avg R =  $\Sigma R$  / number of closed cycles ← the only number that matters
```

Break-even is roughly a 33–35% target rate

At default settings the model is about 2:1, so one winner pays for two losers. That gives a concrete pass mark: if the panel's target rate is comfortably above ~35% and average R is positive, the configuration is working on this instrument. If target rate looks respectable but average R is negative, you are taking too many low- R : R setups — raise the gate (Section 19).

What is NOT included

Commission, spread, slippage and partial fills. The model assumes a fill at the bar's close and clean touches at the levels. On NQ, round-turn costs are roughly 5% of a typical R at 5-minute settings — real but not decisive. On faster charts with smaller stops that share grows, which is one of several reasons Section 16 recommends against 1-minute for live use.

The persistent scoreboard

This is the feature that makes the indicator worth running. Almost every pattern tool that keeps score keeps it in memory, which means the sample dies when you reload the chart — and dies again every time you nudge a setting while exploring. You can never accumulate enough closed cycles to know anything. Here the results are written to disk, so the sample grows across sessions, weeks and months.

≈10
CYCLES ≈ 4 DAYS
ON 5-MINUTE

≈100
CYCLES ≈ 6 WEEKS
THE USEFUL SAMPLE

4,000
ROWS RETAINED
PER RULE SET

±5pt
STANDARD ERROR
AT 100 CYCLES

How it behaves — the two rules that matter

Reloading keeps your sample
Close the chart, restart NinjaTrader, come back next week — the scoreboard is where you left it. Historical replay re-derives the same cycles, and they de-duplicate against what is already stored using their close timestamp, so nothing is ever double-counted. Loading *more* history simply adds the older cycles.

Changing the rules starts a clean file
Every setting that changes *which cycles exist* is hashed into an 8-character rule-set id, which forms part of the filename. Adjust the compression threshold or the R:R gate and you get a fresh, empty sample — correctly, because mixing outcomes from different rules would produce a meaningless average.

Change this...	Effect on the stored sample
Reload the chart · restart NT8 · change instrument tab	Preserved and continues to grow
Panel size · colours · stats period · min sample · alerts	Preserved — display only
Compression · range width · trim · boundary mode · maturity · age	New rule set → new file
Return deadline · depth cap · liquidity · re-arm	New rule set → new file
Stop buffer · fib extension · R:R gate · risk floor · timeout	New rule set → new file
HTF bias · session window	New rule set → new file
Different instrument or different timeframe	Separate file, always

Where the file lives
Documents\NinjaTrader 8\bin\Custom\TraderSuite\stats\po3_NQ_Minute5_a1b2c3d4.csv
One row per closed cycle: timestamp, outcome, R, and whether it came after a re-arm. Written on real-time closes (at most once a minute) and on shutdown. The rule-set id is displayed on the panel's **DIAGNOSTICS** row, so a support screenshot always identifies which configuration produced it. To reset a sample deliberately, delete the file — nothing else clears it.

Recommendation
Settle your configuration *before* you start counting. Spend a few days in Playback exploring settings — those samples are disposable. Once you have chosen, write down the rule-set id from the panel, then leave the settings alone for at least six weeks. That single habit is the difference between having an opinion about this pattern and having evidence.

SECTION 11

The panel

Bottom-left of the chart, always on, refreshed live. Four sections answer four different questions: is the licence healthy, what is the engine doing right now, is this configuration working, and how often does this market play the game at all.

TS

tradersuite.online

TS_AMD_Po3_Pro v1.0.0

--- LICENSE STATUS ---

Status: ✓ LICENSED

Activations: 1 / 2

--- CYCLE ---

Phase: DISTRIBUTION (7b)

Compression: 18p

Long @ 20145.25

Stop: 20132.50

Target: 20170.75

R:R: 2.00R

--- STATISTICS · ALL-TIME ---

Cycles: 118 (ds 14)

Target: 46 (39.0%)

Stop: 58 (49.2%)

Timeout: 14 (11.9%)

Avg R: +0.171R

Form:

--- DIAGNOSTICS ---

Breakout / Fail: 31 / 4

R:R rejected: 22

Ruleset: a1b2c3d4

Stats persisted · Free for Members

Reading each row

Row	What it tells you
Status	Licence state in colour. Red means no cycles are detected at all.
Phase	The engine's current act, with bars spent in it: IDLE · ACCUMULATION · SWEEP? · DISTRIBUTION.
Compression	Live percentile of the 20-bar width. Green when below your threshold — a range could anchor now.
Range / levels	The active boundaries when idle; entry, stop, target and R:R when a distribution is live.
Cycles	Closed cycles in the selected period. ds = how many came after a re-arm.
Target / Stop / Timeout	The three outcomes with percentages, hidden until the minimum sample is reached.
Avg R	The number that matters. Green above +0.15R, amber near zero, red below -0.15R.
Form	The last ten cycles, newest right: target · stop · timeout.
Breakout / Fail	Sweeps that never returned, and confirmations lost to filters. High breakout counts mean this market trends out of its ranges.
R:R rejected	How many confirmations the gate refused. Section 19 explains what to do with this.
Ruleset	The 8-character id of the configuration producing these numbers.

Compact mode

Set **Panel Size** to Compact to drop the level detail, timeout row, form strip and diagnostics. The panel also auto-compacts if the chart window is too short for the full layout, so it never overflows the bottom of the screen.

Settings — license & range

Sections 12–15 document every input in the order NinjaTrader shows them. Defaults are the shipped values, tuned for NQ on 5-minute. **Remember that changing anything in groups 1–4 starts a new statistics sample** — that is correct, but it means tuning and measuring are separate activities.

0. About · 0A. License

Option	Default	What it does
Website	tradersuite.online	Informational.
Show Logo / Show Panel	On / On	The crest, and the master toggle for the statistics panel.
Panel Size	Full	Compact drops levels, timeout, form and diagnostics. Auto-compacts on short charts anyway.
Statistics Period	All-Time	Window the scoreboard aggregates: last 24 hours, last 30 days, or everything ever recorded for this rule set. Display only — never changes the file.
License Key	(empty)	Masked as ●●●●● and deliberately kept out of the chart label and Control Center columns. Blank = trial.

1. Range — where accumulation is found

Option	Default	What it does & how to tune it
Min range maturity	12 bars	A breach before this is an early break, not a sweep. Raise on noisy charts to demand a properly formed range.
Max range age	96 bars	Time budget before the range expires unswept. A re-arm receives a fresh budget from the re-arm bar.
Compression threshold	25	The master sensitivity control. Raise to 35 for more cycles of lower quality; drop to 15 for only the tightest coils. Every point changes how many samples you collect per week.
Distribution window	200 bars	Lookback for the width and depth percentiles, and therefore the warm-up. Larger = more stable thresholds, longer wait before anything draws.
Boundary tolerance	0.10	How far the boundaries may breathe (as a fraction of width) before a move counts as a breach.
Min range width	0.15%	Percentage of price, so it scales across instruments. The setting most likely to need changing on a new market or a faster chart — see Sections 16 and 17.
Boundary method	Pivot	Pivot uses confirmed swings, clamped so a boundary can never sit inside already-printed price. Absolute uses raw window extremes — slightly more aggressive, slightly noisier.
Impulse-tail trim	15%	Drops leading bars while each dominates the width (Section 05). 0 disables it — not recommended.

Settings — manipulation & the reference model

2. Manipulation — what counts as a sweep

Option	Default	What it does & how to tune it
Max bars until return	8 bars	The close must come back inside the range within this many bars or the move is logged as a breakout and nothing is modelled. Raise on slow instruments where the return takes its time; lower it if you are picking up sluggish drift-backs that were never a rejection.
Require liquidity beyond boundary	Off	Demands that two or more confirmed pivots sat within 0.1× width of each other beyond the boundary — the equal highs or equal lows that give the sweep something to hunt. Turning this on typically halves cycle count and raises quality. Worth testing once you have a baseline sample.
Soft depth cap	100 (off)	Caps how far beyond the boundary price may travel, expressed as a True Range percentile. Set to 85–90 to reject deep excursions that are really breakouts wearing a sweep's clothing. Off by default because it also cuts legitimate violent sweeps.
Allow re-arm	On	When a filter refuses the projected direction, the range survives and waits for a sweep of the <i>opposite</i> side rather than dying — the classic double-sweep. Re-armed cycles are counted separately on the panel as ds; a large ds share means this market habitually raids both sides before committing.

3. Reference model — where the three lines go

Option	Default	What it does & how to tune it
Stop buffer past sweep	0.4 × ATR	How far beyond the swept extreme the stop sits, so a retest of the wick does not knock the model out. Raising it lowers R:R on every cycle — a wider stop is a bigger denominator. 0.25–0.6 is the sane band.
Fib extension target	1.5	Extension of the manipulation leg measured from the sweep extreme. 1.0 lands exactly on the far boundary; 1.272, 1.5, 1.618 and 2.0 are the usual choices. Higher = fewer targets hit but each worth more. This is a direct trade of hit-rate against payoff — see Section 08.
Min reward:risk	1.5	FIX 2 — the most important setting in the indicator. Refuses confirmations whose geometry is poor. 0 disables the gate and reproduces the original script's behaviour. See the box below.
Min risk (ticks)	8	FIX 3. Floor on entry-to-stop distance. Without it, a confirming close that happens to print a hair from the sweep extreme produces a microscopic risk and a fictional +14R that poisons the average. Set it to roughly one-third of a typical bar range on your instrument.
Distribution timeout	64 bars	If neither level is touched within this many bars the cycle closes as a timeout, scored at the actual R implied by the close. Scale this with your timeframe — Section 16 has the numbers.
Show Entry / Stop / Target	On	The three dashed lines. Off leaves boxes and markers only.
Show right-edge price tags	On	Rounded price labels at the right edge of the chart for the live cycle's levels.

Why Min reward:risk exists

The entry is the confirming bar's close, and nothing in the model controls *where inside the range* that close lands. Close back just above the swept low and the stop is near while the target is a whole range away — a 3.5R setup. Drift back and close mid-range and you have paid most of the distance to the target while the stop stayed put — a 0.5R setup. **The two look identical on the chart**, and the original script scored both.

Settings — filters & statistics

4. Filters — when a cycle is allowed to exist

Option	Default	What it does & how to tune it
Use HTF bias	Off	Only models a cycle when its direction agrees with the higher-timeframe trend. Enabling it attaches a secondary data series , so the indicator reloads and needs more history. Expect roughly a 40–50% cut in cycle count for a meaningful lift in target rate on trending instruments.
HTF minutes	240	The bias timeframe in minutes. Must be well above the chart timeframe to mean anything: 240 (4-hour) against a 5-minute chart is a sensible ratio; 15 against 5 is noise.
HTF EMA period	50	Bias is simply HTF close above or below this EMA. Longer = slower, stickier bias.
Use session window	Off	Sweeps count only inside the window. Times are chart bar times in your PC time zone , not exchange time — set them by looking at the chart, not by converting from a session table.
Session start / end	09:30 / 16:00	The window itself. A start later than the end is read as an overnight window that wraps midnight, which is what you want for FX and index futures overnight sessions.

Add filters one at a time, and give each one a sample

Every filter you enable changes the rule-set id and therefore starts a fresh scoreboard. Turning on HTF bias, liquidity requirement and a session window together tells you nothing about which of the three helped. Change one, collect at least the minimum sample, write the average R down, then change the next.

5. Statistics — how the scoreboard behaves

Option	Default	What it does & how to tune it
Min sample for % and Avg R	10	Percentages and average R stay hidden until this many cycles have closed. A win rate built on four trades is noise dressed as insight. Raising this to 25 or 30 is honest; lowering it below 10 is not recommended.
Persist statistics to disk	On	FIX 4. Outcomes are written to <code>bin\Custom\TraderSuite\stats</code> , keyed by cycle close time, so the sample survives reloads, template changes and NinjaTrader restarts. Turn it off only if you want a session-local experiment that leaves no trace.

What the file is called

`po3_NQ_09-26_5Min_a1b2c3d4.csv`

Product, instrument, timeframe, and the eight-character rule-set fingerprint. Change any engine setting and the fingerprint changes with it, so a new file starts and the old sample stays intact and reusable.

What is inside it

One line per closed cycle: close time, direction, outcome code, R multiple, bars held, and whether it followed a re-arm. Plain CSV — open it in Excel, pivot it, plot the equity curve of R, do whatever you like with it. The indicator only ever appends and de-duplicates by close time.

Deleting a stats file is permanent

There is no undo and no cloud copy. If you want to keep a baseline before an experiment, copy the file somewhere else first — although in practice the rule-set fingerprint already does this for you, because the experiment writes to a different filename.

Settings — appearance, alerts & the NinjaTrader tabs

6. Appearance

Option	Default	What it does
Accumulation colour	 #5B7FA6	The range box and its boundaries while the range is live.
Manipulation colour	 #E2963A	The sweep wick highlight and the M marker.
Bullish colour	 #00E39A	Long distributions, targets, positive outcome tags.
Bearish colour	 #FF3B6B	Short distributions, stops, negative outcome tags.
Entry line colour	 #9BA4B0	Deliberately neutral so the entry never competes with stop and target for attention.
Show boxes	On	The accumulation and distribution rectangles. Off leaves lines and markers only — useful on a chart already carrying other zone tools.
Show markers	On	The M, D and outcome tags on the bars that produced them.
Fade closed cycles	On	History dims so the live cycle is unmistakable at a glance.
Show last N cycles bright	4	How much recent history stays at full opacity. Raise it when reviewing, lower it when trading.

7. Alerts

Option	Default	What it does
Enable Alerts	On	Three events reach the NinjaTrader Alerts window with sound: sweep detected , distribution opened (with direction and R:R), and cycle closed (with outcome and R). Real-time bars only, fired on the close, so nothing alerts twice on the same bar and nothing fires while a historical chart loads.

NinjaTrader's own tabs

Data Series

Calculate is fixed at **On bar close** and should stay there. Every threshold in the engine is a percentile of closed-bar data, and a tick-by-tick evaluation would flicker phases as intrabar price wandered across a boundary.

Days to load matters more here than in most indicators: the distribution window plus the maximum range age is your minimum, and history is what builds the scoreboard. 30–60 days on a 5-minute chart is a reasonable starting point.

Visual

The plots are transparent by design. Everything you see is drawn in `OnRender` with `Direct2D`, which is why the boxes, gradients and panel look the way they do rather than like standard NinjaScript output. Changing plot colours in this tab does nothing visible.

Auto scale is off so a distant fib target cannot compress your price axis.

Public series for strategy builders

Nine values are exposed for use from a strategy or another indicator: `PhaseSeries`, `CycleDir`, `RangeHighSeries`, `RangeLowSeries`, `ModelEntry`, `ModelStop`, `ModelTarget`, `ModelRewardRisk` and `ClosedCycleCount`. They are written on bar close on the primary series, so a strategy can read them the moment a distribution opens without duplicating any of the detection logic.

Choosing your timeframe

The engine's timing constants are written in bars, not minutes. Change the timeframe and you change what those constants mean in clock time — which is why the same settings behave completely differently on a 1-minute and a 30-minute chart.

The short answer

5-minute is the sweet spot. 15-minute is the practical upper bound and needs one setting changed. 30-minute and above: don't.

What the defaults mean at each timeframe

Chart	Range must mature	Range expires	Distribution timeout	Cycles / day	Time to 100 closed cycles
1 min	12 min	1.6 h	64 min	~12	~8 days
5 min	1 h	8 h	5.3 h	~2.5	~6 weeks
15 min	3 h	24 h	16 h	~0.8	~4 months
30 min	6 h	2 days	32 h	~0.4	~8 months
1 hour	12 h	4 days	2.7 days	~0.2	over a year

Cycle rates are order-of-magnitude figures for a liquid index future at default settings; your market will differ. What does not differ is the shape of the table — sample accrual falls off a cliff above 15 minutes.

Why 5 minutes wins

The whole point of this indicator is the scoreboard, and a scoreboard needs samples. Five minutes gives you a readable statistic in about six weeks while keeping the range boundaries structural rather than microstructural. A one-hour range is a genuine accumulation; so is a five-minute one. A one-hour *sample size* is not a sample.

It also lines up with how the AMD idea is normally traded intraday: a range built over the first hour of a session, swept, and distributed inside the same day.

Why above 15 minutes fails

Not because the pattern stops working — because you will never know whether it does. At 30 minutes you reach the minimum sample around month two and a statistically meaningful 100 cycles some time next year, by which point the market regime that produced them has changed twice. The 64-bar distribution timeout also becomes absurd: on a 1-hour chart it holds a model open for nearly three days, which is not a distribution, it is a swing position with a stop from last Tuesday.

Per-timeframe adjustments

Chart	Change	Why
1 min	Min range width % → 0.05 Distribution window → 300 Min risk (ticks) → 12 Min reward:risk → 2.0	One-minute ranges are genuinely narrow, so the 0.15% floor rejects almost everything. The wider percentile window stabilises thresholds against minute-to-minute noise, the higher tick floor keeps spread out of the R calculation, and the stricter R:R gate compensates for the lower quality of one-minute structure.
5 min	Defaults	Everything is calibrated here.
15 min	Distribution timeout → 26	Keeps the model's holding period at roughly six and a half hours instead of sixteen — one session, not two. Nothing else needs touching.
30 min +	Not recommended. If you must, at least set Distribution timeout to 14 and accept that the scoreboard will stay hidden for months.	

Starting points by market

These are calibrated starting points, not optimised settings — the whole design of the indicator is that *you* find out what works on your market by collecting a sample. Load one, let it run, read the average R.

Index futures — ES, NQ, YM, RTY

5-minute chart, defaults, with two changes:

Min range width %	0.15
Min risk (ticks)	8
Use session window	On, 09:30–16:00

Index futures are what the defaults were built around. The session window matters here more than anywhere else: overnight ranges on ES sweep constantly and distribute rarely, and including them drags the average R down without teaching you anything about the session you actually trade.

FX majors — EURUSD, GBPUSD, USDJPY

5- or 15-minute chart:

Min range width %	0.04
Min risk (ticks)	15
Compression threshold	20
Use session window	On – London or NY only

FX percentage ranges are an order of magnitude tighter than equity index ranges, so the width floor must come down hard or nothing ever qualifies. The tighter compression threshold compensates by demanding better coils. Asian-session ranges swept at the London open are the textbook case this indicator is looking for.

Crypto — BTC, ETH perpetuals

5-minute chart:

Min range width %	0.30
Soft depth cap	88
Stop buffer past sweep	0.6 ATR
Use session window	Off (24/7)

Crypto sweeps run deeper and faster, so the stop buffer needs room and the depth cap earns its keep by throwing out the genuine breakouts. Expect a lower target rate and a higher average R per winner than on index futures — the payoff distribution has a longer tail in both directions.

Energy & metals — CL, GC

5-minute chart:

Min range width %	0.20
Max bars until return	10
Min reward:risk	1.8
Use session window	On, pit hours

Crude in particular produces slow return legs, so the sweep needs a couple more bars to be recognised. The stricter R:R gate is there because a sluggish return means the confirming close often lands mid-range, exactly the geometry the gate exists to refuse.

Every one of these presets is a different rule-set id

Switching from the FX preset to the crypto preset does not merge two samples into a confused average — it opens a separate scoreboard. You can move between presets freely and come back to a preserved history.

Recommended use

This indicator is not a signal service. It is an instrument for finding out whether a specific, precisely defined pattern has an edge on *your* market, at *your* settings — and then for trading it once you know. The recommended workflow follows that order.

- 1 **Install and let it sit for a week.** Chart, defaults, 5-minute, no trades. The scoreboard hides percentages until ten cycles have closed for a reason: anything you conclude before then is noise. On a liquid index future you will pass ten in roughly four days.
- 2 **Read the average R, not the win rate.** With the default 1.5 fib extension and the R:R gate at 1.5, target rate near 35% is break-even. Forty percent with an average R of +0.17 is a good market. Fifty-five percent with an average R of -0.05 means your winners are small and your losers are full-size, and the pattern is not paying you.
- 3 **Give it a real sample before you believe it.** At 100 closed cycles the standard error on the target rate is about ± 5 percentage points. At 30 it is about ± 9 — wide enough that a 42% reading and a 33% reading are the same reading. Six weeks on a 5-minute chart gets you to 100.
- 4 **Change one setting at a time.** Each change opens a new scoreboard under a new rule-set id, so experiments never contaminate each other and old samples stay retrievable. Run each variant to the minimum sample, note its average R, move on. Compression threshold, min reward:risk and fib extension are the three worth testing first.
- 5 **Only then trade it — and trade the model, not the idea.** The entry, stop and target on the chart are the exact geometry the scoreboard was built from. Take a discretionary entry three bars later at a better price and you are trading something the statistics have never measured.

How to actually use it in a session

Watch for the accumulation box to appear and the compression figure to go green. When a sweep prints, do nothing — the sweep is not the trade. When the box turns into a distribution with an R:R above your gate, that is the model opening, and the three lines are your entry, stop and target. Then leave it alone until one of them is touched or the timeout closes it. The single most common way to lose money with this tool is to manage the position by feel and then keep score as though you had followed the model.

Three ways to use it, in order of how much work they ask of you

As a research instrument

Run it across several instruments and settings for a month or two, and read off which markets pay for this pattern. This is what it is best at and costs you nothing but patience.

As a discretionary framework

Trade the model as drawn, sized at a fixed fraction of account per R, with your own filter for news and session. The scoreboard tells you honestly whether your filter is helping.

As a strategy component

Read the nine public series from your own NinjaScript strategy and automate the execution. The detection logic never has to be rewritten.

What would make us stop trading it

An average R below zero across 100+ cycles at settings you have not changed. A breakout count that dwarfs the cycle count — meaning this market leaves its ranges rather than reverting into them. Or a form strip that has been solid $\nearrow$ for a fortnight while the all-time average is still positive: the regime has changed and the all-time figure is describing a market that no longer exists.

Reading the numbers

Four readings you will actually see on the panel, and what each of them is telling you to do.

Healthy — trade it

Cycles 118 · Target 46 (39.0%) · Stop 58 (49.2%) · Timeout 14 (11.9%) · Avg R +0.171R · R:R rejected 22

Target rate comfortably above the ~35% break-even, a positive expectancy of about a sixth of a risk unit per cycle, a sample large enough that the figure means something, and a rejection count showing the R:R gate is working without strangling the sample. This is what a market that pays for this pattern looks like.

Losing — change something or stop

Cycles 96 · Target 28 (29.2%) · Stop 61 (63.5%) · Timeout 7 (7.3%) · Avg R -0.084R · Breakout / Fail 74 / 3

Below break-even on a sample large enough to trust, and the breakout count is telling you why: 74 sweeps never returned. This market breaks out of its ranges rather than reverting into them, so the premise of the pattern does not hold here. No setting fixes that. Move to another instrument, or test a much stricter compression threshold and see whether the tighter coils behave differently.

Over-filtered — loosen the gate

Cycles 19 · Target 9 (47.4%) · Avg R +0.28R · R:R rejected 141

The target rate and average R look excellent, and they are almost meaningless — nineteen cycles carries a standard error of about ± 11 points. Meanwhile 141 confirmations were refused. The gate is set so high that you will wait months for a readable sample. Drop **Min reward:risk** toward 1.2, or the **Fib extension** toward 1.272, and let some volume through.

Too loose — tighten it

Cycles 402 · Target 121 (30.1%) · Timeout 96 (23.9%) · Avg R +0.012R · R:R rejected 4

A huge sample scraping along at zero, and nearly a quarter of cycles closing on timeout — the model is opening on setups that never had the momentum to reach anything. Lower **Compression threshold** toward 15 so only genuine coils qualify, and raise **Min reward:risk**. You will trade a third as often for a materially better average.

The two diagnostics people misread

Reading	What it actually means
High R:R rejected	Not a fault. It is the count of setups that looked identical on the chart but had poor geometry. A ratio of roughly one rejection per two accepted cycles is normal and healthy. Ten rejections per accepted cycle means the gate is too strict for this market.
High Breakout	Not a missed opportunity. These are sweeps that never returned inside the range, so no reversion trade existed. A high count is a statement about the instrument's character, not about the indicator's performance — and it is a reason to look elsewhere for this pattern.
High ds	Many cycles followed a re-arm, i.e. the market swept one side, was refused, then swept the other. This market habitually raids both directions before committing. Useful context for anyone trading the first sweep discretionarily.
Timeout share above ~20%	Models are opening without enough follow-through to resolve. Either the timeout is too short for the timeframe, or the entries are being taken on setups with no momentum behind them.

Licensing & activation

The indicator validates once, then caches the result for 24 hours and keeps working offline within a 7-day grace window — a brief internet drop never interrupts your session, and never costs you a cycle from the scoreboard.

Getting started

- 1 Add **TS_AMD_Po3_Pro** to your chart (Indicators dialog → **TraderSuite** folder).
- 2 Open its settings and find the **0A. License** group.
- 3 Paste your **License Key** — or leave it empty to start an automatic 7-day trial.
- 4 Click OK. The panel's **Status** line confirms your state.

License states you may see

Panel shows	Meaning
✓ LICENSED	Valid product license, activated on this machine.
✓ MEMBER	Active membership license — this indicator is free for Pro Members.
✓ ADMIN	Administrator master key (support / testing).
TRIAL (Xd left)	Free trial, days remaining shown on the panel.
✓ OFFLINE MODE	Server unreachable; running on the offline grace period.
✗ message	Attention needed — invalid key, max activations, expired, or internet required. The full reason is spelled out in the message and in the license overlay on the chart.

Activations

A product license activates on a limited number of machines (shown as **used / max** on the panel). Need to move machines? Deactivate the old one from your account dashboard, then activate the new one.

Your key is never displayed

The License Key field is masked, and the key is deliberately kept out of the chart label and the Control Center columns — so a screenshot or a screen share can never leak it.

Your statistics belong to you

The scoreboard files live on your machine under `bin\Custom\TraderSuite\stats`. Nothing about your cycles, outcomes or settings is transmitted anywhere. Licence validation sends a hashed machine identifier and the product code — nothing else.

What this indicator is not

Every tool has a boundary. Knowing where this one ends is what keeps you from asking it for something it was never built to give.

It is not a backtest

The scoreboard measures a *reference model*: enter at the confirming close, one stop, one target, no management, no slippage, no commission, no position sizing. A real account trading the same cycles will not produce the same numbers. What the scoreboard tells you is whether the underlying pattern has directional edge — not what your P&L would have been.

It is not a trade manager

Nothing here moves a stop to break-even, scales out, trails, or reacts to news. If you do any of those things while trading it, you are trading a different system from the one being scored, and the scoreboard stops describing your results.

It is not a claim about the future

A positive average R over 200 past cycles is evidence, not a promise. Markets change regime; the form strip exists precisely so you can see the recent run separately from the all-time figure and notice when they disagree.

It is not repainting — and here is why

Everything evaluates on bar close. Pivots carry a deliberate three-bar confirmation lag, and the pivot boundaries are clamped so a boundary can never be drawn inside price that has already printed. A cycle drawn three days ago looks exactly as it did when it printed.

On the AMD idea itself

Accumulation, manipulation and distribution is a description of one way markets behave some of the time — not a law. Ranges break out as often as they revert, and this indicator counts both: the breakout tally on the panel is there so you can see how often the premise failed on your market rather than only seeing the times it held.

What it *is*, in one sentence

A precisely specified, non-repainting detector for one market pattern, attached to an honest, persistent, rule-set-aware record of how that pattern has actually resolved on your instrument — so that your opinion about it can eventually be replaced by a measurement.

Troubleshooting & FAQ

► **Nothing draws at all.**

Check the panel first. WARMUP means the distribution window has not filled — you need at least `Distribution window + Max range` bars, so about 300 on defaults. A red licence status hides the visuals by design. If the panel is healthy and idle, this market simply has not compressed enough to anchor a range yet.

► **Ranges appear but never become cycles.**

The most likely cause is `Min range width %` — 0.15 is calibrated for index futures and rejects nearly every FX and fast-chart range. See Sections 16 and 17. After that, check the breakout tally: if it is climbing, sweeps are happening but price is not returning within `Max bars until return`.

► **My statistics reset when I changed a setting.**

They did not reset — they moved. Any engine-affecting setting changes the rule-set fingerprint and opens a new scoreboard, because a win rate that mixes two rule sets is worse than no win rate at all. Set the option back and the old sample reappears exactly as it was.

► **Percentages are hidden even though cycles are showing.**

`Min sample for % and Avg R` is 10 by default. The cycle counter still moves; the derived figures appear once the sample is large enough to be worth reading.

► **The indicator reloaded itself when I enabled HTF bias.**

Expected. Enabling the bias filter attaches a secondary data series, which forces NinjaTrader to reload the indicator with the extra data. It also means you need enough history loaded to cover the higher timeframe's EMA.

► **Can I run it on Renko, Range or tick bars?**

It will run, but the timing constants — maturity, expiry, timeout — are all counted in bars, and on a non-time-based chart a bar is not a unit of time. The percentile logic still works; the interpretation of "64 bars" does not. Time-based charts are strongly recommended.

► **Can I use it on more than one chart at once?**

Yes. Each instrument and timeframe writes to its own scoreboard file, so an NQ 5-minute chart and an ES 5-minute chart accumulate independent samples. Two charts on the same instrument, timeframe and settings share one file and de-duplicate by cycle close time, so nothing is double-counted.

► **Where exactly are my statistics stored?**

`Documents\NinjaTrader 8\bin\Custom\TraderSuite\stats\` — one CSV per instrument, timeframe and rule set, capped at 4,000 rows with the oldest trimmed first.

Changelog

What changed in each release. Update through your account download page — import the new file over the old version (delete the old file first if NinjaTrader reports a duplicate).

Version	Notes
v1.0.0	Current build — initial release. Donchian percentile compression detection with impulse-tail trimming · confirmed-pivot range boundaries clamped to the unconfirmable tail · return-validated liquidity sweep with optional equal-highs/lows requirement, soft depth cap and double-sweep re-arm · fib-extension reference model with ATR stop buffer · minimum reward:risk gate and minimum risk floor · R-multiple outcome scoring with conservative same-bar resolution · persistent, rule-set-fingerprinted walk-forward scoreboard · TraderSuite statistics panel with form strip and diagnostics · three alert events · nine public series for strategy use · License System v3.0 with trial, membership and offline grace.

Improvements over the original concept

Fix	What it corrects
FIX 1	Pivot boundaries are clamped to the unconfirmable tail, so a boundary can never sit inside price that has already printed.
FIX 2	A minimum reward:risk gate, because the confirming close's position inside the range is uncontrolled and produced an R:R scatter from roughly 0.4 to 3.7 on visually identical setups.
FIX 3	A minimum risk floor in ticks, so a degenerate setup cannot register a double-digit +R outlier that flatters the average.
FIX 4	Statistics persist to disk, keyed by cycle close time and fingerprinted by rule set, so the sample survives reloads and experiments never contaminate each other.
Also	The impulse-tail trim no longer goes silently inert at higher maturity settings, the range-start index is off-by-one corrected, and all rendering reads immutable snapshots so the UI thread can never fault on data the engine is mid-update.

Staying current

License holders get every update free. The panel shows your installed version next to the product name; compare it with the version on the product page whenever you like. Note that an engine change in a future version will change the rule-set fingerprint by design — your previous scoreboard is preserved, not overwritten.



TraderSuite

Thank you for trading with AMD Po3 Pro. For the latest build, video walkthroughs and support, visit us online.

Web tradersuite.online

Email completetradersuite@gmail.com