



The Second Dimension of the Opening Range

Why the most popular setup in retail trading is missing half its information — and what a second, structure-based opening range stacked on top of it reveals

DAVID LEE NELSON · FOUNDER, TRADEADAPT

JUNE 2026 · TRADEADAPT.AI

ABSTRACT

The Opening Range Breakout (ORB) is arguably the most widely traded setup in retail futures. It is simple, it is visual, and it captures something real about how markets re-price at the start of a session. It is also, in the form most traders run it, incomplete in two specific ways: it measures the market in only one dimension, and it measures it at only one moment of the day.

This paper proposes an extension rather than a replacement. The classic price-based opening range stays exactly where it is. Alongside it, we compute a **second opening range on structure-driven bars** — bars that print only when the market's structural regime actually changes — over the same window. The two ranges almost never coincide, and that is the point: stacked together, they convert a single breakout line into a **layered map of expected reaction areas** — pullback landing zones, support and resistance bands where price and structure agree, and offset gaps that tell a trader where a breakout is likely to retrace before it continues. We then extend the framework in time, treating every major liquidity handoff of the 24-hour day — Asia, London, the 8:00 ET data window, the NY open, and the 10:00 and 14:00 institutional windows — as its own opening range event.

The result is an ORB that knows not just *where* price broke, but what the market's structure was doing when it broke, and where it is likely to be defended. Nothing in this paper requires abandoning the setup you already trade. It asks a narrower question: if the opening range is worth trading, why read it in one dimension?

SECTION 1 — THE SETUP

The Most Crowded Trade in Retail — and Its Two Blind Spots

Every ORB variant works the same way at its core: define a window after the open, mark the high and low of that window, and treat a break of those levels as the day declaring its intentions. The logic is sound. Session opens concentrate order flow, and the first contested range often becomes the day's reference frame.

But look at how the standard implementation is built, and two assumptions surface that almost nobody questions.

Blind spot one: the chart underneath measures time, not market. A 5-minute ORB chart prints a bar every five minutes whether the market moved forty points or four ticks. During the violent first minutes of the NY open, one bar can swallow three distinct structural events; during the lunch drift, twelve consecutive bars can contain no structural information at all. The trader is asked to judge a *price* event — the break of a level — using bars whose boundaries were chosen by a clock. And there is a second cost hiding inside the first: the price range itself is built from extremes that include every wick, every stop-run, every burst of opening noise. The levels everyone watches are, by construction, the levels most contaminated by the moments when the market was least orderly — and because everyone watches them, they are also the levels most likely to be probed, faded, and trapped against.

Blind spot two: the market opens more than once. The 9:30 ET equity open is one liquidity event among several. The Asian session establishes an overnight range. London's open re-prices it. The 8:00 ET hour brings the data window that routinely sets the day's true high or low before New York is awake. After the open itself, the 10:00 ET window — the first post-open releases, the initial-balance fight, and the boundary of the institutional 4-hour bar — produces the morning's second auction, and the 14:00 ET window, opening the afternoon 4-hour bar, hosts the repositioning flows that decide how the day closes. A trader running a single 9:30 ORB is using a reference frame that may be obsolete by 10:15 — and was arguably incomplete at 9:29.

The rest of this paper addresses these blind spots in order. The first requires a second kind of bar, and a second range built from it. The second requires changing when a range begins.

What If You Could See Structure?

StructureBars are TradeAdapt's custom bar type, built on a simple contract: *a new bar prints when the market's structural regime changes, and not before.*

Under the hood, every incoming tick feeds a volatility-adaptive engine stack that continuously maintains a consensus view of the market's current structural zone — the region of prices the market is treating as "here," bounded by the levels it would have to take out to be "somewhere else." Multiple independent engines, each adapting to volatility on its own horizon, must agree before the zone is considered to have moved. When the consensus shifts, a bar prints. When it doesn't, the chart waits.

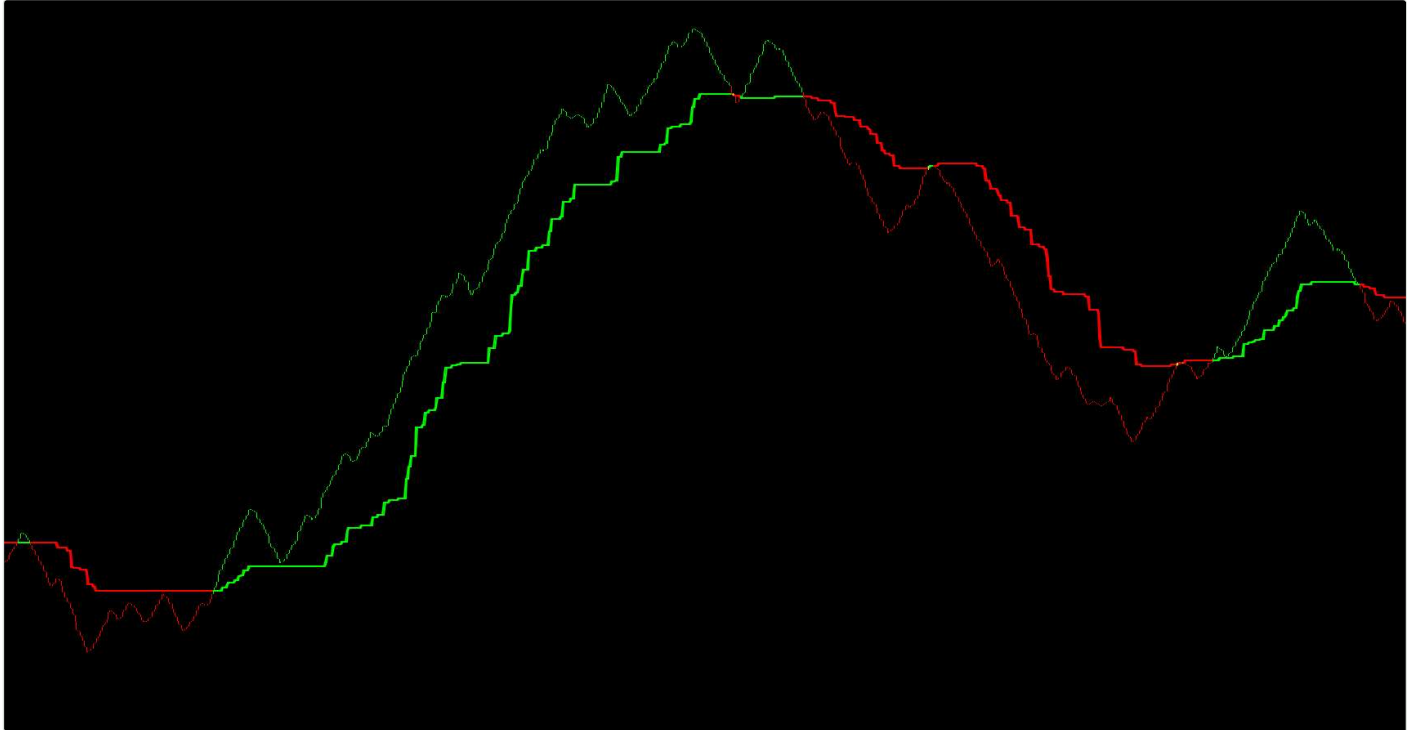


FIGURE 1 StructureBars on NQ futures. In trend the bars stride — each one a confirmed structural step. In chop they compress, because structurally, almost nothing is happening. Every bar carries the same amount of information, whether it formed in nine seconds or nine minutes.

The practical consequences are immediately visible. In trend, StructureBars stride: each bar is a confirmed structural step, and a directional move renders as a clean staircase rather than a smear of mixed candles. In chop, StructureBars compress: forty-five minutes that fill a time chart with alternating noise may produce only a handful of structure bars, because structurally, almost nothing happened. The chart's information density becomes constant — every bar means the same amount.

Because the engine stack adapts to volatility rather than being parameterized to a fixed increment, the same bar type behaves consistently across instruments and regimes. An NQ chart on a CPI morning and a YM chart in a summer drift render structure at the same semantic resolution. The trader stops re-tuning the chart to match conditions; the chart re-tunes itself.

This is the substrate everything else in this paper stands on. A range computed on these bars inherits their property: it marks where structure was contested, not where noise happened to wick.

What If You Applied a System to Structure?

Once bars carry constant structural meaning, systematic logic on top of them becomes dramatically simpler — and more honest.

TradeAdapt's signal layer maintains a continuous directional bias derived from the same structural-zone consensus that drives the bars: price holding above the zone is a long regime, price holding below is a short regime, and price *inside* the zone is explicitly indeterminate — a state the system names rather than hides. Transitions between regimes generate the flip points a discretionary trader sees as arrows and a systematic engine consumes as orders. There are no oscillators tuned to a lookback, no pattern libraries, no per-instrument curve fitting. The system's opinion is a direct restatement of structure.

A note on backtests — and why this paper doesn't lead with one

Run a naive backtest of a flip system on structure-driven bars and the platform will report results that look spectacular — double-digit profit factors, win rates no live account has ever sustained. We know, because we have run them. We are deliberately not publishing those numbers, and the reason matters more than the numbers would.

Backtests executed on synthetic bar types fill orders at bar-transition prices that are not reliably attainable live. The bar closes *because* structure moved; an engine that retroactively buys that close is buying a price the move itself consumed. Add the platform defaults of zero commission and zero slippage, and the output is a statement about chart geometry, not about money. Any vendor showing you a Renko-family backtest with a fantasy profit factor is either unaware of this or hoping you are.

What synthetic-bar studies *are* good for is signal behavior: how often a structural flip is followed by continuation rather than immediate reversal, how flip frequency collapses in compression and expands in trend, and how the indeterminate state filters the chop that destroys time-bar systems. Those properties survive the fill problem, and they are why the signal layer exists. Live, broker-verified forward performance is the only P&L evidence we consider publishable — and the only kind we will ever publish.

A Second Opening Range, Stacked on the First

Here is the synthesis — and the correction to how most traders will first hear this idea. We are not proposing that you compute your ORB on structure bars *instead of* price. We are proposing that you run **both at once**.

The **price ORB** is the range you already trade: the literal high and low of the opening window, wicks and all. It is the level the crowd sees, which makes it indispensable — it is where the breakout orders, the stops, and the traps actually live.

The **structure ORB** is computed in parallel over the same window, on StructureBars. It marks the extremes of the *structural contest* — the levels the market actually fought over and defended, with the opening noise filtered out by construction. Because wicks stretch price extremes beyond where structure held, the structure range typically sits *inside* the price range, offset from it by the exact amount of noise the open produced.



FIGURE 2 The double stack live: price-based and structure-based opening-range levels rendered together with state coloring. The offset between stacked edges marks the pullback shelf — where retracements actually land — while levels where both ranges agree form the day's strongest reference edges.

Stack the two, and a single breakout line becomes a layered map:

Confluence bands. Where a price-ORB edge and a structure-ORB edge sit close together, the level was both *seen* by the crowd and *defended* by structure. These are the day's strongest reference edges — breaks through them carry conviction; rejections at them carry information.

The pullback shelf. Where the two ranges are offset, the gap between the price edge and the structure edge becomes the natural landing area for retracements. A breakout through the price level that pulls back typically finds its first real support not at the broken line — where every textbook says to look — but at the structure edge behind it. The stack tells you, in advance, where "expected support" actually is. The same geometry, inverted, marks expected resistance on failed breaks.

State qualification. Every interaction with either range is read through the structural regime from the previous sections. A break of the price ORB while the structure regime agrees with the direction is expansion; the same break against a contradicting regime is bait. A retest that holds at the structure edge *with* regime agreement is confirmation; the same retest in an indeterminate regime is a coin flip wearing a costume. TradeAdapt's ORB Zone toolset renders this continuously — both ranges and their derived levels live on the chart, colored and qualified by structural state, with dedicated break-quality tracking that distinguishes sustained structural breaks from single-print violations.

Two dimensions, literally stacked: the range the crowd trades, and the range the market defends. ORB traders have been trading the first one alone for decades.

What If the Open Wasn't Just 9:30?

The final extension is temporal. If an opening range is the contested auction that follows a liquidity handoff, then the trading day contains several opens, and each deserves its own stacked range.

The TradeAdapt multi-session ORB framework computes the double-stack at the day's principal liquidity events: the **Asian open**, where the overnight session establishes its frame; the **London open**, which re-prices Asia and frequently sets the European trend; the **8:00 ET window**, where U.S. economic data and pre-market institutional flow stake out the levels New York will inherit; the **NY 9:30 open**, the classic ORB; the **10:00 ET window**, where the first post-open releases, the initial-balance fight, and the open of the morning's institutional 4-hour bar produce the second auction of the day; and the **14:00 ET window**, where the afternoon 4-hour bar opens and repositioning flows — bond-close spillover, late-day institutional execution — contest the session's final range.

The framework's semantics are deliberately simple: the most recently completed range is the active reference frame. The market's newest auction supersedes the old one, exactly as a trader's own attention should. A London-range trader in the pre-dawn hours, a classic 9:30 breakout trader, and an afternoon specialist are all running the same engine pointed at different clocks — and a full-day trader watches the day hand off from one stacked range to the next, each one structure-qualified from the moment it locks.

The practical effect: "the ORB trade" stops being a once-per-day event with a single expiration and becomes a recurring, all-session methodology. Every few hours, the market opens again, fights over a range again, and tips its hand again — and each time, the stack tells you which levels were noise and which were defended.

SECTION 6

Availability

The components described in this paper are layers of one system. StructureBars supply the substrate. The structural-regime engine supplies state. The signal layer translates regime transitions into bias. The multi-session double-stack ORB supplies the day's reference frames.

StructureBars and the ORB Zone multi-session suite ship in TradeAdapt Pro — details and access at tradeadapt.ai/pro. The same structural inputs also feed TradeAdapt's AI desk pipeline, where model-driven desks read this data live and publish scored market calls throughout the session; LiveAI, the free TradeAdapt desktop widget, puts that AI desk on your screen at no cost, with every call scored openly — wins and losses alike. The full product line is at tradeadapt.ai.

ABOUT THE AUTHOR

David Lee Nelson

David Lee Nelson is the founder of TradeAdapt and the architect of the systems described in this paper. Before TradeAdapt, he was principal of CapitalScope Asset Management, a systematic hedge fund built on algorithmic trading, and founder of SmarteTrades, an algo-trading firm whose systems competed across multiple asset classes — including a "Forex 500" competition win. He then spent eleven years at ION Trading building institutional trading technology for tier-1 clients including JPMorgan Chase and Credit Suisse, writing the product and API specifications behind order-management systems processing millions of orders daily. TradeAdapt is the synthesis: institutional-grade structure analysis, built by a trader, running live in his own account every session.

SECTION 7

Risk Disclosure

Futures and forex trading involves substantial risk of loss and is not suitable for all investors. Past performance, whether actual or indicated by historical or simulated studies, is not indicative of future results. The descriptions in this paper are of analytical tooling; nothing herein constitutes trading advice, a recommendation, or an offer to buy or sell any financial instrument. Hypothetical or simulated performance results have inherent limitations, including the fact that they are generally prepared with the benefit of hindsight and do not involve financial risk; no representation is being made that any account will or is likely to achieve profits or losses similar to any discussed. Trade only with capital you can afford to lose.