Energy Consumption Between Trigger Points:
A Diffusion-Based Index for Anticipating Anomalous Moves in Index Futures

Dr. Efrat Levy
CTP Academy

September 2026 · working paper, version 0.2

Abstract

Discretionary traders describe extended consolidation between two support and resistance levels as *energy consumption*: the longer price is held inside a narrow band, the larger the move that eventually resolves it. This paper turns that description into a measurable quantity with a null model behind it. Price is modelled as a diffusion whose scale is set by the daily average true range; under that model the expected range grows with the square root of elapsed time, the expected time to leave a band of height h is proportional to h^2, and the probability of remaining inside the band for T hours decays exponentially in T/h^2. A band that survives far beyond these expectations is improbable under the null, and the improbability is the stored energy. We give four formulations of an *energy index*: a range deficit, a first-passage survival ratio, a confinement surprisal, and the compression work of a gas in a closed syringe; add multipliers for boundary tests and narrowness; and define an asymmetry term that reads the direction from the strength of reactions at the two boundaries. Physical analogies (a compressed spring, a gas under pressure, the Hurst exponent) motivate the functional forms and are labelled as analogies, not laws. The index is evaluated on three years of CME Globex market data for NQ, ES, RTY and YM against three baselines under a pre-registered, calendar-split protocol.

1 Introduction

On 10 September 2026 the Nasdaq future spent roughly twenty hours between two trigger points, 29,428.50 and 29,474.75, then fell through ten mapped support levels to 29,043, a session range of 127 % of its daily ATR. The session review that followed made two claims. First, that the consolidation itself was the warning: a market that has consumed energy for hours will make an anomalous move, and while the direction is not always knowable, the anomaly is. Second, that the direction can sometimes be read from the two boundaries: repeated strong rejections at resistance against a support that is merely respected mean that the more stubborn side is exhausting itself. The practical consequences were stated plainly. On such a day one does not buy every support on the way down, because a support that breaks becomes resistance; and at the extremes one widens the stop and reduces size so that the dollar risk is unchanged.

The same session also offered an example of what a system should do. The CTP signal generator issued exactly two long signals while ten support lines were run through, one losing and one winning near the day's extreme. A model that flags consolidation early and sizes the expected move would let both a discretionary trader and the generator act on that warning rather than react to each level.

This paper asks whether the description can be made operational. We want a single number, computable in real time from market data and the day's trigger points, whose value before the break is informative about the size of the move after it. We call it the energy index, and we call the ratio of the realised move to the stored energy the degree of anomaly.

Test first, build second. The decisive point of the project is its order. The index can be checked statistically before any indicator is written: three years of market data exist for every instrument, so the index can be computed for every historical band and compared with the range of the hours that followed. If the relationship exists, the index enters the NinjaTrader indicator with thresholds derived from the data rather than guessed. If it does not, the project has saved the construction of something beautiful that does not work. The recommendation, and the plan followed here, is therefore: the statistical test on historical data first, and only then the indicator with the index inside it.

2 Background

Trigger points. CTP trigger points are daily support and resistance levels delivered per contract. The live runners consume, for each session, the union of the current day's levels and those of the five preceding sessions. The levels are contract-specific: a level of the September contract is not a level of the December contract, because the roll gap sits between them. All analysis here therefore runs on market data of the same contract the levels were computed on.

Daily ATR. Throughout the CTP system the daily ATR (DAtr) is the mean of the last fourteen daily session ranges, the session running from 18:00 to 17:00 Eastern. It is the unit in which every threshold below is expressed, so that one rule serves instruments whose point values differ by two orders of magnitude. The production figure lags the latest session by one day (the levels bundle is built before the previous session has fully closed); the research reproduces the production figure exactly, so that a percentage of DAtr means the same thing in the paper and on the chart.

Related notions. Compression before expansion is an old observation, expressed as Bollinger squeezes, narrow-range days, Wyckoff accumulation and, in the order-flow literature, as absorption at a level. In the statistical literature the same intuition appears as variance-ratio tests [5] and the Hurst exponent [4], where a series that has moved less than a random walk over some horizon is called anti-persistent. What distinguishes the present formulation is that the band is not a statistical envelope but a pair of pre-existing levels, so the boundaries are known before the consolidation starts and the false break has a reference to be false against.

3 A Diffusion Null Model

Everything that follows is measured against one question: how far would price have travelled if nothing were holding it? The simplest answer is a driftless diffusion [3].

Scale from the daily ATR. Let price over a session follow a Brownian motion with volatility σ per $\sqrt{\text{hour}}$. The expected range of a Brownian path over a horizon T is

$$\mathbb{E}[R(T)] = \sqrt{\frac{8T}{\pi}}\,\sigma \approx 1.596\,\sigma\sqrt{T}, \tag{1}$$

so a session of length T_d hours with mean range DAtr fixes the scale:

$$\sigma = \frac{\mathrm{DAtr}}{1.596\sqrt{T_\mathrm{d}}}. \tag{2}$$

With $T_\mathrm{d} = 23$ hours this gives, for the motivating session (DAtr = 367 points), $\sigma \approx 48$ points per $\sqrt{\text{hour}}$.

Expected range after T hours. Under the null the expected range grows as the square root of elapsed time:

$$\mathbb{E}[R(T)] = \mathrm{DAtr}\sqrt{T/T_\mathrm{d}}. \tag{3}$$

This is the diffusion prior for the function $g(T)$ of the short version of this paper, $g(T) = \sqrt{T/T_\mathrm{d}}$. Real sessions are not homogeneous, so g is also estimated empirically per session profile (Asia, London, US) on the training set, and the two are compared.

First exit from a band. For a Brownian motion started at the centre of a band of height h, the expected time to first reach either boundary is

$$\mathbb{E}[\tau_\mathrm{exit}] = \frac{(h/2)^2}{\sigma^2} = \frac{h^2}{4\sigma^2}. \tag{4}$$

For the motivating band, $h = 46$ points and $\sigma \approx 48$, the expected exit time is about 0.23 hours, fourteen minutes. The band lasted twenty hours.

Probability of confinement. The probability that the path has not left the band by time T decays, for large T, as the first eigenmode of the interval:

$$P(\tau_\mathrm{exit} > T) \approx \frac{4}{\pi}\exp\!\left(-\frac{\pi^2\sigma^2 T}{2h^2}\right). \tag{5}$$

Its negative logarithm, the surprisal of confinement, is linear in T and in $1/h^2$:

$$-\ln P \approx \frac{\pi^2\sigma^2 T}{2h^2} - \ln\frac{4}{\pi}. \tag{6}$$

Figure 1: The proposed indicator on the motivating session (NQ, 5-minute, 10 September 2026): the band between 29,428.50 and 29,474.75 painted with an intensity that grows with age, the reaction read at each boundary, the break marked as the anomalous move, and the extreme flag once the session passes 100 % of DAtr. Illustration; numbers in the tags are not computed.

4 Definitions

Band. Let $L_1 < L_2$ be two adjacent trigger points of the current contract, or the lower and upper edges of a cluster of at most three levels. A band is open at the first candle close inside $[L_1, L_2]$ and stays open while every subsequent close remains inside it. Wicks may extend beyond the boundaries by at most w per cent of DAtr; a close beyond a boundary by more than b per cent of DAtr ends the band and marks the break. Defaults: $w = 5$, $b = 10$, both to be revisited in Section 8.

Duration and height. T is the band's age in hours; $h = L_2 - L_1$ its height. Both are used in dimensionless form: T/T_d and $\eta = h/\mathrm{DAtr}$.

Four formulations of the energy index. Each compares what the null model expects with what the band allowed. They are ordered from the most intuitive to the most principled, and the evaluation decides between them.

The *range deficit*, in per cent of DAtr, is the amount of movement the market owes:

$$E_R(T) = 100\left(\sqrt{T/T_\mathrm{d}} - \eta\right). \tag{7}$$

The *survival ratio* is elapsed time over the expected exit time, a pure number:

$$S(T) = \frac{T}{\mathbb{E}[\tau_\mathrm{exit}]} = \frac{4\sigma^2 T}{h^2} = \frac{4}{1.596^2}\,\frac{T/T_\mathrm{d}}{\eta^2} \approx 1.57\,\frac{T/T_\mathrm{d}}{\eta^2}. \tag{8}$$

The *confinement surprisal*, in nats, is how improbable the band is under the null:

$$E_S(T) = \frac{\pi^2\sigma^2 T}{2h^2} = \frac{\pi^2}{2\cdot 1.596^2}\,\frac{T/T_\mathrm{d}}{\eta^2} \approx 1.94\,\frac{T/T_\mathrm{d}}{\eta^2}. \tag{9}$$

S and E_S are proportional; both say that energy grows linearly with time and with the inverse square of the band height, which is the formal content of "narrow stores more".

The fourth formulation leaves the diffusion picture and takes the one the trader actually uses (Section 5): the band

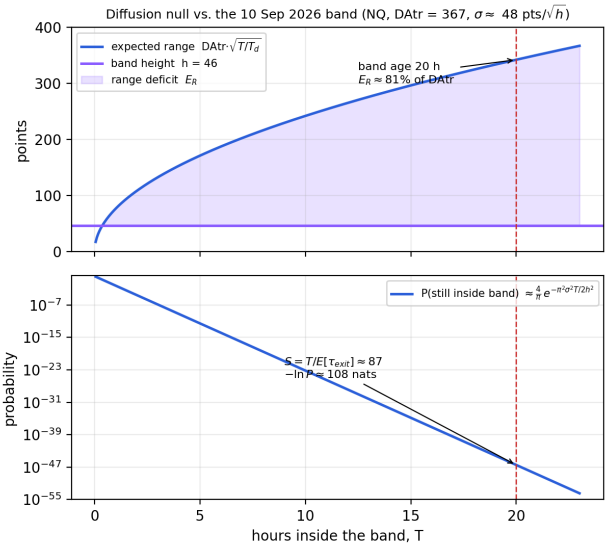

Figure 2: The diffusion null against the motivating band. Top: the expected range $\mathrm{DAtr}\sqrt{T/T_\mathrm{d}}$ rises with elapsed time while the band's height stays fixed at 46 points; the shaded gap is the range deficit E_R. Bottom: the probability of remaining inside the band under the null, on a log scale, with the survival ratio S marked at the band's actual age.

is a closed syringe and the session compresses the gas inside it. The *compression work*, up to a constant, is

$$E_W = n\,\sigma^2 \ln \frac{h_0}{h}, \tag{10}$$

with n the volume traded inside the band, σ^2 the instrument's variance (its temperature) and h_0 a reference height (the training-set median). Unlike the three diffusion bases, E_W carries traded volume, and it grows with the logarithm of the squeeze rather than with its inverse square.

For the motivating band, $T/T_\mathrm{d} = 20/23$ and $\eta = 46/367 = 0.125$, so $E_R \approx 81$ per cent of DAtr, $S \approx 87$ and $E_S \approx 108$ nats: under the null, a band like this happens essentially never, which is the quantitative form of the trader's "it has to explode". Figure 2 shows both readings against elapsed time.

Multipliers. Two refinements are added multiplicatively, each with a physical reading.

Tests at the boundaries. Let n_1 and n_2 be the number of tests at L_1 and L_2 (a test is a wick within w per cent of DAtr of the line, separated from the previous test by at least one candle that closed away from it). Every test places more stops just outside the band, and those stops are the fuel of the break. We use

$$M_n = 1 + \alpha \ln(1 + n_1 + n_2), \tag{11}$$

with α fitted on the training set; the logarithm reflects diminishing returns and keeps a chopping band from dominating.

Narrowness. Already inside S and E_S through $1/\eta^2$; for E_R it is added as

$$M_\eta = (\eta_0/\eta)^\beta, \tag{12}$$

with η_0 a reference height (the training-set median) and β fitted, $\beta = 2$ recovering the diffusion scaling. The combined

Figure 3: The boundary read. Each test of a line is scored by the wick beyond it and the retreat after it; here resistance shows five tests with strong rejection and one false break, support three tests merely respected. The asymmetry A is positive and the verdict is that sellers are stronger, the read the session review gave before the break.

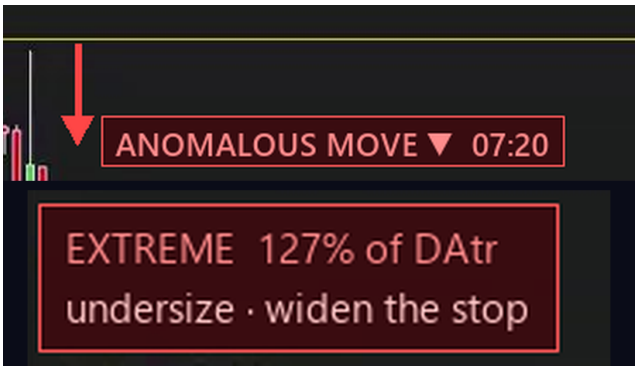

Figure 4: Top: the break, a close more than b per cent of DAtr beyond the band, marked as the anomalous move with its time. Bottom: the extreme flag, shown once the session's range exceeds 100 % of DAtr, carrying the second tip of the session review: undersize and widen the stop.

index is $E = E_* M_n M_\eta$ for whichever base E_* the evaluation selects, and the base without multipliers is always reported alongside.

Asymmetry. For the i-th test of a boundary we record the wick depth beyond the line, d_i, and the retreat that follows within m candles, r_i, both in per cent of DAtr, and define the rejection strength $\rho_i = d_i + r_i$. The asymmetry is the difference of mean rejection strengths at the two boundaries:

$$A = \bar{\rho}(L_2) - \bar{\rho}(L_1). \tag{13}$$

Positive A means resistance rejects harder than support holds, the configuration described in the session review as sellers being stronger; the sign of A is the direction read, and $|A|$ its confidence (Figure 3).

Anomaly. Before the break, the anticipated anomaly is E itself: the size of the move the market owes. After the break, the realised anomaly has two readings. The *release ratio* is the forward range in the k hours after the break over the range deficit,

$$\kappa_k = \frac{R_\mathrm{fwd}(k)}{E_R}, \tag{14}$$

which says how much of the stored energy was released; and the session's total range in per cent of DAtr, which is the figure already displayed on every CTP chart and the one the session review used to define the extreme (127 % on the motivating day). The two readings answer different questions, and both are reported (Figure 4).

5 Physical Analogies, and Their Limits

The vocabulary of the session review is physical, and the analogies are useful for choosing functional forms as long as they are not mistaken for conservation laws. Markets conserve nothing.

A compressed spring. A spring compressed by x stores $U = \frac{1}{2}kx^2$. If the "compression" is the shortfall of realised range against the null, $x \propto \sqrt{T/T_{\mathrm{d}}} - \eta$, the stored energy is quadratic in the range deficit. This suggests testing E_R^2 as well as E_R against the forward range, which the evaluation does.

The closed syringe. Take a syringe without a needle, block the outlet with one thumb and push the piston with the other. The air inside is a closed system and the process is the compression of a gas. Pushed slowly, so that heat leaks through the walls, it is isothermal: Boyle's law $PV = \text{const}$ holds and the work done is

$$W_{\mathrm{iso}} = nRT \ln \frac{V_0}{V}. \tag{15}$$

Pushed fast, so that no heat escapes, it is adiabatic: $PV^\gamma = \text{const}$, the gas heats, and the work is stored as internal energy,

$$W_{\mathrm{ad}} = \frac{P_1V_1 - P_0V_0}{\gamma - 1}. \tag{16}$$

Either way the thumb's work is stored in the gas, and when the thumb lets go the piston is thrown back: stored work becomes motion.

The mapping to the band is direct. In kinetic theory temperature is the mean squared velocity of the molecules; the market's temperature is σ^2. The amount of gas n is the volume traded inside the band. The volume V is the band's height h; the piston is the session pressing the range against two trigger points that do not give. The isothermal work then reads as equation (10): stored energy grows with the traded volume trapped in the band, with the instrument's variance, and with the logarithm of the squeeze h_0/h. Pressure, $P = nRT/V$, is the rate at which the walls are hit, the test rate of Section 4, and it rises when n grows while V is held: the same participants trading the same narrow range press harder against the lines. This is the reading behind the test-count multiplier, and it predicts that volume inside the band and the test rate carry information beyond T and η alone.

Release as kinetic energy. On the break the stored energy becomes motion; the natural measure of the released energy is the squared displacement, $\frac{1}{2}mv^2$ with v the speed of the move. The release ratio κ compares first moments; a second version compares squared quantities, R_{fwd}^2 against E_R^2, and the evaluation reports which relationship is tighter.

The Hurst exponent. For a series whose range scales as $R(T) \propto T^H$, a random walk has $H = 0.5$, a trending series $H > 0.5$, and a mean-reverting one $H < 0.5$. Inside a band the local H estimated from the last T hours falls below 0.5; the null model of Section 3 is exactly the statement $H = 0.5$. The band's persistence can therefore also be scored as the distance $0.5 - \hat{H}$, and the break as the moment $\hat{H}$ crosses back above 0.5. This is a feature, not a competing index, and it is included in the feature set.

What does not carry over. There is no conservation: energy stored may never be released if the day ends, and a break may be larger than anything the band explains because the news, not the band, moved the market. That is why the baselines of Section 7 include the event calendar and the hour of the day, and why the release ratio is reported as a distribution, never as a prediction of a single number.

Table 1: Databento size and cost for one NQ contract over one quarter.

Data	Size	Cost
tick-by-tick trades	1.1 GB	none (current plan)
top of book (mbp-1)	59 GB	none
full order book (mbo)	82 GB	USD 94
one-minute bars	negligible	none

6 Data and Implementation

Instruments and period. NQ, ES, RTY and YM, 28 September 2023 to 11 September 2026: 763 sessions with trigger points per instrument, 48 contract spans in total.

Market data. The study runs on Databento's CME Globex feed (GLBX.MDP3) for the four contracts, matched per contract to the trigger points, and its substance is the deep data: tick-by-tick trades with aggressor side, the top of the book, the full order book (every order added, modified and cancelled), and the CME options on the same roots with definitions, trades and open interest. These are the sources that carry what the session review describes and what a candle cannot show: absorption at a line, aggressive buyers who fail to move price, large orders resting on one boundary and not the other, and the option positioning that turns a resistance into a wall. Volatility indices (VIX, VVIX) come from a separate source. Timestamps are converted to Eastern time and assigned to CME sessions rolling at 18:00.

The first pass uses the coarsest aggregation of the same feed, one-minute bars, because they are available for all instruments over the whole three years at no cost and because they are enough to locate the bands and compute the basic index. They are an entry point, not the data set. Once the bands are located, the deep data are fetched for the band windows, which keeps the volumes tractable. Databento's own estimator gives, for one NQ contract over one quarter (Table 1), that trades for the whole period are about 50 GB at no cost and can be held in full, while the full order book would be about 4 TB and about USD 4,500, and is therefore fetched for band windows only. Options data for gamma walls are in the same dataset and form their own stage.

Order of work. Stage 2 locates bands and computes the basic index and the first feature group. Stage 3 fetches trades

for the band windows and adds delta, absorption and trade-size features; Stage 4 adds the order book for the iceberg features, then the family state, the calendar and options positioning. Each layer is measured against the model without it, so that the contribution of every layer is known rather than assumed.

Levels. Loaded with the same rule as the live runners and verified identical to the runner's output on ordinary days, on a roll week and on the first day of a new contract.

Features. For every band, at every hour of its life, a feature vector is recorded. The features are grouped by the data they need, and each is stated in the vocabulary of the session review.

From price bars (the entry layer, available for the full period).

- Range deficit against DAtr by elapsed time, the survival ratio and the confinement surprisal: the three bases of E.
- Band height η in per cent of DAtr: narrower stores more.
- Duration in hours, and how many session profiles the band has crossed (Asia, London, US): a band that survived London unbroken is stronger than two hours overnight.
- Tests at each boundary, n_1 and n_2, and the time between tests: the rhythm tends to shorten before the break.
- Reaction strength ρ at each boundary and the asymmetry A: *rejected* versus *respected*; the difference gives the direction.
- False breaks: a close outside the band that returned inside within m candles, the manipulation the review describes.
- Position of closes within the band: a rolling mean of $(\text{close} - L_1)/h$; a downward drift before the break supports "the buyers have run out".
- Local Hurst exponent $\hat{H}$ over the band's life, and the ratio of the 15-candle ATR to DAtr: compression that tightens with time.
- Volume: volume at the boundaries against volume inside the band, and volume on the break. On the motivating day the pre-market fall came on thin volume.

From trades and the order book (the deep layer; the CTP system's iceberg engine already reads it).

- Cumulative delta inside the band, $\sum(\text{buy volume} - \text{sell volume})$: aggressive buyers who do not move price are absorption, the reading one participant identified in the review from a footprint image.
- Absorption at the boundaries: traded volume on the level per point of penetration.
- Trade size: the presence of large trades near one boundary and not the other.

From the correlated family (the CTP signal engine's state).

- Parallel bands in ES, RTY and YM: four instruments consolidating at once is family energy. The short of 9 September was taken with a SPECIAL line in all three.
- Who broke first: the leader breaks while the others are still in their bands; by the CTP doctrine the leader is the false-break suspect and the last to arrive decides.
- Family disagreement: one member on support while another sits on resistance during the band.

Daily context (the configuration and the statistics store).

- DAtr already consumed by the session: a band after 80 per cent is not a band at 20 per cent.
- Position of the band in the day: at the high, at the low, or in the middle.
- The opening range: whether the band is the ORB itself, and the session's historical double-break probability.
- The session gap and its fill state.
- The event calendar: FOMC and CPI dates are already configured; 08:30 Eastern, named in the review, is a data release. A band that ends into a scheduled event is almost always the break.
- Day of week and hour, for which statistics already exist.

Requiring a new source (deferred).

- Options positioning: max pain and gamma walls, for which code exists in the Next_Gen project. A gamma wall above the band explains why resistance rejects hard and where the break will come from.
- VIX and VVIX: compression in volatility alongside compression in price.
- Micro against mini volume: the ratio of MNQ to NQ volume as a gauge of small money trapped.

Targets. The forward range in the next 2, 4 and 8 hours from the moment of evaluation; the range from evaluation to the end of the session; and, for bands that break, the range from the break to the end of the session, all in per cent of DAtr; and the sign of the break for the direction test.

Software. A standalone Python project, `CTP_Energy`, with a parquet cache and a per-run record of every threshold used.

7 Evaluation Protocol

The problem. Overfitting is the discovery of a rule that describes the past sample rather than the market. Three features of this study make it a live danger rather than a formality: about 700 sessions per instrument, which is small; sessions that are correlated with their neighbours, so that random splits leak the future into the past; and a hypothesis born from one day, 10 September 2026, which everyone involved has already seen. A rule tuned until it fits that day proves nothing. The protocol below is fixed before any result is examined.

1. **Calendar split, fixed before any result is seen.** Training from October 2023 to March 2025, validation from April to December 2025, test from January to September 2026. The test set is opened once, at the end, with the rule frozen; if it fails, the answer is "it does not work", not "let us adjust". Random splits are forbidden because adjacent days leak.
2. **Hypothesis in writing before the test.** Every candidate rule is written with its thresholds in `output/hypotheses.md` before it is evaluated. A rule changed after a validation look returns to training.
3. **Few knobs, physical units.** Thresholds are percentages of DAtr and hours, never points, and never per instrument. One rule for the whole family, or no rule.
4. **A budget of looks.** Validation answers a handful of candidate rules, and every look is logged.

5. **Baselines that must be beaten.** (a) any consolidation of N hours without reference to levels; (b) the hour of the day alone, which carries the 08:30 and 09:30 effects; (c) the DAtr already consumed. The index is worth building only if it adds to all three.
6. **Effect size, not p-value.** The distribution of the forward range by decile of E, with bootstrap intervals, separately per year and per instrument. Success is a monotone relationship that holds in each year; a relationship present in 2024 and absent in 2025 is not a finding.
7. **Walk-forward at the end.** $g(T)$, α, β and any calibration are fitted on a rolling twelve-month window and applied to the following three months, then rolled.
8. **No leakage.** DAtr for a day is the mean of the prior sessions, and every feature uses only market data up to the moment of evaluation.

Statistics. Within that protocol the tests are fixed as well. For magnitude: Spearman rank correlation between E and $R_{\text{fwd}}(k)$ per year; quantile regression of $R_{\text{fwd}}(k)$ on E with the three baselines as covariates, reporting the coefficient on E at the median and the 90th percentile with block-bootstrap intervals (blocks of five sessions, to respect autocorrelation); and the incremental R^2 of E over the baselines. For direction: the agreement rate between the sign of A and the direction of the break, tested against one half with a binomial test, stratified by $|A|$ tercile. For the three bases of E: the same battery on each, and the base with the largest out-of-sample rank correlation on validation is the one carried to test.

8 Expected Results and Discussion

If the description in the session review is right, four things should appear. The forward range should rise with E, and the rise should survive conditioning on the hour and on the ATR consumed. The number of boundary tests should sharpen the relationship rather than replace it. The asymmetry A should carry sign information, so that the direction of the break agrees with the sign of A more often than chance, especially for large $|A|$. And the confinement surprisal, the most principled of the three bases, should not do worse than the range deficit; if it does, the null model is missing something about how sessions actually unfold, which is itself worth knowing.

Two outcomes would be instructive failures. If E predicts the range but no better than the hour of the day, the phenomenon is the news calendar wearing a consolidation's clothes. If E predicts range only when the band is the day's low or high, the effect is a range-breakout effect, not a level effect, and the trigger points add nothing.

On the motivating session, twenty hours in a 46-point band with a 367-point DAtr gives E_R of roughly 81 per cent before the break and a survival ratio near 87; the session's realised range was 127 per cent, a release ratio κ of about 1.6. One day proves nothing, which is the reason for Section 7.

9 Limitations

The trigger points are a proprietary input, so replication requires the same level source. The three-year sample contains one full rate cycle and one election; regimes outside it are unrepresented. The diffusion null ignores intraday seasonality of volatility, which the empirical $g(T)$ partly repairs, and it ignores jumps, which the event calendar partly repairs. The order-book and options layers are the most expensive to acquire and are evaluated after the trades layer; until then the absorption reading rests on trades alone. The definition of a band depends on two tolerances, w and b, and the sensitivity of every result to them will be reported.

10 Planned Application

The order is fixed: the statistical test first, the indicator second, and the indicator's thresholds are the ones the data selected. If the index passes Section 7, two products follow. A NinjaTrader indicator, CTP Energy, draws the band between the two trigger points with an intensity that grows with E, annotates each boundary with its test count and reaction read, and marks the break as the anomalous move. And the CTP signal generator's regime memory receives E and A as inputs, so that a strong energy reading suppresses countertrend entries on the levels below a breaking band, the behaviour the session review asked of traders.

References

[1] Levy, E. (2026). *Energy Consumption and Extreme Risk Management*, recorded session review, 10 September 2026.

[2] CTP Academy (2026). CTP Signals: correlated trigger points signal engine, internal documentation.

[3] Feller, W. (1971). *An Introduction to Probability Theory and Its Applications*, Vol. II. Wiley.

[4] Hurst, H. E. (1951). Long-term storage capacity of reservoirs. *Transactions of the American Society of Civil Engineers*, 116, 770–799.

[5] Lo, A. W. and MacKinlay, A. C. (1988). Stock market prices do not follow random walks: evidence from a simple specification test. *Review of Financial Studies*, 1(1), 41–66.

[6] Wyckoff, R. D. (1931). *The Richard D. Wyckoff Method of Trading and Investing in Stocks*.

[7] Databento (2026). GLBX.MDP3 dataset documentation.