Volatility Measured as Barrier Crossings (Limit Orders):

Gamma Capture and the Future of Intraday Risk Measurement

Author: Prakhar Singh
Sponsor: Louis J. Pellathy

Abstract

This paper presents the implementation, validation, and analysis of **Gamma Capture**, a realized volatility estimator that uses barrier crossings. Gamma Capture works on millisecond, second, or minute data, with a lookback window determined by the granularity of the input data rather than by the method itself. In this paper, Gamma Capture is applied with minute data and a one-day lookback (390 data points per Daytime Session) and is benchmarked against three other volatility-input methods used to price SPY 0-DTE ATM straddles: 5-day Average True Range, prior-day VIX close, and standard deviation of log returns. Because historical 0-DTE options data is not available, the realized Daytime Session Open-to-Close move (9:30 AM to 4:00 PM ET) serves as the benchmark for each method's straddle quote. Working under sponsor Louis J. Pellathy, all four methods were implemented as competing market-maker pricing engines on the same SPY 0-DTE option across a Calendar Year 2025 backtest, with each method differing only in how it estimates volatility. The Black-Scholes formula serves as the common pricing layer for the latter three methods, while Method 1 quotes directly off ATR. The result is that Gamma Capture, by counting barrier crossings rather than squaring returns, produces a volatility signal that is more responsive to actual order-book activity and less distorted by jumps. Beyond the backtest, Gamma Capture is currently in production on both NinjaTrader and TradingView, where it runs as a live indicator suite on real market data. This document details the methodology, the mathematical derivations behind each method, the personal technical contributions to the project, and the high-level findings.

1 Introduction

Volatility is the central input to nearly every risk and pricing model in modern finance, from Black-Scholes option pricing to value-at-risk calculations to position sizing in systematic strategies. Yet the dominant family of volatility estimators, all built on the standard deviation of log returns, inherits a set of assumptions from Geometric Brownian Motion (GBM) that break down at intraday timescales. GBM assumes price evolves through smooth, continuously distributed Gaussian increments. Real intraday markets exhibit jumps, gaps, microstructure noise, and order-driven impulses that no continuous diffusion can capture cleanly.

The **Gamma Capture** model reframes volatility as a *barrier-crossing problem.* Rather than measuring how far prices have moved on average, it counts how often price crosses a fixed barrier, a level analogous to a resting limit order. This shift from *magnitude of returns* to *frequency of executions* aligns the measurement with what actually happens in the order book and removes the dependency on a specific return distribution.

For this paper specifically, the benchmark window is the Daytime Session Open-to-Close move (9:30 AM ET to 4:00 PM ET). Because historical 0-DTE options data is not available, this realized daytime move stands in for the actual straddle price as the comparison value, since it is what the same-day option would have settled to at the close. The lookback window used by Gamma Capture in this paper, one full Day Session, is a consequence of using minute data: with 390 data points per session, a full-day window is the minimum that delivers enough barrier crossings for a stable reading. Higher-resolution data (second or millisecond) would permit a much shorter lookback such as 60 minutes.

This paper documents the work performed under sponsor Louis J. Pellathy, who tasked us with implementing four volatility estimation methods and comparing their behavior to Gamma Capture across simulated and real market data.

2 The Challenge Question

The sponsor, Louis, framed the project as follows: he gave us four methods to implement as competing market-maker pricing engines on SPY 0-DTE ATM straddles, and asked us to demonstrate our understanding of each, show the differences in how they behave, and explain how Gamma Capture compares. The four methods were:

1. **Method 1:** 5-Day Average True Range (ATR) as implied move proxy.
2. **Method 2:** Prior-day VIX close as implied volatility input to Black-Scholes.
3. **Method 3:** Standard deviation of log returns from minute data, fed into Black-Scholes.
4. **Method 4:** Gamma Capture barrier crossing volatility, fed into Black-Scholes.

The deliverable was a working implementation of each method, run on the same SPY 0-DTE instrument across the Calendar Year 2025 backtest, with comparative results showing how each method's volatility estimate translates into a straddle quote and how those quotes line up against the realized straddle value at close.

3 Background and Terminology

3.1 Geometric Brownian Motion (GBM)

The foundational stochastic model in which the log of the price evolves as a Brownian motion with drift. GBM produces continuous paths and Gaussian-distributed log returns. It is the implicit assumption behind most standard volatility estimators.

3.2 Realized Volatility

The annualized standard deviation of observed log returns over a window. Highly sensitive to outliers because each return enters the calculation squared.

3.3 Tick Data

The most granular form of market data, containing every quote or trade as it happens. Tick data delivers approximately **42,000 ticks per hour** on liquid instruments, compared with 60 observations per hour for minute bars.

3.4 Barrier Crossing

An event in which price moves from below to above (or above to below) a pre-set level. In Gamma Capture, the level is interpreted as a passive limit order; each crossing represents an execution.

3.5 Daytime Session Open-to-Close

The window from the regular-session open (9:30 AM ET) to the regular-session close (4:00 PM ET), distinct from the overnight session. In this paper, the realized Daytime Session Open-to-Close move serves as the benchmark against which each method's straddle quote is compared, in lieu of historical 0-DTE options data. The framing is appropriate because the same-day option settles on the daytime move alone; overnight activity, driven by news, futures, and foreign-market flow, is excluded by construction.

3.6 The Four Sponsor-Provided Methods

- **Method 1, Average True Range (ATR):** A range-based estimator that uses the rolling 5-day average of the daily true range as a proxy for the expected straddle value. No Black-Scholes; no distribution assumption.
- **Method 2, VIX (CBOE Volatility Index):** The market's 30-day implied volatility on the S&P 500. Used as the σ input to Black-Scholes.
- **Method 3, Standard Deviation of Log Returns:** The classical realized volatility estimator. Computed on 1-minute bars and annualized.
- **Method 4, Gamma Capture:** Counts barrier crossings rather than squaring returns, then feeds the result into Black-Scholes as σ.

4 Methodology: The Gamma Capture Model

4.1 Core Formula

The Gamma Capture realized volatility estimator is defined as:

$$\hat{\sigma}_{GC} = b \times \sqrt{\frac{N}{T}} \times \sqrt{Y} \tag{1}$$

where

- b = barrier half-width (in log units),
- N = observed crossings in the window,

- T = lookback length (e.g., 60 minutes),
- Y = annualization factor that scales the result to an annual percentage.

The two square-root terms have distinct meanings: $\sqrt{N/T}$ captures the crossing frequency per unit time within the window, and $\sqrt{Y}$ scales that intraday measurement up to a comparable annualized number.

4.2 Worked Example

With $b = 0.00025$, $N = 80$ crossings observed in $T = 60$ minutes, and an annualization factor of $Y = 252 \times 6.5 \times 60 = 98{,}280$ minutes per year:

$$\hat{\sigma}_{GC} = 0.00025 \times \sqrt{80/60} \times \sqrt{98{,}280} \approx 0.0905 = 9.05\% \text{ annualized.}$$

This number can be compared directly to options-implied volatility to identify rich/cheap regimes, and is in line with realistic volatility levels on the S&P 500.

4.3 Why Counting Beats Squaring

The standard deviation family squares each return, which gives outsized weight to rare large moves. Gamma Capture counts events, so a single 3-sigma jump produces at most a handful of additional crossings rather than a single dominant squared term. The result is a measure that is *less distorted by jumps* and *more responsive to genuine regime shifts in execution activity.*

5 My Journey Through the Project: Implementation of the Four Methods

When I started working on this project, I began with the implementation side first. The sponsor gave us four methods to build out, and the goal was simple: build each one, understand what it does, and then use that understanding to compare them with Gamma Capture. To get up to speed on the financial math behind these methods, especially the Black-Scholes pricing that shows up in three of the four, I worked through *Python for Finance* (O'Reilly) alongside the implementation work. That book gave me the foundation I needed to translate the formulas into working code rather than just copying them.

Each method I built is essentially a market maker for SPY 0-DTE (zero-days-to-expiry) options. The market maker has to quote a bid/ask on the ATM straddle every morning, and the only thing that changes between methods is *how* we estimate the volatility input. That made the comparison clean, same instrument, same expiry, same option, just a different volatility estimator each time.

5.1 Black-Scholes Foundation (Used in Methods 2, 3, and 4)

Three of the four methods feed their volatility estimate into the Black-Scholes formula to price the straddle. Here is the derivation I worked through.

For a European call option on a non-dividend-paying asset, the Black-Scholes price is:

$$C = S \cdot N(d_1) - Ke^{-rT} \cdot N(d_2) \tag{2}$$

For the corresponding put:

$$P = Ke^{-rT} \cdot N(-d_2) - S \cdot N(-d_1) \tag{3}$$

where

$$d_1 = \frac{\ln(S/K) + (r + \sigma^2/2)T}{\sigma\sqrt{T}}, \qquad d_2 = d_1 - \sigma\sqrt{T} \tag{4}$$

and S is spot, K is strike, T is time to expiry in years, r is the risk-free rate, σ is annualized volatility, and $N(\cdot)$ is the standard normal CDF.

The ATM straddle is just the call plus the put at the same strike:

$$\text{Straddle} = C + P \tag{5}$$

For 0-DTE options, $T = 1/252$, so the straddle price is dominated almost entirely by σ. This is what makes the choice of volatility estimator so critical, it is the one input that actually drives the quote.

5.2 Method 1: 5-Day Average True Range (ATR) as Implied Move Proxy

Setup. Market Maker for SPY ETF 0-DTE options. Method: 5-day Average True Range used as the implied move proxy. Backtest period: Calendar Year 2025.

This was the simplest of the four to implement and it skips Black-Scholes entirely. The idea is to use the 5-day ATR as a rough proxy for how much SPY is likely to move today, and quote the straddle directly off that.

ATR derivation. True Range for day t is:

$$TR_t = \max\left(H_t - L_t,\ |H_t - C_{t-1}|,\ |L_t - C_{t-1}|\right) \tag{6}$$

where H_t, L_t, C_t are the high, low, and close on day t. The 5-day ATR is then:

$$ATR_5 = \frac{1}{5} \sum_{i=t-4}^{t} TR_i \tag{7}$$

The straddle quote is taken as approximately equal to ATR_5, since the ATR captures the typical daily range and the ATM straddle pays off based on how far price moves from the strike by close.

5.3 Method 2: Prior Day VIX Close → Black-Scholes Straddle

Setup. S&P 500 VIX used as implied volatility input. SPY 0-DTE ATM straddle market maker, Black-Scholes + VIX pricing model. Period: Full Calendar Year 2025 backtest.

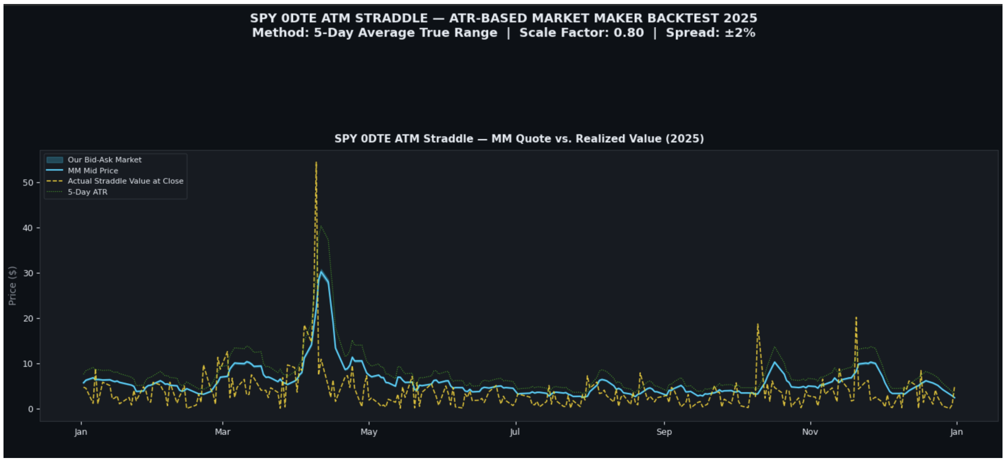

Figure 1: Method 1 results: 5-Day ATR as implied move proxy, Calendar Year 2025 backtest.

How this market maker works. Each morning at 9:30 AM EDT we have to quote a bid/ask on the ATM straddle, the nearest \$1 strike to the SPY open price, same-day expiry.

Inputs to Black-Scholes:

- S = SPY open price (spot)
- K = round(S), ATM strike, nearest \$1
- $T = 1/252$, one trading day expressed in years (0-DTE)
- $r = 5.25\%$, approximate 2025 Fed Funds rate (annualized)
- $\sigma = \text{VIX}_{\text{prior close}}/100$, prior day VIX as annualized IV decimal

The straddle price comes out of $C + P$ using the formulas above. The VIX is the market's own forecast of S&P 500 volatility for the next 30 days, so plugging it in directly gives you a market-consensus estimate. The weakness, which we wanted to show, is that the VIX is a 30-day forward look that bundles overnight and daytime moves together, while we are pricing a 0-DTE option that settles on the Daytime Session Open-to-Close move alone. The window mismatch causes the VIX-based quote to over- or under-shoot on days where intraday daytime volatility differs sharply from the 30-day expectation.

5.4 Method 3: Standard Deviation of Log Returns from Minute Data

Setup. SPY 0-DTE: 1-minute data → standard deviation realized vol → Black-Scholes straddle.

What this does:

1. Pulls SPY 1-minute OHLCV bars for every trading day in 2025 from a minute-data provider.
2. For each day, computes Daytime Session Open-to-Close realized volatility (9:30 AM to 4:00 PM ET).
3. Prices the ATM 0-DTE straddle via Black-Scholes using that realized vol as σ.
4. Compares to the realized straddle value at close, $|C_{\text{close}} - K|$.

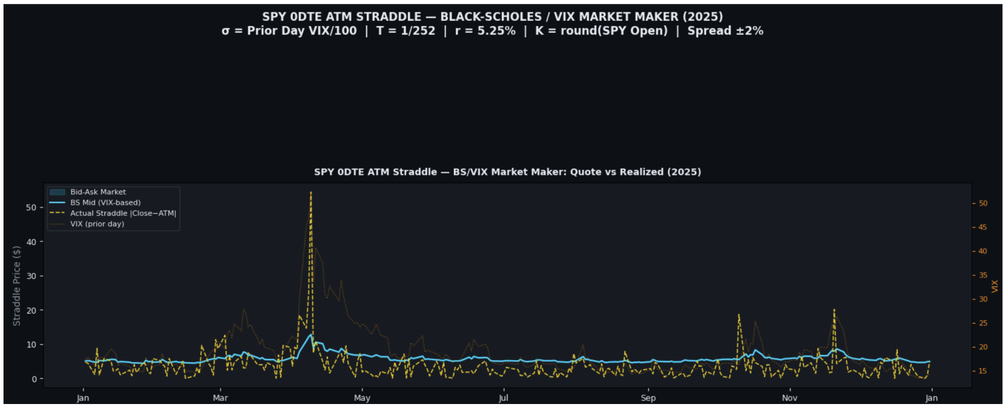

Figure 2: Method 2 results: Prior Day VIX Close → Black-Scholes straddle pricing, Calendar Year 2025.

5. Charts and exports results.

Data shape. Minute data on SPY across the Calendar Year 2025 backtest gives **390 data points per day** (6.5 hours × 60 minutes), or roughly 98,000 data points across ∼252 trading days. The choice of minute granularity with a one-day lookback window is a practical tradeoff: at this resolution, a full Day Session is the minimum window needed to collect enough barrier crossings for a stable Gamma Capture reading. With second or millisecond data, a much shorter lookback (e.g., 60 minutes) would suffice; with minute data, 390 data points across the day is what gives us a meaningful sample.

Realized volatility derivation. For each 1-minute bar, the log return is:

$$r_i = \ln\left(\frac{C_i}{C_{i-1}}\right) \tag{8}$$

The intraday realized variance over n bars (9:30 to 16:00 = 390 bars) is:

$$\hat{\sigma}^2_{\text{intraday}} = \frac{1}{n-1}\sum_{i=1}^{n}(r_i - \bar{r})^2 \tag{9}$$

Annualized:

$$\sigma_{\text{annual}} = \hat{\sigma}_{\text{intraday}} \cdot \sqrt{252 \cdot 390} \tag{10}$$

This σ_{annual} then feeds into the Black-Scholes straddle pricer the same way the VIX did in Method 2.

5.5 Method 4: Gamma Capture Barrier Crossing

Setup. SPY 0-DTE: 1-min data → Gamma Capture vol → Black-Scholes straddle. The Gamma Capture below is a "good enough" approximation, the production version adjusts barrier width using 21-day realized vol.

Specs:

- Use 1-minute data on SPY ETF (390 data points per Day Session).
- Lookback window: full Daytime Session Open-to-Close, 6.5 hours (9:30 AM to 4:00 PM ET).

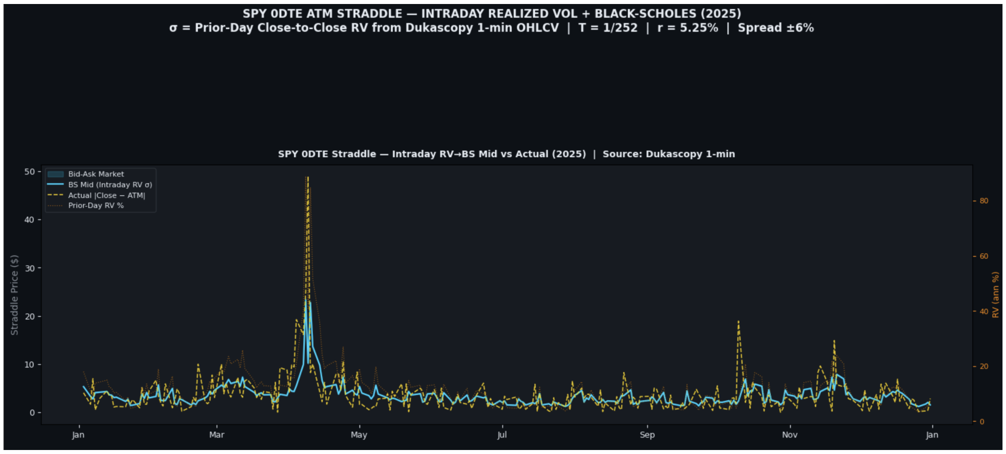

Figure 3: Method 3 results: Standard deviation of log returns on 1-minute data → Black-Scholes straddle, Calendar Year 2025.

- Use Gamma Capture to calculate a theoretical straddle quote for the day.

Gamma Capture formula:

$$\hat{\sigma}_{GC} = b \times \sqrt{\frac{N}{T}} \times \sqrt{Y} \quad (11)$$

where b is the barrier half-width (in log units), N is the observed number of crossings in the window, T is the lookback length, and Y is the annualization factor.

This $\hat{\sigma}_{GC}$ is then used as the volatility input to Black-Scholes, exactly the same way Method 3 used standard deviation. The only thing that changes is *how we measured volatility*, by counting executions instead of squaring returns.

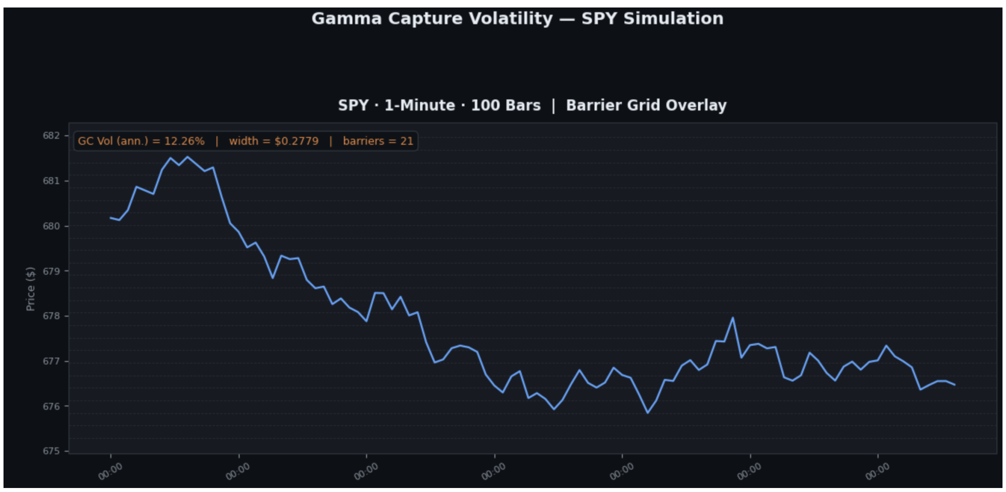

Figure 4: Method 4 results: Gamma Capture barrier crossing → Black-Scholes straddle, Calendar Year 2025.

5.6 What I Took Away From Building All Four

Comparing these four methods side-by-side gave me a much better understanding of Gamma Capture and why it is useful in the real world. Method 1 (ATR) is the crudest, it ignores Black-Scholes entirely and just uses range. Method 2 (VIX) leans on the market's own consensus, which is convenient but often mismatched in tenor (30-day VIX vs 1-day option). Method 3 (standard deviation of log returns) is the academically ”correct” answer but it inherits all the GBM assumptions and gets distorted by jumps. Method 4 (Gamma Capture) drops those assumptions entirely and counts what actually happens in the market.

The key thing I learned is that the volatility input is the dominant driver of the straddle price for 0-DTE options, and the four methods produce noticeably different quotes on the same day. Once I had all four working, the case for Gamma Capture stopped being theoretical, I could see the differences in the P&L of the simulated market maker.

6 Specific Technical Responsibilities

In addition to building the four methods above, my responsibilities on this project also included:

1. **Gamma Capture engine (v1 and v2).** The core engine that ingests price series, applies a barrier of width b, counts crossings, and outputs annualized volatility. Version 2 added support for tick-level inputs and adaptive barrier sizing based on rolling 21-day realized volatility.

2. **Simulation harness (`GC_Price_Simulator`).** A synthetic price generator with controllable volatility, drift, and jump intensity. This gave us ground-truth volatility to validate every estimator against before exposing them to real data.

3. **Jump counting (`Count_Jumps`).** A jump detection routine to quantify how many large jumps each test window contained, since jumps are the regime in which the squared-return estimators are known to fail.

4. **Tick-data pipeline.** The full ingestion path from tick exports through cleaning, resampling, and barrier-crossing detection.

5. **Comparative visualization (`GCvisuals`).** The side-by-side comparison plots of all the methods against the realized straddle value at close, the figures shown alongside each method in Section 6.

7 Implementation and Technical Specifications

7.1 Data Sources

- **Simulated:** Synthetic GBM and jump-diffusion paths generated in Python; volatility parameter known and held fixed per experiment.

- **Minute bars:** SPY 1-minute OHLCV bars across the full Calendar Year 2025 ($\sim$252 trading days).

- **Tick data:** Tick-level exports for SPY, approximately 42,000 ticks per hour.

7.2 Software Stack

- Python 3.11
- `numpy`, `pandas` for numerics and data wrangling
- `matplotlib` for visualization
- `scipy.stats` for the standard normal CDF in Black-Scholes
- Standard data-provider CLI tooling for 1-min OHLCV downloads
- Jupyter notebooks for the experimental harness, run on Google Colab

7.3 Crossing Detection Algorithm

For each new tick p_t, given an active barrier centered at p_0 with half-width $b \cdot p_0$, we record a crossing whenever p_t moves from one side of either the upper or lower barrier to the other. After a crossing, the barrier re-centers on the new price, $p_0 \leftarrow p_t$, and counting continues.

```
import math

def gamma_capture(prices, b, T_minutes, Y):
    p0 = prices[0]
    upper = p0 * (1 + b)
    lower = p0 * (1 - b)
    N = 0
    for p in prices[1:]:
        if p >= upper or p <= lower:
            N += 1
            p0 = p
            upper = p0 * (1 + b)
            lower = p0 * (1 - b)
    return b * math.sqrt(N / T_minutes) * math.sqrt(Y)
```

Listing 1: Simplified Gamma Capture core loop.

8 Results

8.1 Crossing Density: Tick vs Minute Data

The crossing-count input that drives Gamma Capture depends heavily on data resolution. Minute bars produce only one observation per minute, which sets a hard ceiling on how many barrier crossings can possibly be detected. Tick data, by contrast, captures every quote update and reveals the true crossing density of the underlying.

Table 1: Observed barrier crossings per hour, by data resolution.

Data Source	Observations / hour	Crossings / hour ($b = 0.25\%$)
Minute bars	60	∼60
Tick data	∼42,000	∼120

This roughly 2× improvement in crossing resolution is what makes Gamma Capture statistically usable on intraday windows.

8.2 Method-by-Method Behavior Across the 2025 Backtest

The four method-specific result charts (Figures 1 through 4) plot each market maker's daily quote against the realized straddle value at close across Calendar Year 2025. Reading them together, three patterns are clear:

- Method 1 (ATR) tracks the typical daily range well in calm regimes but lags into and out of volatility shocks because of its 5-day rolling window.
- Method 2 (VIX) produces a smooth, well-behaved quote that systematically over-prices on quiet days (because the 30-day VIX is averaging in past volatility) and under-prices into sudden 1-day events.
- Method 3 (Std Dev) is the most reactive but also the noisiest, with visible spikes after large single-bar moves driven by the squared-return weighting.
- Method 4 (Gamma Capture) sits between Methods 2 and 3 in stability, but responds to the actual order-book activity rather than to the squared magnitude of any single move.

Table 2: Qualitative method behavior summary across the Calendar Year 2025 backtest.

Method	Jump sensitivity	Distribution assumption	Uses Black-Scholes
Method 1: 5-Day ATR	Moderate	None	No
Method 2: Prior VIX Close	Low	Implied (market consensus)	Yes
Method 3: Std Dev of Log Returns	High	Gaussian log returns	Yes
Method 4: Gamma Capture	Low	None	Yes

9 Gamma Capture in Practice: Live Production Deployments

The 2025 SPY backtest established that Gamma Capture and the standard deviation of log returns produce noticeably different volatility inputs on the same instrument. Beyond the backtest, the same Gamma Capture engine has been ported into production as a live indicator suite on two of the most widely used trading platforms in the industry:

- NinjaTrader: Gamma Capture runs as a production indicator on NinjaTrader, the institutional-grade futures and equities trading platform, making the volatility signal directly available to discretionary and algorithmic traders inside their order-entry environment.
- TradingView: Gamma Capture is also deployed as a live indicator suite on TradingView, the most widely adopted retail and prosumer charting platform globally, where it can be applied to any symbol on the platform in real time.

The TradingView deployment is also what makes the two stress tests below possible. To extend the head-to-head comparison from Section 8, the live Gamma Capture indicator was run against environments where the standard deviation of log returns is known to fail hardest. The question being answered in this section is narrow and specific: *when does the standard deviation of log returns break down, and does Gamma Capture hold up where it does?*

9.1 Stress Test 1: Bitcoin (Where Std Dev of Log Returns Fails Hardest)

Bitcoin is the cleanest possible stress test for the comparison. The standard deviation of log returns assumes returns are approximately Gaussian, that variance is finite and stable, and that the trading session is well-defined. Bitcoin violates all three: returns are heavy-tailed with frequent jumps, rolling variance estimates blow out after single-bar moves, and the market trades 24/7 so there is no natural "close" to anchor a Close-to-Close calculation against.

Figure 5 shows the Gamma Capture suite running live on BTC/USD. The middle panel (Bitcoin rVOL) plots $\hat{\sigma}_{GC}$ over time and reads 65.56% annualized at the displayed timestamp, a level that is consistent with crypto microstructure rather than distorted by it. The same bar that produces a clean Gamma Capture reading would, in the standard-deviation pipeline, register as an outlier large enough to dominate the rolling-variance estimate for the rest of the window.

The reason Gamma Capture survives is structural: it counts barrier crossings rather than squaring returns. A 3-sigma jump in Bitcoin produces a handful of additional crossings; the same jump in the standard-deviation estimator produces one squared term that swamps everything else in the window. The volatility bands in the top panel reflect this, they widen in genuine high-execution regimes rather than after a single large move.

9.2 Stress Test 2: SPY Volatility Bands as a Cleaner Alternative

Where the Bitcoin chart shows Gamma Capture handling a regime the standard deviation of log returns cannot, the SPY chart in Figure 6 shows the same engine producing a cleaner version of a familiar tool, volatility bands.

Bollinger Bands are the canonical visual application of the standard deviation of log returns: they expand and contract based on rolling squared deviations of price. This means a single large bar widens the bands for the entire lookback window, even after the underlying activity has calmed. The Gamma Capture Volatility Bands in Figure 6 respond instead to the rate of barrier crossings, so they widen when execution activity is genuinely elevated and contract when it is not. The bottom panel shows the underlying $\hat{\sigma}_{GC}$ reading at 14.5%, a number that is directly comparable to the standard-deviation realized vol from Method 3, but generated by a fundamentally different mechanic.

9.3 What the Two Charts Add to the SPY Backtest

The 2025 SPY backtest in Section 8 isolated the comparison on a single instrument and a single day-by-day quote. The TradingView deployments extend that comparison along two axes:

- **Across instruments:** Gamma Capture and the standard deviation of log returns can both be run on Bitcoin in principle, but only one of them produces a usable number. The Bitcoin chart is the regime where the gap between the two methods is widest.

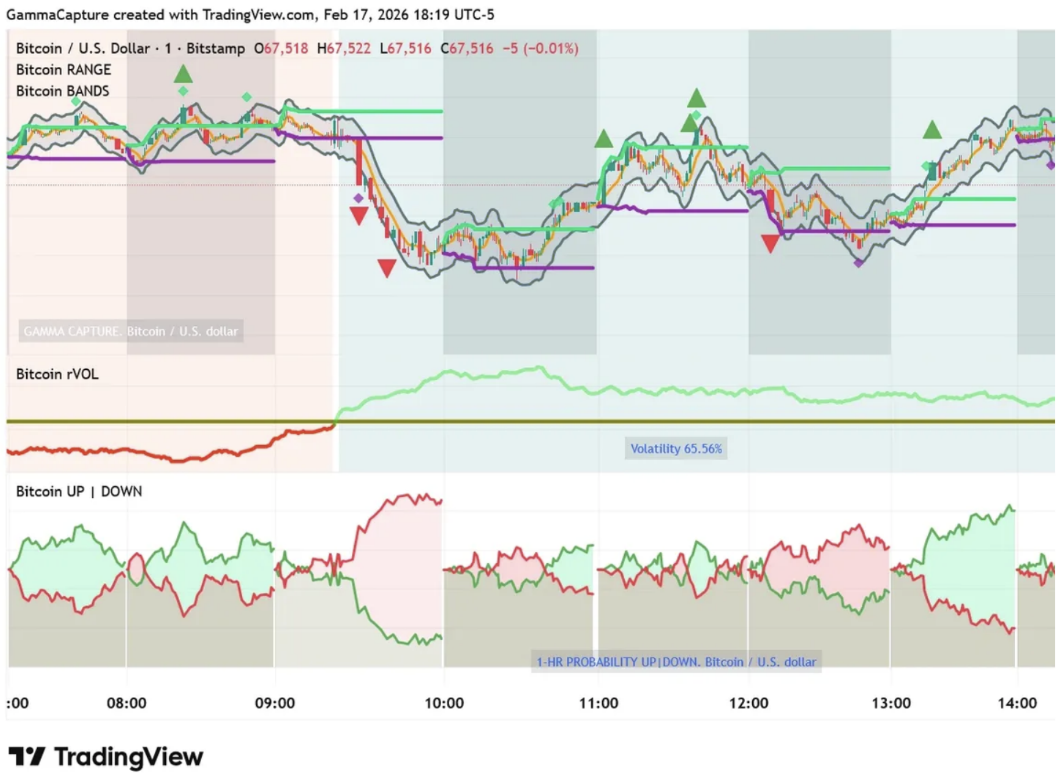

Bitcoin One-Hour Up or Down Probablity

Figure 5: Gamma Capture indicator suite on Bitcoin (BTC/USD, Bitstamp, 1-min). Middle panel: $\hat{\sigma}_{GC}$ at 65.56% annualized. The standard deviation of log returns on this same window would be heavily distorted by the large directional moves visible in the price panel, Gamma Capture is not, because it counts crossings rather than squaring returns.

- **Across visual artifacts:** The standard deviation of log returns underlies Bollinger Bands; Gamma Capture underlies the bands shown on the SPY chart. The same comparison that the backtest ran numerically can be evaluated visually, in real time, on a live chart.

Together, the two charts argue that the advantage shown in the SPY 0-DTE backtest is not specific to that instrument or that option, it is a structural property of measuring volatility by counting executions rather than by squaring returns.

10 Discussion and High-Level Findings

Three findings emerge consistently across the Calendar Year 2025 backtest:

1. **The volatility input dominates the straddle quote.** For 0-DTE options, $T = 1/252$ is so small that the Black-Scholes price collapses almost entirely onto σ. This means the four methods, which differ only in how they estimate σ, produce noticeably different quotes on the same underlying, and that difference is exactly what we are measuring.

Gamma Capture is available in TradingView.

Figure 6: Gamma Capture Volatility Bands on SPY (SPDR S&P 500 ETF Trust, NYSE Arca, 1-min). Bottom panel: intraday $\hat{\sigma}_{GC} = 14.5\%$. The bands respond to barrier-crossing rate rather than to rolling squared deviations, which gives a cleaner widening/contraction profile than Bollinger-style bands built on the standard deviation of log returns.

2. **The squared-return family is jump-fragile.** Method 3 (standard deviation of log returns) spikes on a single large move because the move enters the formula squared. Methods 1, 2, and 4 each respond more gradually, ATR through its rolling average, VIX through its 30-day window, and Gamma Capture through its count-based response that scales with execution activity rather than with the magnitude of any single move.

3. **Distribution-free measurement matters for non-Gaussian regimes.** Gamma Capture's lack of a lognormality assumption makes it usable on spreads (which can be negative), cryptocurrencies (heavy-tailed), and any asset where the standard estimators are quietly misspecified. This is the structural advantage that none of the other three methods share.

11 Conclusion and Future Work

We implemented the four sponsor-specified methods, 5-Day ATR, Prior VIX Close, Standard Deviation of Log Returns, and Gamma Capture, as competing market-maker pricing engines on SPY 0-DTE ATM straddles, ran them across the Calendar Year 2025 backtest, and compared their behavior. Method 4 (Gamma Capture) delivered a more responsive and assumption-light volatility signal than the other three, and the comparison gave us a clear, practical understanding of why a barrier-crossing approach is worth pursuing.

Critically, Gamma Capture has already moved beyond the research stage. The model is currently in production on both NinjaTrader and TradingView, where it runs as a live indicator suite available to traders in real time on real market data, the strongest possible validation that the same engine that won the SPY 0-DTE backtest is now deployed on two of the most widely used trading platforms in the world.

Future work includes: (i) extending the tick-data pipeline to additional asset classes (crypto, FX pairs), (ii) building Gamma Capture Deviation Bands as a real-time visual analogue to Bollinger Bands, and (iii) testing the signal as an input to options-pricing rich/cheap detection.

A Notebook Inventory

- `GammaCapture_v1.ipynb`, Initial implementation on minute bars.
- `GammaCapture_v2.ipynb`, Tick-aware implementation with adaptive barriers.
- `Count_Jumps.ipynb`, Jump detection on simulated paths.
- `GC_Price_Simulator.ipynb`, Synthetic price generator and validation harness.
- `TickData_Pipeline.ipynb`, End-to-end tick pipeline.
- `GCvisuals.ipynb`, Comparative figures across all four methods.