From Barrier Crossings to Terminal Distributions: A Skellam-Based Options Pricing Framework for 0-DTE Markets

Louis J. Pellathy

Founder, Gamma Capture

gammacapture.com

June 2026

Abstract

Prior work established that barrier-crossing frequency is a superior volatility estimator to squared log-returns for 0-DTE index options. This paper extends that framework by replacing the Black-Scholes pricing layer entirely. Rather than feeding a barrier-crossing volatility signal into a lognormal pricing formula, we derive a closed-form pricing model whose terminal distribution is Skellam — the difference of two independent Poisson processes. Each Poisson process represents directional barrier crossings: up-crossings and down-crossings. The resulting pricer generates a full implied volatility surface from a single empirical parameter, the crossing asymmetry ratio $\varphi = N^+/N^-$, without any options data, implied vol fitting, or smile calibration. A Calendar Year 2025 backtest on SPY 0-DTE options using minute data confirms that the Skellam pricer achieves near-zero mean error across 655 ITM expiries (excluding two black-swan days), with mean absolute error of $1.87. The vol skew emerges structurally from φ. Black-Scholes machinery is entirely replaced by a discrete probability distribution. While academic peer review is useful, options market making is the real-life is proof that the Gamma Capture framework works.

1. Introduction

The zero-days-to-expiry (0-DTE) options market has grown to dominate index options volume, exceeding $1.2 trillion in daily notional on SPY alone as of May 2025. This growth has exposed a structural deficiency in the dominant pricing framework. Black-Scholes assumes price evolves according to Geometric Brownian Motion (GBM), a continuous diffusion process with lognormally distributed returns. As time-to-expiry approaches zero, the Black-Scholes formula collapses to a point mass at the current price, producing systematically wrong option values precisely when the market is most active.

Earlier work by the author benchmarked four volatility estimation methods — Average True Range, prior-day VIX, standard deviation of log returns, and Gamma Capture barrier-crossing volatility — as competing market-maker pricing engines on SPY 0-DTE straddles across the Calendar Year 2025 backtest. Gamma Capture outperformed the alternatives by producing a volatility signal responsive to order-book execution activity rather than squared return magnitude. However, all four methods retained Black-Scholes as the pricing layer.

This paper removes Black-Scholes entirely. The terminal price distribution is not assumed to be lognormal. Instead, it is derived directly from the Poisson crossing process: the terminal net displacement is the difference of two independent Poisson counts, a Skellam distribution. The Skellam distribution has fat tails relative to the Gaussian, retains probability mass at the wings as $T \to 0$, and generates a structural vol skew through the crossing asymmetry ratio φ. The result is a framework grounded in observable microstructure rather than statistical assumptions.

2. Prior Literature

The barrier-crossing approach to volatility measurement has antecedents in several independent research lines.

Rice (1944) established the foundational mathematics of crossing rates for continuous random processes. His expected crossing rate formula — the number of times a stationary Gaussian process crosses a fixed level per unit time — provides the theoretical basis for interpreting crossing frequency as a volatility signal. The Gamma Capture framework applies this principle to discretized price data with a fixed barrier grid, making Rice's continuous result operational on market microstructure.

Sinclair and Merrill addressed the connection between barrier crossings and volatility estimation in their paper "Volatility Estimation via First Exit Times." Their work frames volatility in terms of the time required for price to first cross a barrier level, connecting crossing frequency to the parameters of an underlying diffusion.

The Gamma Capture rvol formula:

$$\boldsymbol{\sigma = b \times \sqrt{(N/T)} \times \sqrt{Y}}$$

where b is barrier width, N is crossing count, T is the lookback window length, and Y is the annualization factor. It is structurally consistent with this first-exit-time framework. Gamma Capture vol can be used to calculate an options market makers hedge P&L on a constant gamma book.

Options Market Making P&L:

Options P&L = [σ²_implied_vol − σ²_gamma_capture] * [Γ x Time to Expire]

Andersen, Dobrev, and Schaumburg (2014) developed duration-based volatility estimation, using the time between threshold crossings as the primary measurement object rather than return magnitudes. Their estimator is robust to microstructure noise and jumps for the same structural reason Gamma Capture is: neither method squares returns, so large single-bar moves do not disproportionately inflate the variance estimate.

Perz (2024) documented the profitability of selected 0-DTE index options strategies across empirical backtests, confirming the practical significance of accurate 0-DTE pricing and the material cost of mispricing at short tenors. The consistent underperformance of Black-Scholes-based strategies in Perz's data is consistent with the structural failure mode this paper addresses: a lognormal terminal distribution that collapses to a point mass at $T = 0$.

The present paper differs from all of the above in one critical respect: it does not use crossing-based vol as an input to Black-Scholes. The Poisson crossing process is the price process. The terminal distribution is derived from it directly.

3. The Gamma Capture Pricing Model

3.1 The Crossing Process

Let price evolve on a discrete barrier grid with spacing b dollars. Each time price crosses a barrier level, it is counted as either an up-crossing (N^+) or a down-crossing (N^-). The crossing counts over a trading day are modeled as independent Poisson random variables:

$$N^+ \sim \text{Poisson}(\lambda \cdot T / 2), \quad N^- \sim \text{Poisson}(\lambda \cdot T / 2)$$

where λ is the total crossing intensity per day and $T = 1$ for a 0-DTE option. The equal split of intensity between up and down crossings reflects the symmetric baseline; asymmetry is introduced through φ.

3.2 Terminal Distribution

The terminal net displacement in barrier units is $X = N^+ - N^-$. The difference of two independent Poisson random variables follows a Skellam distribution:

$$X \sim \text{Skellam}(\mu^+, \mu^-), \quad \mu^+ = \mu^- = \lambda(K) \cdot T / 2$$

The Skellam PMF is:

$$P(X = k; \mu^+, \mu^-) = \exp(-(\mu^+ + \mu^-)) \cdot (\mu^+/\mu^-)^{(k/2)} \cdot I_{|k|}(2\sqrt{\mu^+\mu^-})$$

where $I_{|k|}$ is the modified Bessel function of the first kind. Terminal price is $S\_T = S_0 + b \cdot X$, mapping integer net crossings to dollar-denominated price outcomes.

3.3 Strike-Dependent Crossing Intensity

Crossing intensity varies by strike. OTM options require more crossings to reach payoff, and the empirically observed crossing asymmetry generates a structural skew. The intensity function is:

$$\boldsymbol{\lambda(K) = \lambda\_ATM + K\_mult \times |K - ATM| \times (1/\varphi \text{ puts}, \ \varphi \text{ calls})}$$

where λ_ATM is the ATM crossing intensity (calibrated empirically as 130 crossings/day for SPY 2025), K_mult controls the slope of intensity with moneyness, and $\varphi = N^+/N^-$ is the 5-day rolling crossing asymmetry ratio. When $\varphi < 1$, down-crossings dominate and put intensity auto-amplifies. When $\varphi > 1$, the reverse holds. The vol skew is produced without any options data.

3.4 Option Pricing Formula

Option value is the probability-weighted sum of payoffs across all terminal states:

$$\textbf{Call} = \Sigma_{i=-I\_max}^{+I\_max} \max(S_0 + b \cdot i - K, 0) \times P\_Skellam(i; \mu^+, \mu^-)$$

$$\textbf{Put} = \Sigma_{i=-I\_max}^{+I\_max} \max(K - S_0 - b \cdot i, 0) \times P\_Skellam(i; \mu^+, \mu^-)$$

No present value discounting is applied at 0-DTE tenors. The summation range I_max = 50 covers ±50 barrier widths from current price, capturing effectively all Skellam probability mass at empirical λ values.

4. Empirical Calibration

4.1 Data

The backtest uses S&P 500 SPY ETF minute-bar data for Calendar Year 2025 (239 trading days, 93,449 M1 bars). Daily bars provide the 21-day rolling realized volatility used to compute the dynamic barrier width b(t). The barrier grid is anchored to the daily open price.

4.2 Parameter Estimates

Empirical crossing statistics for SPY 2025:

λ_daily: mean 126.7, median 113, max 420 (two black-swan days at λ > 280)

φ_calls (5-day rolling): mean 0.90, range [0.40, 1.50]

φ_puts (5-day rolling): mean 1.17, range [0.67, 2.48]

b(t): dynamic, mean ~$0.35, driven by 21-day realized vol inversion

Calibrated pricing parameters: λ_ATM = 130, K_MULTIPLY_CALL = 4, K_MULTIPLY_PUT = 8. These were selected to minimize mean pricing error across the 2025 backtest without reference to options prices.

5. Backtest Results

5.1 Setup

For each of 233 trading days with sufficient φ data, the model priced calls and puts at moneyness ±1 through ±5 (integer strikes relative to ATM). Realized payoff is the intrinsic value at close. Error is defined as gc_price − payoff. Two days (November 20 and October 10, 2025) are excluded as black-swan events: both had λ > 280 and realized moves exceeding 4σ. These days are retained in the dataset for vol calibration purposes but excluded from pricing error analysis.

5.2 Results

Across 655 ITM expiries (ex black-swan):

Overall mean error: −$0.02 (effectively zero)

Overall mean absolute error: $1.87

Error by side and moneyness (call / put):

Calls: +0 (ATM) to +5 moneyness → mean errors −$0.10 to +$0.22, std $1.74–$2.29

Puts: −1 moneyness → +$0.75; −5 moneyness → −$0.87

The residual put slope (deeper OTM puts slightly underpriced) reflects the limits of a single φ parameter to capture the full skew term structure. Calls are nearly unbiased at all moneyness levels. The mean error of −$0.02 compares favorably to any Black-Scholes implementation, which requires a separately fitted IV at each strike to achieve comparable accuracy.

5.3 Black-Swan Days

November 20 ($\lambda = 287$, realized move $16–$20 ITM on OTM puts) and October 10 ($\lambda = 300$, similar profile) produced errors of −$12 to −$16. These are not model failures — no model calibrated on prior-day data prices a 4σ intraday move correctly at open. Their exclusion from error analysis is appropriate; their retention in the vol calibration data is equally appropriate.

6. Discussion

The result that a Skellam pricer achieves near-zero mean error without using any options data carries a strong implication: the vol surface implied by options markets is, on average, consistent with the structure of the underlying crossing process. φ is not a free parameter tuned to fit the smile — it is measured directly from price crossings on the prior day and rolled forward. The fact that it generates a structural skew that matches realized payoffs is evidence that the smile is a consequence of crossing asymmetry, not an independent phenomenon requiring separate calibration.

The comparison to Black-Scholes is instructive. BSM requires a separately implied vol at every strike and expiry. The implied vol surface is refitted daily. The vol smile is not produced by the model — it is patched onto it by hand. In the Gamma Capture framework, one parameter (φ) replaces the entire smile calibration machinery. The vol skew emerges from the data.

The Skellam distribution retains tail probability as $T \to 0$ because it is a count distribution, not a continuous diffusion. As the number of expected crossings decreases with time, the distribution does not collapse to a point mass — it assigns finite probability to small positive net crossing counts. This is the structural fix for the 0-DTE failure mode in Black-Scholes.

7. Conclusion

We have presented a complete options pricing framework built on the Poisson barrier-crossing process, with no Black-Scholes machinery retained. The terminal distribution is Skellam. The vol surface is structural. The calibration requires no options data. A 2025 SPY backtest confirms near-zero mean error across 655 ITM expiries.

The framework is currently in production on NinjaTrader and TradingView as the Gamma Capture indicator suite. Institutional applications include intraday vol arb (identifying mispriced wings from the structural surface), execution algo integration, and extension to crypto and energy options where lognormality assumptions are most severely violated.

Future work: tick-level calibration of λ (current implementation uses minute bars, which undercount crossings by approximately 2×); multi-day term structure extension; and academic publications.

A note on validation: this paper describes what is already happening. Every delta hedge executed by a market maker is a barrier crossing of width b. Every trade execution is a crossing event. Market makers count N automatically through their fill blotter. The global options market is a continuous, real-money empirical validation of the Gamma Capture framework — operating at scale, every trading day, across every listed strike and expiry. Market Makers place limit order bids and offers (up/down barriers) each day to gamma hedge. While academic peer review is useful, options market making is the real-life is proof.

References

Andersen, T.G., Dobrev, D., & Schaumburg, E. (2014). Duration-Based Volatility Estimation. Federal Reserve Bank of New York Staff Reports.

Perz, P. (2024). Profitability of Selected 0-DTE Index Options Strategies. Faculty of Management, Rzeszow University of Technology.

Rice, S.O. (1944). Mathematical Analysis of Random Noise. Bell System Technical Journal, 23(3), 282–332.

Sinclair, E. & Merrill, C. Volatility Estimation via First Exit Times. Unpublished manuscript.

Pellathy, L.J. (2026). Gamma Capture: Barrier-Crossing Volatility. Lehigh University MFE Collaboration Working Paper.