NINJA CODE SOLUTIONS · NINJATRADER 8

Market Structure Copilot

Reads market structure across six degrees of swing, states a conditional thesis with the levels that would confirm or break it, and packages the whole picture for an AI second opinion.

`Indicator + GPT workflow` `Version 2.0.0` `Futures · index CFDs`

What it does

AND WHAT IT DELIBERATELY DOESN'T

The Copilot runs six ZigZag instances at different retracement thresholds over one chart. Each is a **degree of structure** — the same price action measured at a different swing size — not a separate timeframe. From those it derives a bias, a market phase, and the specific prices that would confirm or invalidate the current read.

It then writes that entire picture — structure, zones, session levels, swing volume, instrument specs — into a single snapshot you paste into a purpose-built GPT for a plain-English trade plan.

> THE INDICATOR DOES NOT
>
> Generate entries, size positions, or place orders. Every output is conditional: *if price does X at level Y, the read becomes Z.* The decision stays yours.

Quick start

ORDER MATTERS HERE

01 Add the indicator

Apply `Market Structure Copilot` to a chart. A floating panel opens alongside it, titled with the instrument.

02 Turn on higher-timeframe evidence

Set `Enable Higher Timeframe Evidence` to **true**. It ships off. Left off, the Copilot has no genuine higher timeframe and says so on every read.

03
Load enough history

The track record is rebuilt from whatever bars the chart holds. A few days gives thin numbers; several weeks gives useful ones.

04
Set up the GPT once

Open the Market Structure Copilot GPT `(https://chatgpt.com/g/g-68519540b4588191a123bac70e33e1f0-ninja-code-solutions-market-structure-copilot)`, or build your own from the supplied instructions and schema reference.

05
Copy and paste

Press **Copy Snapshot** at the bottom of the panel, paste into the GPT, send. You get a trade plan or a reasoned no-trade.

Chart overlay

WHAT GETS DRAWN ON PRICE

The badge

A compact block in one corner carrying the whole read: bias, phase, confidence, then the three decision prices. Move it with `Overlay Badge Corner`.

HEADER	BEARISH | TREND UNDER PRESSURE | LOW 44
PRESSURE	Selling pressure easing
LEVELS	C 7721.75 X 7739.75 NEXT 7704.50

Bias colours the badge: `Bullish` `Bearish` `Neutral`. If confidence reads `(FLOOR)`, the evidence produced no usable edge at all — treat it as a hard warning rather than a weak signal.

Decision lines

Three horizontal lines mark the prices that matter. Their names change with the read, which is the point:

DIRECTIONAL BIAS	NEUTRAL BIAS	MEANING
CONFIRM	UPPER	Acceptance here confirms the thesis, or breaks balance upward

DIRECTIONAL BIAS	NEUTRAL BIAS	MEANING
INVALID	LOWER	Acceptance here breaks the thesis, or breaks balance downward
NEXT	BEYOND	The next structural level past confirmation

WHY THE NAMES CHANGE

When bias is neutral there is no thesis to confirm or invalidate — only a range with two edges. Labelling those edges CONFIRM and INVALID would imply a long position that the evidence never supported, so they are drawn in grey as boundaries instead.

Zones and session levels

The nearest relevant supply or demand zone is shaded as a band labelled ACTIVE SUPPORT or ACTIVE RESISTANCE. Turning on `Show Detailed References` adds session landmarks — PDH, PDL, ONH, ONL, ORH, ORL — plus the raw structural break references.

The panel

TOP TO BOTTOM

Current read

A headline and a short narrative, then four metrics:

METRIC	WHAT IT MEANS
Bias	Weighted vote across higher structure, execution structure, EMA stack, structural events, bar pressure and the higher timeframe
Phase	Balance, breakout attempt, initiative expansion, trend continuation, pullback, trend under pressure, or structural transition
Pressure	Who has controlled the last few bars, and whether that control is strengthening or easing
Confidence	Structure agreement blended with setup quality, 15–95

What changed

Only material transitions appear, ranked so a bias change always outranks a wording change in pressure. A read that flips repeatedly in a short window is telling you it is unstable.

Scenarios and levels

Two conditional paths — continuation and failure — followed by the three decision levels. Each level shows its price, what it is, its distance in ticks and in ATR, and whether price has already **tested**, **rejected** or **accepted** it.

Why this read / Cautions

The evidence behind the call, and anything arguing against it. Cautions are worth more attention than the headline: they name the specific weakness.

Advanced

Five tabs. **Thesis** holds session phase, location and the change log. **Levels** lists every reference and the nearest zones. **Evidence** carries structure alignment, EMA and volume context, the recent sequence and raw details. **Confidence** shows the score's components and every adjustment applied — the audit trail for the number. **Logs** covers saved snapshots.

AI workflow

THE REASON THE SNAPSHOT EXISTS

Copy Snapshot puts a readable brief plus a complete JSON payload on the clipboard. Paste the whole thing into the GPT.

The GPT is instructed to form its own read from the raw evidence *before* looking at the indicator's conclusion, then say whether it agrees. That disagreement is informative — it is a second opinion, not an echo.

What comes back

BOTTOM LINE	Two sentences, no jargon
DECISION	TRADE PLAN or NO TRADE
SETUP TYPE	Continuation / Mean reversion / Breakout
CONFIDENCE	0–100, under 50 is no trade
TRACK RECORD	Expectancy for this phase, this session
TARGETS & RISK	Named levels with reward-to-risk
INDICATOR CHECK	Agree or disagree, and why

On a no-trade it drops targets and stop entirely and gives you **What Would Change This** instead — the conditions to watch, each with a named level.

Rules the GPT enforces

- A stop must be anchored to a level that exists in the snapshot — never a price invented to flatter the ratio
- Risk must clear both 0.25× ATR and an eight-tick floor
- Reward-to-risk is stated every time, measured from the intended entry
- Breakouts need above-average participation; mean reversion needs evidence that momentum is fading
- Outside regular hours, thin volume is normal and cannot be read as exhaustion

> BUILDING YOUR OWN GPT
>
> Paste `GPT-Instructions.md` into the Instructions box — it is tuned to sit just under ChatGPT's 8,000-character limit — and upload `MarketStructureCopilot-Schema-v2.md` as a Knowledge file. The schema file documents every field and is what the GPT consults for anything not in the instructions.

Track record

EXPECTANCY, NOT HIT RATE

Every time the structure changes, the Copilot records the thesis and its two levels, then watches forward to see which one price reaches first. Confirmation first is a win paying its reward-to-risk; invalidation first loses exactly 1R.

The headline number is **expectancy in R** — average gain per thesis:

RTH	+0.18R across 205 resolved · 81% hit
RTH BALANCE	+0.46R over 66 resolved
ETH	−0.11R across 182 resolved

DO NOT QUOTE THE HIT RATE

A hit rate mostly reflects how close the confirmation level sits relative to the invalidation level. Four wins at 0.2R and one loss is an 80% hit rate that loses money. Expectancy is the number that answers whether a read is worth taking.

Regular and extended hours are tracked separately because overnight liquidity behaves differently enough that a blended average means nothing. The GPT uses only the bucket matching the current session, and lowers its confidence when the phase in play has a weak record.

Nothing is stored on disk. The record is rebuilt from the loaded chart history every time, so it always matches the chart in front of you and can never drift out of date.

Settings

GROUPED AS THEY APPEAR IN THE PROPERTIES DIALOG

Track Record

SETTING	DEFAULT	NOTES
`Enable Track Record`	`true`	Turn off to remove all outcome scoring and its load-time cost

Parameters

SETTING	DEFAULT	NOTES
`Retracements`	`0.25, 0.5, 1, 2, 3, 5`	The six structural degrees. Largest is primary context, smallest is timing
`Equal Threshold %`	`0.03`	How close two swings must be to count as equal
`Context Lookback Bars`	`50`	Window for range location, zone touches and reactions
`Fast EMA / Slow EMA`	`9 / 21`	Drives the EMA stack vote
`ATR Period`	`14`	Every distance is judged against this
`Volume Avg`	`20`	Baseline for participation
`Enable Higher Timeframe Evidence`	`false`	**Turn this on.** Without it there is no true higher timeframe
`Higher Timeframe Minutes`	`60`	Must exceed the chart's own period

Visual

SETTING	DEFAULT	NOTES
`Show Copilot Overlay`	`true`	Badge, decision lines and zone bands
`Overlay Badge Corner`	`TopLeft`	Move it off your price action
`Show Detailed References`	`false`	Adds session landmarks and break references
`Ultra Compact Mode`	`false`	Smaller panel for tight screens
`Show Filled FVGs`	`false`	Keeps spent gaps in the zone list
`Clipboard Popups`	`true`	Confirmation toast on copy
`Enable Evidence Logging`	`false`	Writes snapshots to disk for later review

Evidence Tuning

SETTING	DEFAULT	NOTES
`Structure Score Threshold`	18	Bias score needed to call a direction. Lower is more decisive, and noisier
`Zone Candidates To Score`	12	How many zones survive to the final ranking

Snapshot & Performance

SETTING	DEFAULT	NOTES
`Snapshot Swing Count`	6	Swings included per degree in the export
`Snapshot Event Count`	10	Structural events included per degree
`History-Heavy Export Mode`	false	Adds a full event tape and swing ladders. Much larger paste
`Historical UI Lookback`	250	Bars the panel will refresh over during historical processing
`Update UI During Historical`	false	Leave off unless debugging
`Max Tracked Zones`	200	Cap on stored gaps and order blocks
`Historical Zone Retention`	120	Bars a spent zone survives before removal

Advanced Tuning

These shape how the engine reasons. The defaults are calibrated for index futures on intraday charts; change them only with a reason.

SETTING	DEFAULT	NOTES
`Bias Hysteresis Ratio`	0.65	A bias holds until the score falls to this fraction of the threshold. Set to 1.0 to disable
`Pressure Lookback Bars`	3	Horizon for reading bar pressure
`Pressure ATR Threshold`	0.18	Move size, in ATR, that counts as directional
`Low Participation Ratio`	0.9	Below this is weak volume
`High Participation Ratio`	1.15	Above this is strong volume

SETTING	DEFAULT	NOTES
`Min Level Distance (ATR)`	`0.25`	Levels nearer than this are treated as noise
`Acceptance Lookback Bars`	`20`	Window for judging tested / rejected / accepted
`Acceptance Closes Required`	`2`	Consecutive closes beyond a level to count as acceptance
`Outcome Horizon Bars`	`60`	Bars a thesis has to resolve before it expires unscored

Replay & review

STUDYING PAST READS

Five controls at the panel's foot step the snapshot backwards through history: **Back** and **Forward** move one bar, **Prev Event** and **Next Event** jump between structural events, and **Live** returns to the current bar. A gold vertical line marks the replayed bar on the chart.

Replay is read-only. Transition tracking pauses, and nothing recorded in replay enters the track record.

With `Enable Evidence Logging` on, snapshots are written to `Documents\NinjaTrader 8\log\MarketStructure200`, each paired with an outcome template you can fill in afterwards. The **Review tools** expander copies the latest log, the log index, or an outcome template, and can jump the replay straight to the last logged bar.

Troubleshooting

COMMON SITUATIONS

SYMPTOM	CAUSE AND FIX
"No true higher timeframe is available"	Set `Enable Higher Timeframe Evidence` to true. It defaults off, and must exceed the chart's period
Track record says "no resolved theses"	Not enough loaded history for any thesis to resolve. Each needs up to 60 bars. Load more data
Slow chart load	Outcome scoring runs across all loaded bars. Set `Enable Track Record` to false to measure the difference, or load fewer days
GPT says fields are missing	The snapshot came from an older build, or was pasted in part. Recompile, reload the chart, and copy again

SYMPTOM	CAUSE AND FIX
Confidence reads “(floor)”	Working as designed. The evidence supports no usable edge; the number is pinned at its minimum
Bias flips back and forth	The score is sitting near the threshold. Lower `Bias Hysteresis Ratio` to make an established bias stickier
Panel closed by accident	Remove and re-apply the indicator; the panel opens with it

Glossary

TERMS USED THROUGHOUT

TERM	MEANING
`Structural degree`	The same chart measured at a different swing size. Not a separate timeframe
`BOS`	Break of structure. Price passed a prior swing point, confirming the trend continues
`CHoCH`	Change of character. The first break the other way — an early warning, not yet a reversal
`HH / HL / LH / LL`	Higher high, higher low, lower high, lower low. HH with HL is an uptrend; LH with LL is a downtrend
`FVG`	Fair value gap. An imbalance left behind by a fast move
`Order block`	The last opposing candle before an impulsive move
`Premium / Discount`	Upper or lower part of a range. Premium favours selling, discount favours buying
`Liquidity sweep`	Price ran past a high or low to trigger stops, then reversed
`Acceptance`	Price held beyond a level for consecutive closes, rather than poking through and returning
`Invalidation`	The price that proves the read wrong
`R`	One unit of risk. A 2R win pays twice what the stop risked
`Expectancy`	Average R per thesis. Positive means the read made money historically
`RTH / ETH`	Regular trading hours, 09:30–16:15, versus everything else