

VOLATILITY, RETRACEMENT & MEAN REVERSION INSIDE NASDAQ FUTURES

Rishabh Vengavillagam

rvengavillagam@gmail.com | September 20th, 2026

Abstract

Markets don't move for just one reason, since CPI prints, GDP revisions, geopolitics, corporate news, and the fight of supply and demand are all pulling on price at once. Price, however, is manipulated more by certain factors based on the levels. In the small, second to second micro-movements of given markets for example, we see more supply-demand and potential imbalance. This study builds on a narrower topic relating to supply and demand, in the form of information asymmetry imbalances. These imbalances are discontinuities that, in theory, retrace back to price regions the market aggressively moved past. This theory is tested to see what holds by building a detection algorithm that could find these price regions (defined by aggressive order flow) and running it against a full year of data. What came back was 1,209 qualifying three bar microstructure sequences over that year, and a 96.1% eventual mean reversion rate, with a 95% confidence interval from 95.0% to 97.2%.

Two important numbers to look at to contextualize that mean reversion rate are time to mitigation and maximum adverse excursion(MAE), since MAE is what tracks the worst unrealized loss a position hits between entry and whenever it finally closes. While 96.1% of imbalances eventually fill, 23.4% of them drag through a drawdown past 1% prior, and the median time to mitigation is 4.7 hours.

There is a real positive correlation between how often these imbalances form and where the VIX sits, and it holds up even though the reversion rate itself barely moves across volatility regimes; only the speed of mitigation does, dropping sharply once things get more volatile. What all of this adds up to is that these imbalances revert in a way that does not seem to care much about directional bias one way or the other. But the distinction between an imbalance eventually reverting and an imbalance being something you could build a viable, risk-adjusted strategy around turns out to matter quite a lot

INTRODUCTION

Futures markets discover price continuously, which mostly means the process looks smooth from a distance even though it is really just a long sequence of small negotiations happening one resting order at a time, and every so often that negotiation gets interrupted by a burst of aggressive market orders that consumes several resting liquidity levels faster than the rest of the book can react. What is left behind is a dislocation, a stretch of price where almost nothing traded, and this paper is going to keep calling that an asymmetry imbalance or an information asymmetry imbalance throughout, since that framing is closer to what is happening mechanically than just calling it a gap. These imbalances are a temporary departure from equilibrium, and that departure is worth studying because it might say something about where liquidity wants to sit and how price discovery unfolds once the dust settles.

To find out, the formation, characteristics, and eventual mitigation of these imbalances got examined using minute-level data across the entirety of 2025 and all 24 hours of the CME session, and every three-bar sequence that satisfied the formation criteria was pulled out and tracked forward to see what happened to price afterward.

Three questions organize everything that follows. First, what fraction of these imbalances actually experience mean reversion, meaning price comes back and revisits them, and does that fraction move around depending on the broader market regime. Second, once an imbalance does mitigate, what does the distribution of time-to-mitigation and maximum adverse excursion look like, and what does that distribution imply about whether trading this pattern is economically feasible rather than just statistically real. Third, how does the sheer frequency of imbalance formation relate to VIX, and is that relationship the same one governing whether imbalances eventually fill, or are those two things actually independent of each other.

Prior trading literature has recognized gap-filling tendencies for a long time, informally at least, but a rigorous, intraday treatment of this specific microstructure pattern in futures markets has mostly been missing, and this paper is trying to fill that gap, no pun particularly intended. The contribution here is empirical evidence that these imbalances revert at a high rate regardless of trend direction, that there is real, substantial variation underneath that headline number once you look at time-to-mitigation and drawdown exposure, and that imbalance formation frequency tracks volatility closely even while the reversion rate itself stays remarkably indifferent to it.

LITERATURE REVIEW

MARKET MICROSTRUCTURE AND PRICE DISCOVERY

The starting point here is Kyle's (1985) sequential trading model, which shows how an informed trader can exploit temporary information asymmetries by spreading their activity out over time rather than revealing everything in one shot, and Glosten and Milgrom (1985) build on a related idea by showing that bid-ask spreads exist largely to compensate market makers for the risk of trading against someone who knows more than they do. Parlor and Seppi (2008) push this further into limit order book environments specifically, where the choice between submitting a limit order and a market order ends up shaping where liquidity clusters at any given price level.

Hasbrouck (1991, 1995) is really the empirical backbone underneath all of this, since his work established that trades themselves carry information about fundamental value and that markets absorb that information through fast, ongoing adjustment. What this paper does differently is shift attention away from that kind of continuous, incremental adjustment and toward discrete events instead, moments where order flow is aggressive enough that it skips several price levels entirely rather than working through them one at a time.

IMBALANCE DYNAMICS AND MEAN REVERSION

Most of the early empirical work on price imbalances focused on overnight gaps in equity markets rather than anything intraday. Blume, MacKinlay, and Terker (1989) found that opening imbalances following the October 1987 crash partially reversed over the course of the trading day, and Huang and Tsai (2008), looking at Taiwanese futures markets, found that roughly 70% of opening imbalances get partially filled within five sessions, though their analysis stayed entirely within the overnight-gap framework rather than looking at the kind of intraday microstructure phenomenon this paper is built around.

HIGH-FREQUENCY TRADING AND LIQUIDITY PROVISION

Brogaard, Hendershott, and Riordan (2014) show that high-frequency traders end up playing two roles at once, supplying liquidity in some moments and consuming it in others, with that balance tipping toward consumption specifically during periods of elevated volatility and fresh information arrival. Menkveld (2013) adds a useful mechanism here too, since his model has designated market makers holding continuous two-sided quotes most of the time but periodically building up

inventory imbalances that force them into rapid, aggressive repositioning.

That mechanism offers a fairly natural explanation for both halves of what this paper documents, since aggressive HFT order flow consuming resting limit orders is exactly what would create the kind of temporary price dislocation being studied here, and the mean reversion that follows could just as easily reflect the same market makers, or a competing liquidity provider, quoting back into the region once the initial inventory pressure eases. A 96.1% eventual reversion rate is a strong signal that something like this is happening structurally rather than by coincidence.

VOLATILITY AND MICROSTRUCTURE PATTERNS

Fleming and Kirby (2011) establish that trading activity and return volatility move together closely in futures markets, largely because volatility itself tends to cluster around bursts of information arrival, and Andersen (1996) shows that scheduled macroeconomic announcements reliably produce both a spike in volatility and a disruption to liquidity at the same time. The 0.80 correlation documented later in this paper, between imbalance formation frequency and VIX, fits comfortably into that same literature, though it goes a step further by showing that the reversion rate itself barely moves even as formation frequency swings wildly with volatility, which suggests something structural is driving mitigation rather than something regime-dependent.

METHODOLOGY

FORMAL DEFINITION OF INFORMATION ASYMMETRY IMBALANCES

Let $\{t_{-1}, t_0, t_{+1}\}$ denote three consecutive one-minute bars, and for each bar t_i let $H(t_i)$ and $L(t_i)$ denote its high and low, giving a range of

$$R(t_i) = H(t_i) - L(t_i).$$

A bullish imbalance exists across $\{t_{-1}, t_0, t_{+1}\}$ once three things are all true at once: the low of the third bar sits above the high of the first, meaning

$$\begin{aligned} L(t_{+1}) &> H(t_{-1}) \\ R(t_0) &> R(t_{-1}) \text{ and } R(t_0) > R(t_{+1}) \\ |L(t_{+1}) - H(t_{-1})| &\geq \tau \end{aligned}$$

where τ is a minimum size filter set to 0.1 index points, and a bearish imbalance mirrors this exactly with the inequalities flipped.

Each of these three conditions is doing something specific. The first makes sure an aggressively move exists rather than just a small overlap. The second makes sure it was the middle bar, not its neighbors, that produced the violent move responsible for the gap, which rules out cases where the surrounding bars are just naturally choppy. The third exists purely to filter out tick-level noise that would otherwise flood the sample with imbalances too small to mean anything economically.

For a bullish imbalance the upper boundary sits at $B_u = L(t_{+1})$ and the lower boundary at $B_l = H(t_{-1})$, and for a bearish imbalance those flip to $B_u = L(t_{-1})$ and $B_l = H(t_{+1})$.

MEAN REVERSION AND MITIGATION

An imbalance V formed at t_f counts as mitigated once price fully crosses back through it and reaches the opposite boundary at some later time $t_m > t_f$, which for a bullish imbalance means

$$\begin{aligned} \text{Low}(t_m) &\leq B_l \\ \text{High}(t_m) &\geq B_u \quad (\text{bearish case}) \end{aligned}$$

and this is a strict definition on purpose, since a partial retracement does not count as mitigation here, only a full one does.

RISK-ADJUSTED METRICS

A reversion rate on its own cannot really tell you whether a pattern is economically viable, which is why two additional metrics get tracked for every mitigated imbalance. The first is maximum adverse excursion, defined for an imbalance V with formation time t_f and mitigation time t_m as

$$\text{MAE}(V) = \max_{t \in [t_f, t_m]} \frac{P(t) - P(t_f)}{P(t_f)}$$

where $P(t)$ is the midpoint price at time t , capturing the worst point the position would have reached before eventually being proven right. The second is **time-to-mitigation**, $\Delta T(V) = t_m - t_f$, measured in hours.

STATISTICAL FRAMEWORK

Confidence intervals for the reversion rates use the Wilson score method,

$$\text{CI} = \frac{\hat{p} + \frac{z^2}{2n} \pm z \sqrt{\frac{\hat{p}(1-\hat{p})}{n} + \frac{z^2}{4n^2}}}{1 + \frac{z^2}{n}}$$

with $\hat{p}$ the observed rate, n the sample size, and $z = 1.96$ for 95% confidence, chosen over the simpler normal

approximation specifically because that approximation tends to break down at proportions this close to 100%. The relationship between imbalance formation frequency and VIX gets tested with a Pearson correlation between daily counts and the same-day VIX close, and significance is assessed with a standard t-test.

ALGORITHM IMPLEMENTATION

Detection runs in Pine Script v5 on TradingView, scanning every new one-minute bar against the three conditions above and keeping a live registry of every imbalance that has not yet mitigated, updating that registry's status continuously as new bars arrive. For each imbalance the algorithm logs the formation timestamp, the boundary prices, whether it was bullish or bearish, its size, the mitigation timestamp once one exists, and the realized MAE along the way.

DATA AND EMPIRICAL IMPLEMENTATION

DATA DESCRIPTION

Everything here uses one-minute OHLC bars for the E-mini Nasdaq-100 futures contract on the CME, covering January 1 through December 31, 2025, across the full 24-hour Globex session, which comes out to roughly 525,600 bars once holidays and maintenance windows are accounted for.

Restricting this to U.S. regular trading hours would have been a mistake, since NQ trades nearly continuously and a meaningful share of imbalance formation happens during Asian and European hours, particularly around scheduled U.S. data releases that land well outside the RTH window, so using the full session avoids introducing a sample-selection bias against exactly the kind of information-driven displacement this paper cares about.

SAMPLE SELECTION AND FILTERING

Running the detection criteria against the full dataset produced 1,209 imbalances in total, split into 569 bullish (47.1%) and 640 bearish (52.9%).

Category	Count
Total Detected Imbalances	1,209
Bullish Imbalances	569 (47.1%)
Bearish Imbalances	640 (52.9%)

The 0.1-point threshold was not chosen arbitrarily. Dropping it to zero increases the raw imbalance count roughly tenfold, but nearly all of that extra volume turns out to be tick-level noise that obscures the underlying structural pattern rather than revealing more of it, so 0.1 points represents a deliberate

compromise between catching real displacements and filtering out noise.

IMBALANCE CHARACTERISTICS

Imbalance size is right-skewed, with a median around 2.15 points, though occasional displacements exceed 40 points during extreme events, the April 2025 tariff announcement being the clearest example in this sample.

VOLATILITY PROXY

The CBOE VIX serves as the volatility proxy throughout, matched to imbalance counts by date. It technically tracks S&P 500 rather than Nasdaq-100 implied volatility, but the two move together closely enough, with a correlation above 0.90 across the sample period, that using VIX here seems like a reasonable choice rather than a stretch.

RESULTS

BASELINE REVERSION RATES

Of the 1,209 imbalances detected, 1,162 mitigated within the observation window, putting the overall reversion rate at 96.1% with a 95% confidence interval running from 95.0% to 97.2%. Bullish imbalances mitigated slightly less often, at 95.6% (544 of 569), while bearish imbalances mitigated at 96.6% (618 of 640), and while that gap might look like it is pointing toward the familiar "down faster than up" pattern in equity indices, it does not clear conventional significance thresholds under a two-proportion z-test.

What stands out more than the directional split is how stable the overall rate stayed across the year, with all 24 of the non-overlapping fifteen-day windows in the sample clearing a 90% reversion threshold, which reads less like a lucky year and more like a structural feature of the microstructure itself.

TIME-TO-MITIGATION ANALYSIS

A 96.1% eventual reversion rate only tells part of the story, since how long it takes to get there matters just as much for anyone trying to trade this.

Percentile	Time to Mitigation
25th	1.2 hours
50th	4.7 hours
75th	18.3 hours
90th	67.4 hours
95th	142.6 hours

Roughly 40% of imbalances mitigate within four hours, but the tail stretches out considerably further than that, with 17% taking longer than a full day and the 95th percentile reaching nearly six days. Splitting by VIX quartile makes the pattern clearer still, since median time-to-mitigation runs 8.9 hours in the calmest quartile but compresses down to just 1.1 hours in the most volatile one, suggesting elevated volatility accelerates how fast the order book replenishes itself rather than just coinciding with more imbalances forming in the first place.

MAXIMUM ADVERSE EXCURSION ANALYSIS

The 96.1% figure hides a fair amount of interim pain, which is exactly what MAE is meant to surface.

MAE Threshold	% Exceeding
0.25%	51.3%
0.50%	37.8%
1.00%	23.4%
2.00%	11.7%
3.00%	2.0%
5.00%	0.8%

Nearly a quarter of imbalances that eventually mitigate move against the expected direction by more than 1% first, and at the kind of 10x leverage that is common in NQ futures, a 1% adverse move against a \$100,000 position is a \$10,000 unrealized loss, which is more than enough to trigger a margin call well before the eventual reversion ever shows up. Imbalances that take longer to mitigate also tend to hurt more along the way, which is exactly the combination that makes a naive "fade every imbalance" strategy dangerous even with a 96.1% eventual win rate sitting behind it.

VOLATILITY CORRELATION ANALYSIS

Monthly imbalance counts swing from 25 in September up to 299 in April, a twelve-fold range driven mostly by the April tariff shock, when VIX averaged 24.70 with intraday spikes above 60.

Month	Imbalances	Avg VIX	Reversion %
Jan	125	16.43	92.0%
Feb	77	19.63	96.1%
Mar	129	22.28	94.6%
Apr	299	24.70	95.0%
May	99	18.57	94.9%
Jun	56	16.73	94.6%
Jul	33	16.72	93.9%

Month	Imbalances	Avg VIX	Reversion %
Aug	43	15.36	95.3%
Sep	25	16.28	92.0%
Oct	64	17.44	95.3%
Nov	123	16.35	95.9%
Dec	61	14.95	96.7%

The Pearson correlation between monthly count and average VIX comes out to 0.80 ($p = 0.002$), confirming that elevated volatility really does drive more imbalance formation. But the reversion rate over those same months only ranges from 92.0% to 96.7%, a standard deviation of just 1.40%, and the correlation between reversion rate and VIX is statistically indistinguishable from zero. Volatility clearly determines how many imbalances show up, but it has almost no bearing on whether they eventually fill, which reads as evidence that order-book replenishment is a structural feature of this market rather than something that comes and goes with the regime.

A smaller pattern worth mentioning: the Sep-Oct window showed a fairly even split between bullish and bearish imbalances (45.6% and 54.4% respectively), a less extreme asymmetry than an earlier pass at this data had suggested, occurring during a stretch of low-volatility, AI-driven grind higher in the underlying index.

ROBUSTNESS CHECKS

ALTERNATIVE MITIGATION DEFINITIONS

Relaxing the mitigation threshold to 50% or 75% retracement instead of full retracement pushes the rate up to 98.7% and 97.4% respectively, compared to 96.1% under the strict definition, and the fact that all three sit comfortably above 96% suggests the headline result is not particularly sensitive to how strictly mitigation gets defined.

IMBALANCE SIZE SENSITIVITY

Quintile	Size Range (pts)	Reversion %
Q1	0.10–0.95	97.3%
Q2	0.96–1.87	96.8%
Q3	1.88–3.14	96.0%
Q4	3.14–5.73	95.1%
Q5	5.74+	94.6%

Reversion declines a little as imbalances get larger, though every single quintile still clears 94.5%, which might just mean bigger displacements carry more information and are therefore a bit less likely to be purely mechanical.

FORMATION TIME-OF-DAY EFFECTS

No statistically meaningful differences in reversion rate turned up across Asian hours, European hours, U.S. RTH, or U.S. after-hours, which supports using the full 24-hour dataset in the first place and suggests this pattern is really about microstructure mechanics rather than anything specific to who happens to be trading at a given hour.

DISCUSSION

ECONOMIC INTERPRETATION

A 96.1% mitigation rate that barely moves across volatility regimes points toward these imbalances being temporary disequilibria inside the limit order book rather than shifts in fundamental value. Aggressive order flow consumes resting limit orders and creates the imbalance, liquidity gets replenished as market makers refill the book or contrarian traders fade the move, and the eventual reversion is really just the order book normalizing rather than the market absorbing new information. Menkveld's (2013) model of HFT inventory cycles fits this story well, and the high reversion rate suggests there is a economic incentive pulling prices back toward these depleted regions, presumably because liquidity providers see profit in quoting back into a temporarily vacant price zone.

THE MARTINGALE CRITIQUE AND PRACTICAL VIABILITY

A 96.1% hit rate sounds like a strategy until you start accounting for what it costs to sit through. Nearly a quarter of imbalances exceed 1% MAE before mitigating, which is a problem under the 5 to 10x leverage common in futures trading. A median 4.7-hour hold ties up capital in a way that gets expensive during trending regimes, where the opportunity cost of waiting can eat into or exceed the eventual gain. And none of this analysis accounts for the frictionless entry and exit it assumes, meaning spreads, slippage, and commission would all cut into whatever edge remains. The gap between a pattern being statistically real and a pattern being something you can trade profitably is really the whole point of this section.

COMPARISON TO EXISTING LITERATURE

The 96.1% intraday rate documented here sits well above the roughly 70% overnight gap-fill rate Huang and Tsai (2008) found in futures markets, and a few things probably explain that gap. Overnight imbalances often carry actual news, earnings or macro data, which can justify a permanent repricing rather than a mechanical fill, whereas intraday imbalances mostly just reflect order flow. Continuous intraday trading also allows for much faster order-book replenishment

than overnight gaps get, since those have to cross a discrete session boundary first. And the three-bar filter used here is specifically built to select for liquidity-driven displacement, while overnight-gap studies tend to sweep in every kind of price discontinuity regardless of what caused it.

LIMITATIONS AND FUTURE RESEARCH DIRECTIONS

SINGLE-YEAR SAMPLE PERIOD

This analysis only covers 2025, and confirming that the 96.1% figure holds up as a regime-independent property rather than something specific to this particular year would require multi-year data, ideally including a period of stress comparable to March 2020.

TRADE EXECUTION ASSUMPTIONS

What gets tracked here is imbalance mitigation, not executed trades, so a strategy backtest would still need explicit entry and exit rules, slippage modeling, position sizing, and stop-loss discipline, none of which are built into this analysis.

THRESHOLD PARAMETER SELECTION

The 0.1-point threshold was picked pragmatically rather than derived from anything more principled, and a more rigorous selection process, using information-theoretic criteria or variance-ratio tests, is left for later work.

DIRECTIONAL INDEPENDENCE ASSUMPTION

Bullish and bearish imbalances get treated symmetrically throughout this analysis, but the Sep-Oct asymmetry noted earlier suggests that symmetry might not hold uniformly across every regime, and a model that conditions on prevailing trend could probably refine these baseline estimates further.

ORDER BOOK DEPTH

This dataset only has OHLC bars, not order book depth or volume-at-price, and Level 2 data would let a future version of this study test directly whether mitigation probability depends on depth asymmetry around the imbalance boundary.

COMPARISON ACROSS ASSET CLASSES

Everything here is specific to NQ, so replicating this across other index futures like ES or YM, plus commodities and FX, would be a test of whether this framework generalizes beyond the Nasdaq-100 complex specifically.

CONCLUSION

Intraday information asymmetry imbalances in NQ futures revert at a high, stable rate, 96.1% across 1,209 sequences spanning the full 2025 trading year, and that rate barely moves across volatility regimes despite massive swings in how often these imbalances form. Underneath that headline number sits heterogeneity in time-to-mitigation, with a median of 4.7 hours stretching out to nearly six days at the tail, and in maximum adverse excursion, where 23.4% of imbalances exceed a 1% drawdown before they resolve. That decomposition is really the point of this paper, since a high eventual-reversion probability can still hide enough interim risk to make a strategy non-viable under realistic capital constraints, a distinction most retail-facing treatments of this pattern tend to skip entirely.

Formation frequency correlates strongly with VIX, at 0.80, while the reversion rate itself does not, and that decoupling is arguably the most interesting single finding here, since it points toward something structural driving mitigation rather than something that comes and goes with the broader market regime. Anyone trying to trade this needs to take that seriously: the setups are high-probability, but they demand risk management built around the wide TTM and MAE distributions documented above, not just confidence in the headline hit rate.

REFERENCES

1. Andersen, T. G. (1996). Return volatility and trading volume: An information flow interpretation of stochastic volatility. *Journal of Finance*, 51(1), 169–204.
2. Blume, L., MacKinlay, A. C., & Terker, B. (1989). Order imbalances and stock price movements on October 19 and 20, 1987. *Journal of Finance*, 44(4), 827–848.
3. Brogaard, J., Hendershott, T., & Riordan, R. (2014). High-frequency trading and price discovery. *Review of Financial Studies*, 27(8), 2267–2306.
4. Fleming, J., & Kirby, C. (2011). Long memory in volatility and trading volume. *Journal of Banking & Finance*, 35(7), 1714–1726.
5. Glosten, L. R., & Milgrom, P. R. (1985). Bid, ask, and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1), 71–100.
6. Hasbrouck, J. (1991). Measuring the information content of stock trades. *Journal of Finance*, 46(1), 179–207.
7. Hasbrouck, J. (1995). One security, many markets: Determining the contributions to price discovery. *Journal of Finance*, 50(4), 1175–1199.
8. Huang, Y. S., & Tsai, P. L. (2008). On the decision to trade and the evolution of price imbalances in Taiwan futures markets. *Applied Economics*, 40(13), 1683–1695.
9. Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica*, 53(6), 1315–1335.
10. Menkveld, A. J. (2013). High-frequency trading and the new market makers. *Journal of Financial Markets*, 16(4), 712–740.
11. Parlour, C. A., & Seppi, D. J. (2008). Limit order markets: A survey. *Handbook of Financial Intermediation and Banking*, 5, 63–95.