

Introducing **floe**: Static and Real-time Calculation of Options Pricing, Flows, and Dealer Gamma, Vanna, and Charm Exposures via Broker Data

Full Stack Craft LLC

March 4, 2026

Repositories: TypeScript (**floe**) — Go (**floe-go**)

Contents

1	Motivation	4
2	Introducing floe	4
3	Paper Structure	4
4	Real-time Dealer Exposure Calculations	5
4.1	Initialization Phase: Capturing Open Interest	5
4.2	Implied Volatility Surface Construction	5
4.3	Total Variance Smoothing	6
4.4	Dealer Position Assumption	6
4.5	Calculate Greeks for Each Option	6
4.6	Exposure Formulas	6
4.7	Total Exposures	7
4.8	Real-Time Update Process	7
4.9	Live Open Interest Tracking	7
4.10	Minimum Broker Requirements	8
5	Black-Scholes Greeks Calculation	8
5.1	Core Parameters	9
5.2	First-Order Greeks	9
5.3	Second-Order Greeks	9

6	Implied Probability Distribution	10
6.1	Theoretical Foundation	10
6.2	Numerical Implementation	10
7	Exposure-Adjusted Implied Probability Distribution	10
7.1	Theoretical Motivation	11
7.2	Adjustment Framework	11
7.3	Gamma Modifier: Kurtosis Adjustment	11
7.4	Vanna Modifier: Tail Adjustment	12
7.5	Charm Adjustment: Mean Shift	13
7.6	Summary Statistics	14
7.7	Configuration by Market Regime	14
7.8	Interpretation and Application	15
8	Hedge Flow Analysis: Impulse Curve and Charm Integral	15
8.1	Motivation: Dollar Exposures as Sufficient Statistics	15
8.2	The Hedge Impulse: Combining Gamma and Vanna	16
8.3	Regime Derivation from the IV Surface	17
8.4	Deriving the Spot-Vol Coupling from the IV Surface	18
8.5	Mapping from Strike Space to Price Space	18
8.5.1	Adaptive Kernel Width	19
8.6	Curve Analysis	19
8.6.1	Zero Crossings	20
8.6.2	Basins and Peaks	20
8.6.3	Directional Asymmetry	20
8.6.4	Regime Classification	21
8.7	Charm Integral: Unconditional Time Pressure	21
8.8	PressureField: Reachable Stability and Acceleration Zones	22
8.9	Real-Time Recalculation on Open Interest Changes	23
9	Model-Free Implied Volatility and Realized Volatility	24
9.1	Model-Free Implied Volatility	24
9.1.1	Strike Spacing	25
9.1.2	Single-Expiration Mode	25
9.1.3	Two-Term Interpolation	25
9.2	Tick-Based Realized Volatility	26
9.2.1	Design: Tick-Level, Not Windowed	26
9.3	The Intraday Variance Risk Premium	26
9.4	Term Structure Signals	27
9.5	Vol Response Residual Model	27
10	Shares Needed to Cover	28

11 Future Work	29
12 Summary	29
A Repository	31

1 Motivation

In recent years, the interest and involvement of retail traders in global markets has surged, including the use of more complex instruments such as options and futures. For the most part, many brokerages have embraced this secular trend by offering their customers a way to stream real-time data directly on their personal computers. However, the standard interface between each brokerage varies widely, and retail traders are left with challenging documentation and inconsistent APIs to work with. We attempt to alleviate this problem by introducing a zero-dependency analytics library that can compute a variety of both static and real-time options pricing and flow data, including dealer gamma, vanna, and charm exposures from streaming broker data, with both TypeScript and Go implementations.

2 Introducing floe

We introduce `floe`: a zero-dependency TypeScript library for computing real-time options flow data, including dealer gamma, vanna, and charm exposures that can be derived purely from real-time streamed broker data alone. We also provide `floe-go`, a companion Go port for backend and server-side analytics workflows.

The mathematical foundation of `floe` rests upon the seminal work of Black and Scholes [1] and Merton [2], whose option pricing framework enables the computation of the Greeks used throughout this paper.

3 Paper Structure

The remainder of this paper is organized as follows. Section 4, the main focus of this paper, describes our process of calculating in real-time a smoothed implied volatility surface, live open interest, and the dealer exposures derived from them. Subsequent sections detail the Black-Scholes Greeks calculations, implied probability distributions, and the exposure-adjusted PDF framework. Section 8 introduces hedge flow analysis: the hedge impulse curve for instantaneous dealer response modeling and the charm integral for time-decay pressure quantification. Section 9 develops model-free implied volatility using the variance swap methodology and tick-based realized volatility for intraday variance risk premium analysis. All following sections detail additional analytics that `floe` can compute.

4 Real-time Dealer Exposure Calculations

The real-time dealer exposure calculation consists of two phases: an initialization phase that captures open interest at market open ($t = 0$), and a continuous update phase that recalculates exposures as new spot prices and option quotes arrive. These calculations can be done over a unified interface regardless of the given broker or data source.

4.1 Initialization Phase: Capturing Open Interest

At or before market open, we fetch the complete option chain \mathcal{O} for the underlying symbol. Each option $o \in \mathcal{O}$ contains:

$$o = \langle K, T, \phi, \text{bid}, \text{ask}, \text{OI}_0 \rangle$$

Where:

- K = strike price
- T = expiration timestamp (milliseconds)
- $\phi \in \{\text{call}, \text{put}\}$ = option type
- bid, ask = current bid/ask prices
- OI_0 = open interest at $t = 0$

The market context includes the current spot price S_0 , risk-free rate r , and dividend yield q .

4.2 Implied Volatility Surface Construction

Before calculating the actual dealer exposures, we first need to construct a smooth implied volatility (IV) surface for each expiration T . These per-strike IVs later have a direct affect on the vanna exposure calculation, but also through the gamma and charm calculations themselves via the greeks formulas.

For each option with observed market price P_{mkt} , we solve for σ_{IV} :

$$\text{BS}(S, K, \tau, \sigma_{\text{IV}}, r, q, \phi) = P_{\text{mkt}} \quad (1)$$

Using bisection search with bounds $\sigma \in [0.0001, 5.0]$ (0.01% to 500% volatility).

4.3 Total Variance Smoothing

To ensure arbitrage-free and smooth IV surfaces, we apply total variance smoothing:

1. Convert IV to total variance: $w(K) = \sigma^2\tau$
2. Apply cubic spline interpolation to $w(K)$
3. Enforce convexity via convex hull projection
4. Convert back to IV: $\sigma_{\text{smooth}}(K) = \sqrt{w(K)/\tau}$

The convexity constraint ensures no calendar spread arbitrage exists in the surface.

4.4 Dealer Position Assumption

We assume dealers are net short options (standard market-maker hedging assumption):

- **Short calls:** Dealers sold calls to retail buyers
- **Long puts:** Dealers bought puts from retail sellers (equivalently, short put exposure is negative)

4.5 Calculate Greeks for Each Option

For each option o with strike K , expiration T , and smoothed IV $\sigma_{\text{IV}}(K)$, we compute the following Greeks using the Black-Scholes-Merton formulas (see Section 5).

We now have all necessary components to compute dealer exposures.

4.6 Exposure Formulas

For each strike K with call open interest OI_C and put open interest OI_P :

Gamma Exposure (GEX):

$$\text{GEX}_K = (-\text{OI}_C \cdot \Gamma_C + \text{OI}_P \cdot \Gamma_P) \cdot (S \cdot 100) \cdot S \cdot 0.01 \quad (2)$$

Because the share multiplier and 1% move resolve to unity, this simplifies to:

$$\text{GEX}_K = (-\text{OI}_C \cdot \Gamma_C + \text{OI}_P \cdot \Gamma_P) \cdot S^2 \quad (3)$$

Where Γ_C and Γ_P are the gamma of the call and put contracts at that expiration and strike, respectively.

Vanna Exposure (VEX):

$$\text{VEX}_K = (-\text{OI}_C \cdot \text{Vanna}_C + \text{OI}_P \cdot \text{Vanna}_P) \cdot (S \cdot 100) \cdot 0.01 \quad (4)$$

Reducing constants:

$$\text{VEX}_K = (-\text{OI}_C \cdot \text{Vanna}_C + \text{OI}_P \cdot \text{Vanna}_P) \cdot S \quad (5)$$

Where Vanna_C and Vanna_P are the vanna of the call and put contracts at that expiration and strike, respectively. The factor 0.01 converts the measure to dollars per 1 volatility point ($\Delta\sigma = 0.01$).

Charm Exposure (CEX):

$$\text{CEX}_K = (-\text{OI}_C \cdot \text{Charm}_C + \text{OI}_P \cdot \text{Charm}_P) \cdot (S \cdot 100) \quad (6)$$

The factor of 100 accounts for contract multiplier. Since charm is computed as delta decay per day, canonical CEX is dollars per 1 day of time passage (without additional τ scaling).

4.7 Total Exposures

Sum across all strikes for each expiration:

$$\text{GEX}_{\text{total}} = \sum_K \text{GEX}_K \quad (7)$$

$$\text{VEX}_{\text{total}} = \sum_K \text{VEX}_K \quad (8)$$

$$\text{CEX}_{\text{total}} = \sum_K \text{CEX}_K \quad (9)$$

$$\text{Net Exposure} = \text{NEX}_{\text{total}} = \text{GEX}_{\text{total}} + \text{VEX}_{\text{total}} + \text{CEX}_{\text{total}} \quad (10)$$

4.8 Real-Time Update Process

The system subscribes to streaming quote data from brokers. On each update event:

4.9 Live Open Interest Tracking

When intraday trade data is available, we estimate live open interest:

$$\text{OI}_{\text{live}}(t) = \text{OI}_0 + \sum_{i=1}^n \delta_i \quad (11)$$

Algorithm 1 Real-Time Exposure Update

```
1: Input: New quote event (spot price  $S'$  or option quote)
2: Update spot price  $S \leftarrow S'$ 
3: Recalculate IV surface  $\Sigma$  for expiration  $T$  if option quote received
4: Update live open interest if trade data available
5: for each expiration  $T$  do
6:   for each strike  $K$  do
7:      $\sigma \leftarrow \text{getIVForStrike}(\Sigma, T, K)$ 
8:      $\tau \leftarrow (T - \text{now})/\text{MS\_PER\_YEAR}$ 
9:     Compute  $\Gamma$ , Vanna, Charm using updated  $S, \sigma, \tau$ 
10:    Compute  $\text{GEX}_K, \text{VEX}_K, \text{CEX}_K$ 
11:   end for
12:   Aggregate total exposures for expiration  $T$ 
13: end for
14: Output: Updated exposure metrics
```

Where δ_i represents the estimated OI change from trade i , inferred by comparing trade price to NBBO:

- Trade at ask \Rightarrow buyer-initiated \Rightarrow potential OI increase
- Trade at bid \Rightarrow seller-initiated \Rightarrow potential OI decrease or close

4.10 Minimum Broker Requirements

Note for this process to function effectively, brokers must provide a minimum:

- Real-time streaming quotes for underlying and options
- Open interest data before or at market open
- Trade prints with timestamps to estimate live OI changes

`floe` itself could potentially be used to do the rest of all calculations: IV surface construction, Greeks calculation, exposure aggregation, and real-time updates.

5 Black-Scholes Greeks Calculation

For any option, one can compute Greeks with `floe` using the Black-Scholes-Merton model with continuous dividend yield [1, 2].

5.1 Core Parameters

Given spot S , strike K , time to expiry τ (in years), volatility σ , risk-free rate r , and dividend yield q :

$$d_1 = \frac{\ln(S/K) + (r - q + \sigma^2/2)\tau}{\sigma\sqrt{\tau}} \quad (12)$$

$$d_2 = d_1 - \sigma\sqrt{\tau} \quad (13)$$

5.2 First-Order Greeks

For a **call** option:

$$\Delta_C = e^{-q\tau} N(d_1) \quad (14)$$

$$\Gamma = \frac{e^{-q\tau} n(d_1)}{S\sigma\sqrt{\tau}} \quad (15)$$

$$\Theta_C = -\frac{S\sigma e^{-q\tau} n(d_1)}{2\sqrt{\tau}} - rKe^{-r\tau} N(d_2) + qSe^{-q\tau} N(d_1) \quad (16)$$

$$\mathcal{V} = Se^{-q\tau} \sqrt{\tau} \cdot n(d_1) \quad (17)$$

For a **put** option:

$$\Delta_P = -e^{-q\tau} N(-d_1) \quad (18)$$

$$\Theta_P = -\frac{S\sigma e^{-q\tau} n(d_1)}{2\sqrt{\tau}} + rKe^{-r\tau} N(-d_2) - qSe^{-q\tau} N(-d_1) \quad (19)$$

Where $N(\cdot)$ is the cumulative normal distribution and $n(\cdot)$ is the standard normal PDF:

$$n(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2} \quad (20)$$

$$N(x) \approx 1 - n(x) \cdot t \cdot (a_1 + t(a_2 + t(a_3 + t(a_4 + t \cdot a_5)))) \quad \text{for } x > 0 \quad (21)$$

using the Abramowitz-Stegun approximation [4] with $t = 1/(1+0.2316419|x|)$.

5.3 Second-Order Greeks

$$\text{Vanna} = -e^{-q\tau} n(d_1) \frac{d_2}{\sigma} \quad (22)$$

$$\text{Charm}_C = -qe^{-q\tau} N(d_1) - \frac{e^{-q\tau} n(d_1) (2(r - q)\tau - d_2\sigma\sqrt{\tau})}{2\tau\sigma\sqrt{\tau}} \quad (23)$$

$$(24)$$

6 Implied Probability Distribution

`floe` provides functionality to extract the risk-neutral probability distribution implied by option prices, following the Breeden-Litzenberger approach [3].

6.1 Theoretical Foundation

Breeden and Litzenberger demonstrated that the risk-neutral probability density function $f(K)$ of the underlying asset at expiration can be recovered from the second derivative of call option prices with respect to strike:

$$f(K) = e^{r\tau} \frac{\partial^2 C}{\partial K^2} \quad (25)$$

where $C(K)$ is the call option price as a function of strike K , r is the risk-free rate, and τ is time to expiration.

6.2 Numerical Implementation

In practice, we estimate the second derivative using central finite differences on the mid-prices of observed call options:

$$\left. \frac{\partial^2 C}{\partial K^2} \right|_{K_i} \approx \frac{C_{i+1} - 2C_i + C_{i-1}}{(\Delta K)^2} \quad (26)$$

where C_i is the mid-price at strike K_i and $\Delta K = K_{i+1} - K_{i-1}$.

The resulting density values are normalized to sum to unity, yielding a proper probability distribution. From this distribution, we compute summary statistics including the mode (most likely price), median, expected value, and expected move (standard deviation).

7 Exposure-Adjusted Implied Probability Distribution

While the Breeden-Litzenberger approach recovers the market-implied probability distribution under risk-neutral pricing assumptions, this distribution implicitly assumes frictionless markets with continuous hedging and no feedback effects from dealer positioning. In practice, the mechanical realities of dealer hedging create systematic deviations from

these idealized conditions. We introduce an exposure-adjusted probability density function that modifies the baseline distribution to account for the mechanical effects of gamma, vanna, and charm exposures.

7.1 Theoretical Motivation

Standard option pricing theory posits that asset prices follow:

$$dS_t = \mu S_t dt + \sigma S_t dW_t \quad (27)$$

However, this formulation neglects the endogenous price pressure created by hedging flows. A more complete model incorporates these flows explicitly:

$$S_{t+\Delta t} = S_t + \mathcal{F}(\Gamma, \text{Vanna}, \text{Charm}) + \epsilon_t \quad (28)$$

where $\mathcal{F}(\cdot)$ represents the aggregate hedging flow function and ϵ_t captures residual noise. The exposure-adjusted PDF attempts to capture how \mathcal{F} distorts the probability distribution relative to the market-implied baseline. This approach builds upon empirical findings that option order flow and dealer positioning significantly impact both implied volatility surfaces [5] and underlying asset prices [6, 7].

7.2 Adjustment Framework

Let $p_0(K)$ denote the baseline Breeden-Litzenberger probability density at strike K . The exposure-adjusted density $p^*(K)$ is defined as:

$$p^*(K) = \frac{p_0(K) \cdot M_\Gamma(K) \cdot M_V(K)}{\int p_0(K') \cdot M_\Gamma(K') \cdot M_V(K') dK'} \quad (29)$$

where $M_\Gamma(K)$ is the gamma modifier, $M_V(K)$ is the vanna modifier, and the denominator ensures normalization. Charm effects are incorporated as a translation of the distribution mean rather than a multiplicative modifier.

7.3 Gamma Modifier: Kurtosis Adjustment

Gamma exposure creates zones of price “stickiness” (positive GEX) where dealer counter-trading suppresses volatility, and zones of “slipperiness” (negative GEX) where dealer hedging amplifies directional moves.

The gamma modifier at each price level K is computed by aggregating the influence of GEX at all strikes K_i :

$$M_\Gamma(K) = \prod_i \left(1 + \alpha_\Gamma \cdot \text{sgn}(\text{GEX}_{K_i}) \cdot \frac{|\text{GEX}_{K_i}|}{\max_j |\text{GEX}_{K_j}|} \cdot \psi(K, K_i) \right) \quad (30)$$

where $\alpha_\Gamma \in [0, 1]$ is the adjustment strength parameter and $\psi(K, K_i)$ is a spatial decay function:

$$\psi(K, K_i) = \frac{1}{1 + \lambda \left(\frac{K - K_i}{S} \right)^2} \quad (31)$$

with decay rate $\lambda > 0$. This formulation treats GEX concentrations analogously to charges in an electrostatic field, where influence decays with the square of normalized distance—consistent with a localized influence model where nearby exposures dominate.

For strikes with substantial positive GEX, the modifier increases probability density (attractor effect), while negative GEX decreases density (repellent effect). The net result is increased kurtosis around major GEX concentrations and reduced probability mass in negative gamma regions.

7.4 Vanna Modifier: Tail Adjustment

Vanna measures the sensitivity of delta to implied volatility ($\partial\Delta/\partial\sigma$). In declining markets, implied volatility typically spikes, creating a feedback loop: spot decline \rightarrow IV increase \rightarrow negative vanna forces dealer selling \rightarrow further spot decline.

We model this cascade by estimating the IV response to a hypothetical move to strike K :

$$\Delta\sigma(K) = \beta_{S\sigma} \cdot \frac{K - S}{S} \quad (32)$$

where $\beta_{S\sigma}$ is the spot-volatility beta (empirically, $\beta_{S\sigma} \approx -2$ to -4 for equity indices, reflecting the leverage effect).

The vanna-induced flow for a move to $K < S$ is:

$$\mathcal{F}_V(K) = \text{VEX}_{<S} \cdot \Delta\sigma(K) \quad (33)$$

where $\text{VEX}_{<S}$ is the cumulative vanna exposure below current spot. To capture the cascade dynamics, we iterate the feedback process n times with geometric decay:

$$\mathcal{E}_V(K) = \sum_{i=0}^{n-1} |\mathcal{F}_V(K)| \cdot \rho^i \quad (34)$$

where $\rho \in (0, 1)$ is the feedback dampening factor (typically $\rho = 0.5$). The vanna modifier is then:

$$M_V(K) = 1 + \min\left(\alpha_V^{\max} - 1, \frac{\mathcal{E}_V(K)}{S \cdot \kappa}\right) \quad (35)$$

where α_V^{\max} bounds the maximum tail fattening and κ is a normalization constant.

This adjustment systematically fattens the left tail of the distribution when negative vanna exposure is present below spot, reflecting the mechanical reality that downside cascades are more probable than the baseline distribution suggests.

7.5 Charm Adjustment: Mean Shift

Charm ($\partial\Delta/\partial t$) represents the deterministic component of delta evolution due to time decay. Net charm exposure translates directly to expected hedging flow over a given time horizon:

$$\mathcal{F}_C = \text{CEX}_{\text{total}} \cdot \tau_H \quad (36)$$

where τ_H is the time horizon multiplier (e.g., $\tau_H = 0.25$ for intraday, $\tau_H = 1$ for daily).

Rather than modifying the probability density shape, charm induces a parallel shift in the distribution mean:

$$\mu^* = \mu_0 + \eta \cdot \frac{\mathcal{F}_C}{S \cdot \xi} \quad (37)$$

where μ_0 is the baseline expected value, η is the shift scale parameter, and ξ is a flow-to-price-impact conversion factor. This shift is implemented by translating the strike axis of the probability distribution.

7.6 Summary Statistics

From the adjusted distribution $p^*(K)$, we recompute the standard summary statistics:

$$\mu^* = \int K \cdot p^*(K) dK \quad (38)$$

$$(\sigma^*)^2 = \int (K - \mu^*)^2 \cdot p^*(K) dK \quad (39)$$

$$\gamma_1^* = \frac{1}{(\sigma^*)^3} \int (K - \mu^*)^3 \cdot p^*(K) dK \quad (40)$$

$$\gamma_2^* = \frac{1}{(\sigma^*)^4} \int (K - \mu^*)^4 \cdot p^*(K) dK - 3 \quad (41)$$

where γ_1^* is skewness and γ_2^* is excess kurtosis.

The comparison between baseline and adjusted distributions yields actionable metrics:

- **Mean shift:** $\Delta\mu = \mu^* - \mu_0$ indicates charm-driven directional bias
- **Tail ratios:** $p^*(K_{0.05})/p_0(K_{0.05})$ quantifies left tail fattening from vanna
- **Kurtosis change:** $\Delta\gamma_2 = \gamma_2^* - \gamma_{2,0}$ reflects gamma-induced peakedness

7.7 Configuration by Market Regime

The adjustment parameters $(\alpha_\Gamma, \lambda, \beta_{S\sigma}, n, \alpha_V^{\max}, \tau_H, \eta)$ should be calibrated to prevailing market conditions. We identify four canonical regimes:

Regime	Gamma Effect	Vanna Effect	Charm Effect
Low Volatility	Strong pinning	Moderate	Elevated
Normal	Balanced	Balanced	Balanced
Crisis	Weak pinning	Amplified cascade	Reduced
OPEX Week	Strong pinning	Moderate	Accelerated

In low volatility environments, positive GEX strikes act as strong attractors due to reduced noise overwhelming hedging flows. During crisis periods, the vanna cascade dominates as IV-spot correlations strengthen and feedback iterations compound. Approaching options expiration, both gamma pinning effects and charm decay accelerate.

7.8 Interpretation and Application

The exposure-adjusted PDF enables several analytical applications:

1. **Edge identification:** Where $p^*(K) > p_0(K)$, mechanical flows suggest the market is underpricing the probability of reaching level K .
2. **Risk-adjusted VaR:** The adjusted 5th percentile provides a flow-informed downside risk estimate.
3. **Pin probability:** Elevated density at major GEX strikes quantifies magnetic price effects.
4. **Cascade risk:** Significant vanna-driven left tail fattening signals elevated crash risk.

It bears emphasis that this adjustment framework represents a heuristic correction rather than a rigorous arbitrage-free pricing model. The parameters require empirical calibration, and the assumption of static exposures over the adjustment horizon introduces model error. Nevertheless, the framework provides a principled approach to incorporating dealer positioning into probability assessments.

8 Hedge Flow Analysis: Impulse Curve and Charm Integral

The exposure calculations of Section 4 produce per-strike dollar exposures for gamma, vanna, and charm. The question that follows naturally is: how should these be combined into an actionable representation of dealer positioning? This section develops two complementary metrics that answer orthogonal questions. The *hedge impulse curve* combines gamma and vanna into a single price-space response function via the empirical spot-volatility coupling, answering “if price moves to level S , do dealers amplify or dampen the move?” The *charm integral* accumulates the expected delta change from time passage alone, answering “if price does nothing, where does time decay push things?” The separation is deliberate: these represent conditional and unconditional forces that a trader can reason about independently.

8.1 Motivation: Dollar Exposures as Sufficient Statistics

A common approach to constructing composite positioning indicators involves applying time-dependent weighting functions to the different Greek exposures. For instance, one might assign increasing weight to

charm and decreasing weight to vanna as expiration approaches. However, this double-counts information already present in the exposures themselves.

The dollar-denominated GEX, VEX, and CEX computed in Section 4 already incorporate all time-to-expiry effects through the Black-Scholes Greeks. Gamma exposure explodes near expiry because $\Gamma \propto 1/(S\sigma\sqrt{\tau})$; vanna exposure collapses because vega decays as $\mathcal{V} \propto S\sqrt{\tau}$; charm acceleration is inherent in $\partial\Delta/\partial t \propto 1/\sqrt{\tau}$. Additional time weighting on top of these values is redundant.

This observation motivates a framework that operates directly on dollar exposures without auxiliary weighting functions, combining them only where a physical relationship justifies the combination.

8.2 The Hedge Impulse: Combining Gamma and Vanna

Consider a dealer with aggregate gamma exposure Γ and vanna exposure V at a given strike. If spot moves by ΔS , the change in the dealer's delta is approximated by Taylor expansion:

$$\Delta\delta \approx \Gamma \cdot \Delta S + V \cdot \Delta\sigma \quad (42)$$

where $\Delta\sigma$ is the implied volatility change accompanying the spot move. Intraday, the spot-volatility relationship is well-approximated by an empirical coupling:

$$\Delta\sigma \approx -k \cdot \frac{\Delta S}{S} \quad (43)$$

where $k > 0$ is the spot-vol coupling coefficient. Substituting:

$$\Delta\delta \approx \Delta S \left(\Gamma - \frac{k}{S} \cdot V \right) \quad (44)$$

The quantity in parentheses is the *hedge impulse*:

$$\boxed{H = \Gamma - \frac{k}{S} \cdot V} \quad (45)$$

This single scalar captures the net dealer delta hedge change per unit spot move, with vanna’s contribution folded in through the spot-vol coupling rather than treated as an independent additive term. When $H > 0$, a spot move triggers dealer hedging that *dampens* the move (mean-reversion). When $H < 0$, dealer hedging *amplifies* the move (trend acceleration).

8.3 Regime Derivation from the IV Surface

The hedge impulse and its derived parameters depend on quantities extracted from the IV surface itself, ensuring portability across underlyings without requiring external data sources (such as VIX levels). Given an IV surface $\sigma(K)$ with strikes $\{K_1, \dots, K_n\}$, we compute:

ATM Implied Volatility. The base volatility level σ_{ATM} is obtained by linear interpolation at spot S :

$$\sigma_{\text{ATM}} = \sigma(S) \approx \sigma(K_i) + \frac{S - K_i}{K_{i+1} - K_i} (\sigma(K_{i+1}) - \sigma(K_i)) \quad (46)$$

Skew and Spot-Vol Correlation. The IV skew at ATM encodes the implied spot-volatility correlation $\rho_{S\sigma}$:

$$\text{Skew} = \left. \frac{\partial \sigma}{\partial K} \right|_{K=S} \cdot S \implies \rho_{S\sigma} \approx \kappa_{\text{skew}} \cdot \text{Skew} \quad (47)$$

where $\kappa_{\text{skew}} \approx 0.15$ is calibrated from stochastic volatility models (SABR, Heston). For equity indices, $\rho_{S\sigma}$ is typically in $[-0.9, -0.5]$.

Curvature and Vol-of-Vol. The smile curvature encodes the implied volatility-of-volatility ν :

$$\text{Curvature} = \left. \frac{\partial^2 \sigma}{\partial K^2} \right|_{K=S} \cdot S^2, \quad \nu \approx \kappa_{\text{curv}} \cdot \sigma_{\text{ATM}} \cdot \sqrt{|\text{Curvature}|} \quad (48)$$

with $\kappa_{\text{curv}} \approx 2.0$.

Regime Classification. The ATM IV level determines the market regime:

Regime	σ_{ATM}	Interpretation
Calm	< 0.15	Low volatility, strong gamma pinning
Normal	$[0.15, 0.20)$	Balanced Greek effects
Stressed	$[0.20, 0.35)$	Elevated vanna importance
Crisis	≥ 0.35	Vanna-dominated dynamics

The expected daily moves used in subsequent calculations are:

$$\mathbb{E}[\Delta S_{\text{daily}}] = \frac{\sigma_{\text{ATM}}}{\sqrt{252}}, \quad \mathbb{E}[\Delta \sigma_{\text{daily}}] = \frac{\nu}{\sqrt{252}} \quad (49)$$

8.4 Deriving the Spot-Vol Coupling from the IV Surface

Rather than treating k as a hardcoded parameter, we derive it from observable IV surface characteristics. From stochastic volatility models (e.g., SABR, Heston), the IV skew encodes the instantaneous spot-volatility correlation $\rho_{S\sigma}$:

$$\text{Skew} = \left. \frac{\partial \sigma}{\partial K} \right|_{K=S} \cdot S \implies \rho_{S\sigma} \approx \kappa_{\text{skew}} \cdot \text{Skew} \quad (50)$$

The spot-vol coupling coefficient is then:

$$k = -\rho_{S\sigma} \cdot \sigma_{\text{ATM}} \cdot \sqrt{252} \quad (51)$$

where the annualization factor converts from instantaneous to daily coupling. For equity indices with typical skew parameters, this yields $k \in [4, 12]$, consistent with empirical observations. Crucially, because k is derived from the IV surface itself, it automatically adapts to different underlying and market regimes without manual calibration.

8.5 Mapping from Strike Space to Price Space

The per-strike exposures GEX_K and VEX_K are defined at discrete listed strikes. To evaluate the hedge impulse at arbitrary price levels—necessary for identifying zero crossings and computing directional asymmetry—we smooth these into continuous functions of price using Gaussian kernel interpolation:

$$\text{GEX}(S') = \frac{\sum_K \text{GEX}_K \cdot w(K, S')}{\sum_K w(K, S')} \quad (52)$$

$$\text{VEX}(S') = \frac{\sum_K \text{VEX}_K \cdot w(K, S')}{\sum_K w(K, S')} \quad (53)$$

where the kernel weight function is:

$$w(K, S') = \exp\left(-\left(\frac{K - S'}{\lambda}\right)^2\right) \quad (54)$$

8.5.1 Adaptive Kernel Width

A fixed percentage kernel width (e.g., $\lambda = 0.5\%$ of spot) produces inconsistent smoothing across underlyings with different strike spacings. For SPX with 5-point strikes at $S = 6000$, a 0.5% kernel spans 6 strikes. For NQ with 25-point strikes at $S = 21500$, the same percentage spans fewer than 2 strikes.

Instead, we define the kernel width in terms of the underlying's strike spacing:

$$\lambda = n_\lambda \cdot \Delta K_{\text{modal}} \quad (55)$$

where ΔK_{modal} is the most common (modal) strike spacing in the option chain and n_λ is a configurable multiplier (default $n_\lambda = 2$). This ensures consistent smoothing behavior—approximately 2 strikes of meaningful blending—regardless of the underlying's contract specifications.

The hedge impulse at each price level S' is then:

$$H(S') = \text{GEX}(S') - \frac{k}{S'} \cdot \text{VEX}(S') \quad (56)$$

evaluated on a fine grid $S' \in [S(1-r), S(1+r)]$ with step size δ (defaults: $r = 3\%$, $\delta = 0.05\%$ of spot).

8.6 Curve Analysis

The hedge impulse curve $H(S')$ encodes the complete instantaneous dealer response landscape. We extract the following features:

8.6.1 Zero Crossings

Points where $H(S') = 0$ are *gamma-vanna flip levels*. These are found by linear interpolation between adjacent grid points of opposite sign:

$$S'_{\text{flip}} = S'_i + \frac{|H(S'_i)|}{|H(S'_i)| + |H(S'_{i+1})|} \cdot (S'_{i+1} - S'_i) \quad \text{where } H(S'_i) \cdot H(S'_{i+1}) < 0 \quad (57)$$

Each crossing is classified as *rising* (H goes negative to positive) or *falling* (H goes positive to negative). Rising crossings mark transitions from acceleration zones into dampening zones; falling crossings mark the reverse.

8.6.2 Basins and Peaks

Local maxima of H where $H > 0$ are *basins* (attractors): price levels where dealer hedging creates maximal mean-reversion. These correspond to “gamma walls” or “pin levels” in trader parlance.

Local minima where $H < 0$ are *peaks* (accelerators): levels of maximal trend amplification, corresponding to “liquidity vacuums” where dealer hedging chases price away.

8.6.3 Directional Asymmetry

The directional bias is determined by integrating H over symmetric intervals above and below spot:

$$\mathcal{I}_{\text{up}} = \int_S^{S(1+\alpha)} H(S') dS' \quad (58)$$

$$\mathcal{I}_{\text{down}} = \int_{S(1-\alpha)}^S H(S') dS' \quad (59)$$

where α is the integration range (default 0.5%). The side with more negative integral is the *path of least resistance*: dealer hedging provides less resistance (or more amplification) in that direction. Note that this is the opposite sign convention from what intuition might suggest—more negative impulse means more acceleration, which means price travels more easily in that direction.

8.6.4 Regime Classification

The curve shape at current spot determines the regime:

Regime	Condition	Interpretation
Pinned	$H(S) \gg 0$	Strong mean-reversion at current level
Expansion	$H(S) \ll 0$, symmetric	Price will break away, direction uncertain
Squeeze-up	$H(S) < 0$, $\mathcal{I}_{\text{up}} < \mathcal{I}_{\text{down}}$	Upside acceleration dominates
Squeeze-down	$H(S) < 0$, $\mathcal{I}_{\text{down}} < \mathcal{I}_{\text{up}}$	Downside acceleration dominates
Neutral	$ H(S) \approx 0$	Weak or mixed signals

8.7 Charm Integral: Unconditional Time Pressure

While the hedge impulse answers “what happens if price moves,” the charm integral answers the complementary question: “what happens from time passage alone.” These two analyses are intentionally separated because they describe orthogonal forces.

The charm integral from the current time t to expiration T is:

$$\mathcal{C}(t, T) = \int_t^T \text{CEX}(u) du \quad (60)$$

where $\text{CEX}(u)$ is the charm exposure at time u . Because charm accelerates as $\partial\Delta/\partial t \propto 1/\sqrt{\tau}$, the instantaneous CEX at time u (with $\tau_u = T-u$ remaining) scales relative to the current CEX (with $\tau_t = T-t$ remaining) as:

$$\text{CEX}(u) \approx \text{CEX}(t) \cdot \sqrt{\frac{\tau_t}{\tau_u}} \quad (61)$$

In practice, we discretize into time buckets of width Δt (default 15 minutes) and compute the cumulative sum:

$$\mathcal{C}_n = \sum_{j=0}^n \text{CEX}(t_j) \cdot \frac{\Delta t}{T_{\text{session}}} \quad (62)$$

where $T_{\text{session}} = 390$ minutes (standard U.S. equity session). The sign of \mathcal{C} indicates the direction of charm-driven pressure: positive implies net buying pressure (dealers must buy to offset delta decay), negative implies net selling pressure.

8.8 PressureField: Reachable Stability and Acceleration Zones

The hedge impulse curve $H(S')$ is rich but can be difficult to consume directly in real-time. We therefore derive a compressed representation called the *PressureField* (also referred to in implementation as a pressure cloud), which maps the curve into actionable zones and regime boundaries.

Reachability Weighting. Not all levels are equally relevant intraday. A large impulse far from spot should carry less operational weight than a moderate impulse near spot. We define:

$$R = m \cdot \mathbb{E}[\Delta S_{\text{daily}}] \cdot S \quad (63)$$

$$w_{\text{reach}}(S') = \exp\left(-\left(\frac{|S' - S|}{R}\right)^2\right) \quad (64)$$

where m is a configurable reachability multiple (default $m = 2$). This Gaussian penalty suppresses distant levels while preserving local structure.

Signed Pressure Scores. For each price level:

$$\text{Stability}(S') = \max(H(S'), 0) \cdot w_{\text{reach}}(S') \quad (65)$$

$$\text{Acceleration}(S') = \max(-H(S'), 0) \cdot w_{\text{reach}}(S') \quad (66)$$

Positive impulse contributes to mean-reversion (stability); negative impulse contributes to trend amplification (acceleration).

Zone Extraction. We identify:

- **Stability zones** from local positive extrema (basins): support below spot, resistance above spot.
- **Acceleration zones** from local negative extrema (peaks/troughs): waterfall risk below spot, squeeze risk above spot.

Zone strength is normalized by the strongest impulse extremum and multiplied by reachability weight:

$$\text{Strength}_i = \frac{|H_i|}{\max_j |H_j|} \cdot w_{\text{reach}}(S_i) \quad (67)$$

Zone bounds are obtained by scanning outward from each center until impulse magnitude decays below a fixed fraction of the local extremum (implementation default: 50%).

Regime Edges. Zero crossings from Section 8.5 become explicit *regime edges*: price levels where flow behavior flips between dampening and amplifying.

Contract-Level Hedge Estimates. To translate pressure into executable scale, the per-level impulse is converted to contracts per one-point move:

$$N_X(S') = \frac{H(S')}{M_X \cdot S \cdot 0.01} \quad (68)$$

where M_X is the product multiplier (e.g., $M_{NQ} = 20$, $M_{MNQ} = 2$, $M_{ES} = 50$, $M_{MES} = 5$). This yields a signed estimate: positive implies dealer buying flow, negative implies dealer selling flow.

Trading Interpretation. The PressureField converts abstract Greek interactions into a small set of operational cues:

- Stability zones: reversal/bounce candidates.
- Acceleration zones: continuation/breakout risk.
- Regime edges: behavioral flip levels where market character changes.

8.9 Real-Time Recalculation on Open Interest Changes

Both the impulse curve and charm integral are functions of the current exposure landscape, which depends on open interest. When live OI tracking (Section 4) detects a change—for instance, a large block of puts opening at a new strike—the exposures at that strike change, propagating through both analyses:

1. New GEX/VEX at the affected strike shifts the impulse curve locally, potentially creating or eliminating zero crossings and altering the directional asymmetry.
2. New CEX at the affected strike modifies the charm integral, changing both the total charm-to-close and the shape of the cumulative curve.

This self-updating property means the analysis does not require a separate “confirmation layer”—the model incorporates new information organically as positioning evolves through the session.

9 Model-Free Implied Volatility and Realized Volatility

The dealer exposure and hedge flow frameworks developed in preceding sections operate on the positioning landscape at a single point in time. A complementary question concerns the aggregate volatility expectations embedded in option prices and how they compare to the volatility actually being realized in the underlying. This section develops two metrics: a *model-free implied volatility* computed from the full option chain using the CBOE variance swap methodology, and a *tick-based realized volatility* computed from streaming price observations. Together, these enable real-time monitoring of the variance risk premium—the spread between what the market expects and what is occurring.

9.1 Model-Free Implied Volatility

Traditional implied volatility is extracted from individual option prices by inverting the Black-Scholes formula. This produces a strike-dependent and model-dependent measure. In contrast, the model-free approach aggregates information from the entire option chain to produce a single implied variance that does not depend on any pricing model.

The methodology follows the CBOE VIX calculation [8]. For a given expiration with time to expiry T , the implied variance is:

$$\sigma^2 = \frac{2}{T} \sum_i \frac{\Delta K_i}{K_i^2} e^{rT} Q(K_i) - \frac{1}{T} \left(\frac{F}{K_0} - 1 \right)^2 \quad (69)$$

where:

- K_0 is the at-the-money strike, defined as the strike where $|C(K) - P(K)|$ is minimized
- $F = K_0 + e^{rT}(C(K_0) - P(K_0))$ is the forward price derived from put-call parity
- $Q(K_i)$ is the midpoint price of the out-of-the-money option at strike K_i : put prices for $K_i < K_0$, call prices for $K_i > K_0$, and the average of put and call at K_0
- $\Delta K_i = (K_{i+1} - K_{i-1})/2$ is the half-distance between adjacent strikes (computed from the actual chain, not assumed constant)
- r is the risk-free rate and T is time to expiry in years

The summation walks outward from K_0 in both directions, terminating after two consecutive zero-bid strikes per the CBOE convention. The annualized implied volatility is $\sigma_{IV} = \sqrt{\sigma^2}$.

9.1.1 Strike Spacing

A common implementation shortcut assumes uniform strike spacing ($\Delta K = 1$ for SPY). Our implementation computes ΔK_i from the actual strike array, making the formula applicable to any optionable underlying regardless of its listed strike structure—from SPY with \$1 strikes to AMZN with \$5 strikes to index options with varying increments across the chain.

9.1.2 Single-Expiration Mode

When applied to a single expiration, equation (69) produces the model-free implied volatility for that specific term. This is the primary mode for ODTE analysis: computing IV from today’s expiring options gives the market’s real-time pricing of expected intraday movement.

Near expiry, T becomes small, which amplifies both the variance sum (via $2/T$) and the forward price correction (via $1/T$). The measure therefore becomes increasingly sensitive to the quality of option quotes as expiration approaches—a feature, not a bug, since this sensitivity captures the market’s rapidly updating expectations for the remaining session.

9.1.3 Two-Term Interpolation

When two expirations are provided (near-term and far-term), the implied variance can be interpolated to a target constant maturity, following the CBOE VIX interpolation formula:

$$\sigma_{\text{target}}^2 = \left[T_1 \sigma_1^2 \cdot \frac{N_{T_2} - N_{\text{target}}}{N_{T_2} - N_{T_1}} + T_2 \sigma_2^2 \cdot \frac{N_{\text{target}} - N_{T_1}}{N_{T_2} - N_{T_1}} \right] \cdot \frac{N_{365}}{N_{\text{target}}} \quad (70)$$

where N_{T_1} , N_{T_2} , N_{target} , and N_{365} are the respective durations in milliseconds, and σ_1^2 , σ_2^2 are the per-term variances from equation (69).

This mode enables reconstruction of the standard CBOE VIX (using ~ 23 -day and ~ 37 -day terms to target 30 days) or any custom constant-maturity measure. For ODTE analysis, interpolating between ODTE and 1DTE options produces a constant-maturity “overnight” volatility measure.

9.2 Tick-Based Realized Volatility

Realized volatility is the empirical counterpart to implied volatility: instead of asking “what does the market expect?” it answers “what is actually happening?”

Given a sequence of n price observations $\{P_1, P_2, \dots, P_n\}$ with timestamps $\{t_1, t_2, \dots, t_n\}$, the realized quadratic variation is:

$$QV = \sum_{i=1}^{n-1} \left(\ln \frac{P_{i+1}}{P_i} \right)^2 \quad (71)$$

This is annualized by the elapsed observation period:

$$\sigma_{RV} = \sqrt{\frac{QV}{\Delta t_{\text{elapsed}}} \cdot N_{\text{year}}} \quad (72)$$

where $\Delta t_{\text{elapsed}} = t_n - t_1$ and N_{year} is the number of milliseconds per year.

9.2.1 Design: Tick-Level, Not Windowed

A common approach to intraday realized volatility uses fixed-interval sampling (e.g., 1-minute or 5-minute returns). In a streaming context where price updates arrive on every trade or quote change, fixed windows discard information between samples. Our implementation operates on all available price observations without windowing: the consumer passes the complete set of observations accumulated since session open, and the function computes from the full series.

Early in the session with few observations, the estimate is noisy. As the session progresses and thousands of ticks accumulate, the law of large numbers ensures convergence. This natural stabilization matches the behavior traders expect: uncertainty about realized vol is high at the open and resolves through the day.

9.3 The Intraday Variance Risk Premium

The spread between implied and realized volatility:

$$\text{VRP}_{\text{intraday}} = \sigma_{IV}^{\text{ODTE}} - \sigma_{RV} \quad (73)$$

is the *intraday variance risk premium*. When $\text{VRP} > 0$, options are pricing in more volatility than is being realized (options are “expensive”). When $\text{VRP} < 0$, realized movement is exceeding expectations (options are “cheap”).

This metric updates in real time as both components respond to market activity: the 0DTE IV adjusts with every option quote change, and the RV accumulates with every underlying price tick. The convergence or divergence of these two measures through the session provides a continuous signal about whether option prices are fair relative to the actual dynamics of the underlying.

9.4 Term Structure Signals

When computing IV for multiple expirations (e.g., 0DTE and 1DTE), the term structure slope carries additional information:

Condition	Interpretation
$\sigma_{\text{IV}}^{0\text{DTE}} \gg \sigma_{\text{IV}}^{1\text{DTE}}$	Event-driven session (FOMC, CPI, earnings)
$\sigma_{\text{IV}}^{0\text{DTE}} \ll \sigma_{\text{IV}}^{1\text{DTE}}$	Quiet session; volatility expected tomorrow
$\sigma_{\text{IV}}^{0\text{DTE}} \approx \sigma_{\text{IV}}^{1\text{DTE}}$	No unusual term structure; typical conditions

This term structure signal, combined with the variance risk premium, provides a two-dimensional view of the volatility landscape: whether today is special relative to tomorrow, and whether the market’s pricing is accurate relative to reality.

9.5 Vol Response Residual Model

The raw IV-RV spread tells you *whether* implied volatility exceeds realized, but not *whether this is abnormal* given the current price path. The canonical microstructure relationship—spot down, vol up—means that much of intraday IV movement is a mechanical response to underlying returns. The real signal is whether volatility is bid or offered *relative to what the price action warrants*.

We formalize this with an expanding-window OLS regression:

$$\Delta\text{IV}(t) = \alpha + \beta_1 \cdot r_t + \beta_2 \cdot |r_t| + \beta_3 \cdot \sigma_{\text{RV}} + \beta_4 \cdot \sigma_{\text{IV}} + \varepsilon_t \quad (74)$$

where $r_t = \ln(S_t/S_{t-1})$ is the log return of the underlying, $|r_t|$ captures the convexity (vol-of-vol) response, σ_{RV} is the current realized volatility level, and σ_{IV} is the current implied volatility level. The coefficients are estimated via ordinary least squares with a small ridge penalty

($\lambda = 10^{-8}$) on non-intercept terms for numerical stability of the normal equations.

The residual of the most recent observation is normalized by the residual standard deviation to produce a z-score:

$$z_t = \frac{\varepsilon_t}{\hat{\sigma}_\varepsilon} = \frac{\Delta IV(t) - \widehat{\Delta IV}(t)}{\hat{\sigma}_\varepsilon} \quad (75)$$

The z-score admits a natural classification:

Condition	Signal	Interpretation
$z > 1.5$	Vol Bid	IV rising faster than price action justifies (stress/demand)
$z < -1.5$	Vol Offered	IV falling faster than expected (supply/crush)
$ z \leq 1.5$	Neutral	Normal vol response given the return path

The regression uses an expanding window from session open: all observations accumulated so far are included in each fit. With few observations early in the session, the model returns an `insufficient_data` signal until a configurable minimum (default 30) is reached. As the session progresses, the coefficients stabilize and the z-score becomes increasingly reliable.

This approach is particularly suited to 0DTE monitoring, where there is no meaningful historical seasonal baseline—each session’s vol dynamics are learned in real time from the session itself. The R^2 of the regression provides a goodness-of-fit check: high R^2 means the price path explains most IV movement and residuals carry strong signal; low R^2 suggests structural breaks or unusual dynamics.

10 Shares Needed to Cover

As a sort of global heuristic across all exposures at all strikes for a given expiration, one can estimate the total hedging flow required to neutralize dealer exposure:

$$\text{Shares to Cover} = \frac{-\text{Net Exposure}}{S} \quad (76)$$

$$\text{Implied Move} = \frac{\text{Shares to Cover}}{\text{Shares Outstanding}} \times 100\% \quad (77)$$

The sign indicates directional pressure:

- Negative net exposure \Rightarrow dealers must buy \Rightarrow upward price pressure
- Positive net exposure \Rightarrow dealers must sell \Rightarrow downward price pressure

This method assumes equal weighting of exposure across all strikes, and that the hedge will be perfectly (and exclusively) executed by buying or selling shares of the underlying.

11 Future Work

- Calibrate the spot-vol coupling coefficient k (Section 8.4) empirically across asset classes and volatility regimes. The current derivation from skew slope is theoretically motivated but would benefit from validation against observed intraday $\Delta\sigma/\Delta S$ relationships.
- Improve nuances of live open interest estimation by weighing the confidence of each change in OI by how far the trade price is from the mid. Trades clearly at the bid or ask have higher signal; trades near mid price are ambiguous as to aggressor side.
- Develop IV surface evolution prediction: given dealer positioning, model how the IV surface itself changes in response to spot moves and time passage. This would involve estimating vanna-driven IV pressure, vomma effects, and iterative cascade simulation where IV changes trigger further hedging flows.
- Backtest the hedge impulse curve's regime classifications and directional asymmetry signals against realized intraday price action to validate predictive utility across different market regimes and underlyings.
- Backtest the exposure-adjusted PDF against realized price distributions to calibrate adjustment parameters and validate the framework's predictive utility.

12 Summary

The complete pipeline for real-time dealer exposure calculation:

1. **Initialize:** Fetch option chain with OI_0 at market open
2. **Build IV Surface:** Calculate IV for each option, apply total variance smoothing
3. **Compute Greeks:** For each option using smoothed IV and current spot

4. **Aggregate Exposures:** Calculate GEX, VEX, CEX, and sum them for NEX across strikes per expiration
5. **Stream Updates:** On each new quote, recalculate IV \rightarrow Greeks \rightarrow Exposures
6. **Track Live OI:** Adjust open interest based on observed trades
7. **Analyze Hedge Flow:** Compute impulse curve and charm integral from updated exposures
8. **Compute PressureField:** Extract stability/acceleration zones, regime edges, and contract-scale hedge estimates
9. **Compute IV/RV:** Model-free implied volatility from option chains and tick-based realized volatility for variance risk premium monitoring
10. **Vol Response:** Fit expanding-window regression on IV changes, compute z-score to classify vol as bid, offered, or neutral relative to the price path

This methodology enables sub-second exposure updates, providing actionable insight into market-maker hedging dynamics as market conditions evolve.

References

- [1] F. Black and M. Scholes, “The pricing of options and corporate liabilities,” *Journal of Political Economy*, vol. 81, no. 3, pp. 637–654, 1973.
- [2] R. C. Merton, “Theory of rational option pricing,” *Bell Journal of Economics and Management Science*, vol. 4, no. 1, pp. 141–183, 1973.
- [3] D. T. Breeden and R. H. Litzenberger, “Prices of state-contingent claims implicit in option prices,” *Journal of Business*, vol. 51, no. 4, pp. 621–651, 1978.
- [4] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Washington, DC: National Bureau of Standards, Applied Mathematics Series 55, 1964.
- [5] N. P. B. Bollen and R. E. Whaley, “Does net buying pressure affect the shape of implied volatility functions?” *Journal of Finance*, vol. 59, no. 2, pp. 711–753, 2004.

- [6] N. Gârleanu, L. H. Pedersen, and A. M. Poteshman, “Demand-based option pricing,” *Review of Financial Studies*, vol. 22, no. 10, pp. 4259–4299, 2009.
- [7] S. X. Ni, J. Pan, and A. M. Poteshman, “Volatility information trading in the option market,” *Journal of Finance*, vol. 63, no. 3, pp. 1059–1091, 2008.
- [8] Cboe Global Markets, “VIX White Paper: Cboe Volatility Index,” 2023. Available: https://www.cboe.com/tradable_products/vix/vix_white_paper/

A Repository

For the full source, examples, and license, see:

- TypeScript (`floe`): <https://github.com/FullStackCraft/floe>
- Go (`floe-go`): <https://github.com/FullStackCraft/floe-go>