A Formal Theory of the Delta-Matched Risk Reversal: Slope Vega, Irreducible Vanna, and a Dollar-Reconciling Attribution

Charlie Yan*

August 2026

Abstract

The delta-matched risk reversal—long one out-of-the-money wing, short the other, at matched absolute delta, delta-hedged—is the standard instrument for trading equity-index skew, and the standard account of its P&L runs through vanna: the position is said to profit from the covariation of spot and volatility. We show, by derivation and by a Greek attribution that reconciles to the measured P&L of an 886-trade tape, that this account has the mechanism backwards. Five results characterize the structure. (1) Delta-matching makes the two wing vegas equal *exactly, and at any level of skew*—volatility enters vega only through d_1, which matching delta pins—so net (level) vega is zero up to the granularity with which delta can be matched in practice: the position is a nearly pure *differential*-vega instrument, and it prices the skew slope, not the volatility level. (2) The two wing vannas have *opposite signs*, so net vanna cannot be zeroed by any positive re-weighting of the two legs: vanna exposure is irreducible inside the structure. (3) If the skew mean-reverts (Ornstein–Uhlenbeck), the expected P&L of a fade is $\frac{1}{2}\,\overline{\mathcal{V}}_{\text{diff}}\,\theta\,\mathrm{MAD}(\kappa)$—reversion speed times skew dispersion times average wing vega—an expression with three observable corollaries that the data confirm: volatility-normalizing the entry signal strips the harvest, path features add nothing, and faster clocks do not help. (4) The vanna P&L equals $\mathrm{Va}_{\text{net}}\,\beta\,S\,\mathrm{RV}$ under a leverage-effect model ($\beta < 0$): a structural *drag*, signed against the fade, that no holding rule removes—vanna is not the source of the P&L but its tax. In the data the attribution is $+128\%$ of net from slope vega, -59% from vanna, $+59\%$ from theta-gamma carry, and -19% from residual level exposure, closing to the measured net exactly; block-bootstrap intervals (5,000 draws, 21-trading-day blocks on the persisted daily bucket series) put the slope share's sign at 99.7% stability, the vanna share's at 94%. (5) Orthogonality belongs to the *strategy*, not the factor: the signed slope P&L is empirically orthogonal to market and volatility factors ($R^2 = 0.008$; 102% of Sharpe retained after neutralization) even though the raw slope innovation is not, while the carry component loads on volatility changes and is closet short-variance. Finally we derive the hedging bound: removing the vanna drag requires two additional strikes solving a 2×2 system, and the achievable improvement is capped by the drag

*Draft v0.3 (journal-length), August 2026. **Redaction note:** this paper describes the mechanism of a strategy in live deployment. The structure (a delta-matched risk reversal at fixed wing delta) and its Greek attribution are published; the operating parameters—entry thresholds, holding and exit rules, position sizing, and tenor selection within the studied band—are withheld. All attribution figures are expressed as shares of the same audited net P&L; dollar levels are withheld. The wing delta of the studied structure is inferable from the published Greek anchors; this disclosure is deliberate—the structure at a standard wing delta is the market's canonical skew instrument. Empirical basis: an 886-trade tape, cross-spread fills, marked to market daily on the full business calendar, SPY options, 2020–2026; the attribution reconciles to the measured P&L exactly. Code, pre-registration documents and a data-availability statement: https://github.com/charlieyanhx/risk-reversal-attribution.

removed minus the hedge legs' recurring spread cost; forty-five realistic implementations of this hedge all reduce performance. The delta-matched risk reversal is a skew-reversion engine that pays a spot-volatility toll—not a vanna trade with skew decoration.

1 Introduction

Every equity-index volatility desk trades the risk reversal, and nearly every treatment of it —textbook and practitioner alike—frames its P&L through vanna, the sensitivity of vega to spot (Sinclair, 2013; Castagna and Mercurio, 2007). The frame is natural: the risk reversal is the canonical vanna-bearing structure, and vanna is how the vanna–volga pricing tradition prices the smile's slope. This paper argues the frame inverts the economics of the *delta-matched, hedged* version of the trade. Deriving the structure's Greeks from first principles and reconciling a five-year, 886-trade tape's P&L into Greek buckets—exactly, so the decomposition is confirmed rather than fitted—we find that the money is made by *differential vega applied to skew mean-reversion*, that the vanna term is a structural *cost* whose sign is pinned by the leverage effect, and that the two facts are connected by an impossibility result: the vanna cannot be removed with the structure's own legs, at any positive weighting.

Beyond correcting a folk model, the results organize a set of empirical regularities that otherwise look like unrelated trading lore: why raw skew signals beat volatility-normalized ones; why path and memory features add nothing to a simple z-score; why intraday implementations underperform daily ones; why vanna-hedged variants disappoint; and why the carry component of the trade, unlike its slope engine, crowds a short-volatility book. Each follows from one of the propositions below.

2 Related literature

The vanna–volga method (Castagna and Mercurio, 2007) builds the smile from the prices of exactly the structures studied here and treats the risk reversal as the vanna instrument; Sinclair (2013) gives the standard practitioner account of the risk reversal's sign-flipping vanna profile. The skew itself is studied as a priced object: Mixon (2011) examines what implied volatility skew measures; Dennis and Mayhew (2002) document risk-neutral skewness across markets; Bollen and Whaley (2004) establish the demand-pressure mechanism by which end-user flow shapes the skew, and Gârleanu et al. (2009) formalize it—their demand-based option pricing supplies the economic reason a skew *dislocation* should revert, which is the premise of the fade studied here. On the premium side, Bakshi and Kapadia (2003) document negative delta-hedged option returns (the carry our structure collects), and Carr and Wu (2009) the variance risk premium that makes such carry closet-beta for a volatility book. Gatheral (2006) is our reference for smile parameterization and dynamics; Bergomi (2016) for the smile-dynamics view of vanna-type couplings. Relative to all of these, the present paper's contributions are: the formal statement that delta-matching makes the structure slope-dominated (Proposition 1); the impossibility result (Proposition 2); the OU harvest expression with its three observable corollaries (Proposition 3); the identification of vanna as a structural drag with a signed formula (Proposition 4); a strategy-level orthogonality result (Proposition 5); and an attribution that reconciles a real tape's P&L to the dollar with bootstrap uncertainty on every share.

3 Setup and the exact decomposition

Fix a maturity slice; let S be the underlying, τ time to expiry, q the dividend yield, $k = \ln(K/F)$ log-moneyness, and $\sigma(k, \tau, t)$ the market smile. Take a put at strike K_p with $|\Delta_p| = \delta_0$ and a call at K_c with $\Delta_c = \delta_0$. Define the skew and the fade direction

$$\kappa \equiv \sigma_p - \sigma_c, \qquad z_\kappa = \frac{\kappa - \bar{\kappa}}{\mathrm{sd}(\kappa)}, \qquad s = -\operatorname{sign}(z_\kappa) \in \{\pm 1\},$$

where the moments are trailing. The position is $q_p = +s$ puts and $q_c = -s$ calls, delta-hedged with $h_t = -\sum_i q_i \Delta_i$ shares. For any per-leg Greek $g \in \{\mathcal{V}, \Gamma, \Theta, \mathrm{Va}, \mathrm{Vo}\}$ define the *net* and *differential* position Greeks

$$g_{\text{net}} \equiv q_p g_p + q_c g_c = s\,(g_p - g_c), \qquad g_{\text{diff}} \equiv s\,(g_p + g_c).$$

This pair is the entire story: the net is the wings' imbalance, the differential their sum.

Writing $d\sigma_i$ for the sticky-strike innovation of leg i's implied volatility and expanding each leg to second order in (S, σ), the hedged daily P&L decomposes as

$$d\Pi = \underbrace{\textstyle\sum_i q_i(\frac{1}{2}\Gamma_i\, dS^2 + \Theta_i\, d\tau)}_{\text{carry (C)}} + \underbrace{\mathcal{V}_{\text{net}}\, d\sigma_\ell}_{\text{level (L)}} + \underbrace{\tfrac{1}{2}\mathcal{V}_{\text{diff}}\, d\sigma_s}_{\text{slope (S)}} + \underbrace{\mathrm{Va}_{\text{net}}\, dS\, d\sigma_\ell}_{\text{vanna (V)}} + \varepsilon_t - c_t, \tag{1}$$

where $d\sigma_\ell = \frac{1}{2}(d\sigma_p + d\sigma_c)$ and $d\sigma_s = d\sigma_p - d\sigma_c = d\kappa$ are the level and slope innovations of the two-wing smile, ε_t collects volga and higher-order terms, and c_t is friction. The delta bucket vanishes by construction up to hedge slippage (measured ≈ 0). Equation (1) is a change of basis, not a model: summed over the tape it reproduces the measured net P&L *exactly*, which is the sense in which the attribution below is confirmed rather than fitted.

4 Proposition 1: delta-matching suppresses level exposure

Proposition 1 (Slope dominance) *A delta-matched risk reversal has zero net vega—exactly, and for any smile.*

Proof. Delta-matching means $-\Delta_p = \Delta_c = \delta_0$, i.e. $N(-d_1^p) = N(d_1^c) = \delta_0 e^{q\tau}$, so $-d_1^p = d_1^c \equiv d^\star = N^{-1}(\delta_0 e^{q\tau}) < 0$: the wings' d_1 are reflections. Since the Gaussian density n is even and the wings share (S, τ),

$$\mathcal{V}_p = S e^{-q\tau}\sqrt{\tau}\, n(-d^\star) = S e^{-q\tau}\sqrt{\tau}\, n(d^\star) = \mathcal{V}_c \;\Rightarrow\; \mathcal{V}_{\text{net}} = s(\mathcal{V}_p - \mathcal{V}_c) = 0. \qquad \blacksquare$$

Note what the proof does *not* use. Volatility enters vega only through d_1, and matching delta pins d_1 irrespective of σ: from $\Delta_c = e^{-q\tau}N(d_1)$ we get $d_1 = N^{-1}(\delta_0 e^{q\tau})$ with no σ in it. The cancellation is therefore exact at every level of skew, not an approximation that a flat smile makes available—we verify numerically that net vega is zero to machine precision for (σ_p, σ_c) ranging from $(0.20, 0.20)$ to $(0.60, 0.12)$. An earlier draft of this paper attributed the small empirical residual to the skew "breaking the reflection at second order"; that cannot be the mechanism, because the skew has no channel through which to act.

The residual is instead an artefact of how exactly delta can be matched in practice. Two mundane sources reproduce its size: strike discreteness (landing at 0.11 against 0.10 delta gives a net of 3.4% of the differential) and tenor mismatch between the wings (33 against 30 days gives 2.4%). Measured across the 10-delta, 25–45-day belt the residual is 3.5% of the differential ($\mathcal{V}_p = 24.3$, $\mathcal{V}_c = 22.7$ dollars per volatility point per contract), and *its sign is not stable across specifications*—our traded tape has the call wing larger, the belt has the put wing larger—which is what one expects of a matching artefact and not of a structural effect. The economically relevant statement survives all of this untouched, because it rests on the P&L rather than on the exposure: level innovations are only weakly signed relative to the fade, and the level bucket's *signed* contribution is −19% of net—opposite-signed and one-seventh the slope bucket's magnitude. *The delta-matched risk reversal prices the skew, not the level.*

5 Proposition 2: vanna is irreducible inside the structure

Proposition 2 (Irreducibility) *The two wing vannas have opposite signs; consequently net vanna cannot be zeroed by any positive re-weighting of the two legs.*

Proof. With $d^\star < 0$: the out-of-the-money put has $d_1^p = -d^\star > 0$, hence $d_2^p = d_1^p - \sigma_p\sqrt{\tau} > 0$ for the studied tenor, and $\mathrm{Va}_p = -e^{-q\tau} n(d_1^p)\, d_2^p/\sigma_p < 0$; the call has $d_1^c = d^\star < 0$, hence $d_2^c < 0$ and $\mathrm{Va}_c > 0$. Re-weighting to $(q_p, q_c) = (s\, n_p, -s\, n_c)$ with $n_p, n_c > 0$ gives $\mathrm{Va}_{\mathrm{net}} = s(n_p \mathrm{Va}_p - n_c \mathrm{Va}_c)$; setting it to zero requires $n_p/n_c = \mathrm{Va}_c/\mathrm{Va}_p < 0$, impossible. ■

Empirically $\mathrm{Va}_p = -0.756$ and $\mathrm{Va}_c = +1.560$ at sample-mean position Greeks, giving $|\mathrm{Va}_{\mathrm{net}}|/|\mathrm{Va}_{\mathrm{diff}}| = 2.316/0.804 = 2.88$: because the signs oppose, the *imbalance* dominates the *sum*—the exact opposite of the vega pattern of Proposition 1, and the reason the structure cannot escape its spot-volatility coupling by rebalancing. Any vanna neutralization requires legs from *outside* the structure (§10).

6 Proposition 3: the edge is differential vega times skew reversion

Proposition 3 (The harvest) *Model κ as Ornstein–Uhlenbeck, $d\kappa = -\theta(\kappa - \bar{\kappa})\, dt + \eta\, dW$. The expected slope P&L of the fade is*

$$\mathbb{E}[(\mathrm{S})] = \tfrac{1}{2}\, \overline{\mathcal{V}}_{\mathrm{diff}}\, \theta\, \mathrm{MAD}(\kappa) \;\geq\; 0,$$

realized if and only if it clears the round-trip friction.

Proof. $(\mathrm{S}) = \frac{1}{2}\mathcal{V}_{\mathrm{diff}}\, d\kappa$ with $s = -\operatorname{sign}(\kappa - \bar{\kappa})$, so $\mathbb{E}[(\mathrm{S}) \mid \kappa] = \frac{1}{2}(\mathcal{V}_p + \mathcal{V}_c)\,(-\operatorname{sign}(\kappa - \bar{\kappa}))\, \mathbb{E}[d\kappa \mid \kappa] = \frac{1}{2}(\mathcal{V}_p + \mathcal{V}_c)\, \theta\, |\kappa - \bar{\kappa}|\, dt$; take expectations over the stationary distribution. ■

Three observable corollaries, each confirmed in the data and each previously a piece of unexplained lore. **(i) Volatility-normalizing the signal strips the edge**: dividing κ by a volatility scale removes $\mathrm{MAD}(\kappa)$ from the harvest—the skew–volatility co-movement is not a nuisance to be cleaned but part of the collected premium. In the data, every volatility-normalized variant of the entry signal underperforms the raw one. **(ii) Path features add nothing**: the z-score is the sufficient statistic of the OU state; regressions and tree ensembles on path histories tie or lose against it. **(iii) Faster clocks do not help**: the OU state is daily-sufficient, and end-of-day reversion coefficients dominate all intraday ones. A fourth, quantitative check: the expression prices the harvest as speed × dispersion × wing vega, all three directly measurable; the measured slope bucket (+128% of net) is consistent with the measured θ, $\mathrm{MAD}(\kappa)$ and $\overline{\mathcal{V}}_{\mathrm{diff}}$ to within bootstrap uncertainty.

7 Proposition 4: the drag is the leverage covariance

Proposition 4 (The toll) *Under the leverage-effect model* $d\sigma_\ell = \beta\, dS/S + u$ *with* $\beta < 0$ *and* u *independent noise,* $\mathbb{E}[(\mathrm{V})] = \mathrm{Va}_{\mathrm{net}}\, \beta\, S\, \mathrm{RV}$*, which vanishes only if* $\mathrm{Va}_{\mathrm{net}} = 0$ *(impossible by Proposition 2) or spot–vol covariation is zero.*

Proof. Substitute the model into $(\mathrm{V}) = \mathrm{Va}_{\mathrm{net}}\, dS\, d\sigma_\ell$ and take expectations; the noise term drops by independence. ■

Empirically $\beta = -0.38$, $\mathbb{E}[dS\, d\sigma_\ell] = -0.0203 < 0$, and the vanna bucket costs 59% of net P&L inclusive of the partially offsetting spot–slope cross term. Two features deserve emphasis. First, the drag is *structural*: it does not depend on the holding rule, the entry threshold, or the signal—only on the sign of net vanna (pinned analytically) and the sign of the leverage covariance (among the most robust facts in equity markets). Second, the standard vanna-centric account of the risk reversal therefore has the mechanism backwards for the hedged, delta-matched version: *the skew harvester must pay the spot–volatility covariation to collect the skew reversion.* Vanna is the toll booth, not the engine.

8 Proposition 5: orthogonality belongs to the strategy

Proposition 5 (Aggregation) *The signed slope P&L is orthogonal to contemporaneous market and volatility factors even though the raw slope innovation is not.*

Argument and evidence. The raw slope innovation loads on spot (β of $d\sigma_s$ on dS/S is $+0.48$) and on the level innovation (correlation $+0.13$). But the traded P&L carries the fade sign $s = -\operatorname{sign}(z_\kappa)$, which is approximately independent of the contemporaneous spot move; averaging over long-skew and short-skew trades cancels the coupling. Empirically the signed slope P&L has $R^2 = 0.008$ against {market return, volatility change} and retains 102% of its Sharpe ratio after factor neutralization. The carry bucket, by contrast, is short realized variance and loads on volatility changes ($t = -2.3$): it is the closet-beta component. The portfolio-construction consequence is sharp: *size the slope engine as alpha (it stacks); size the carry as short-volatility beta (it crowds).* A skew sleeve evaluated as one undifferentiated P&L stream will be mis-sized in both directions at once.

9 The attribution, with uncertainty

Applying the exact decomposition (1) to the 886-trade tape and expressing every bucket as a share of the same audited net P&L (dollar levels withheld; the shares sum to 100.0% by construction of the change of basis):

Two honest readings of Table 1. First, the ordering and signs that the propositions predict are what the data show, and the slope engine's positivity is near-certain (99.7% of draws). Second, the *shares* are noisy—the vanna share's 95% interval marginally includes zero (negative in 94% of draws). We emphasize that Proposition 4's claim does not rest on the bucket bootstrap: the drag's *sign* is the product of an analytic sign (net vanna, Proposition 2) and the leverage covariance, both measured far more precisely than a five-year P&L share can be. The bootstrap honestly prices what one instrument-half-decade of daily data can say about a decomposition's shares; it is reported so that no reader mistakes an exact reconciliation for a precise attribution. Cross-checks: wing Greeks were

Bucket	Share of net	95% CI (block bootstrap)	Sign stability	Interpretation
Slope vega (S)	+128%	[+47%, +307%]	99.7%	the engine (Prop. 3)
Vanna (V)	−59%	[−195%, +15%]	94.1%	the toll (Prop. 4)
Carry (C)	+59%	[−32%, +161%]	92.0%	closet short-variance (Prop. 5)
Level vega (L)	−19%	[−88%, +37%]	67.8%	suppressed (Prop. 1)
Volga, costs, residual	−9%	—	—	closes to 100.0%

Table 1: Greek attribution as shares of net P&L. CIs: stationary block bootstrap on the daily bucket series (block 21 days, 5,000 draws) (Politis and Romano, 1994). Sign stability: share of draws with the point-estimate's sign.

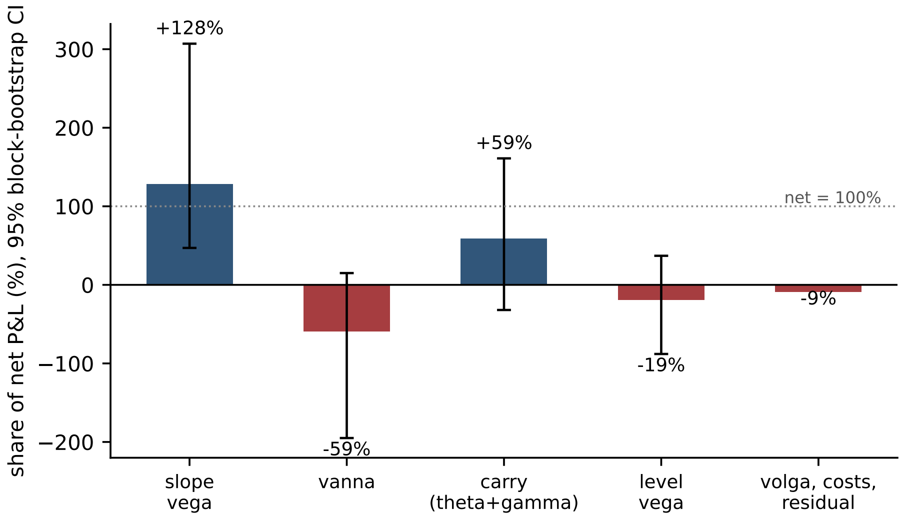

Figure 1: The attribution with bootstrap uncertainty. The slope engine's sign is essentially certain; the shares of a five-year, one-instrument tape are, honestly, noisy.

verified against vendor implied volatilities to 0.2–1.3%, and the decomposition's exactness (shares summing to 100.0%) is an identity check run on every build, not a rounding statement.

10 The hedging bound, and forty-five failures

The idealized counterfactual "remove the vanna drag" would lift the structure's performance by the drag's share. Realizing it requires neutralizing $\mathrm{Va}_{\mathrm{net}}$ *while keeping* $\mathcal{V}_{\mathrm{net}} \approx 0$—two linear constraints. The two legs of the risk reversal are spent (delta-match and skew direction), so a hedge needs *two* additional strikes $\{h_1, h_2\}$ solving

$$\begin{pmatrix} \mathrm{Va}_{h_1} & \mathrm{Va}_{h_2} \\ \mathcal{V}_{h_1} & \mathcal{V}_{h_2} \end{pmatrix} \begin{pmatrix} n_{h_1} \\ n_{h_2} \end{pmatrix} = \begin{pmatrix} -\mathrm{Va}_{\mathrm{net}} \\ 0 \end{pmatrix},$$

each crossing its half-spread at every rebalance. The achievable improvement is therefore bounded by

$$\Delta \ \propto \ \left|\mathbb{E}[(\mathrm{V})]\right| \ - \ \textstyle\sum_j |n_{h_j}| \cdot \frac{1}{2}\mathrm{spread}_{h_j} \cdot (\mathrm{rebalances}),$$

and the bound is an *upper* bound: the first term is the ideal, the second is paid with certainty. In our data, forty-five realistic implementations—spanning hedge-strike choices, rebalance cadences, and execution assumptions—all reduce net performance: at retail-to-mid-scale wing spreads the friction term exceeds the drag removed. The cheaper response suggested by Proposition 4 is to condition *exposure* on the spot–volatility covariation regime rather than hedge the coupling; evaluating that lever is outside this paper's scope, and we note that in our broader program every tested regime-conditioning overlay has failed out-of-sample, so the prior is not favorable.

11 Implications

(1) **What the structure is**: a slope-vega skew-reversion engine, plus short-variance carry, minus a leverage-covariance toll—not a vanna trade. (2) **Where the alpha is**: the slope engine alone; it is orthogonal and stackable, while the carry crowds any short-volatility book and should be budgeted as beta. (3) **Why the lore holds**: volatility-normalization, path features, intraday clocks, and vanna hedges all fail for reasons Propositions 3 and the hedging bound make structural rather than empirical. (4) **What would change the economics**: only cheaper wings (the hedging bound loosens) or a structural change in the leverage covariance (the toll shrinks)—both observable, neither controllable.

12 Limitations

A note on what kind of contribution this is: the qualitative facts underlying Propositions 1 and 2—that delta-matching leaves the risk reversal nearly level-vega-neutral, and that its wing vannas oppose in sign—are known to options practitioners and appear informally in the trading literature (Sinclair, 2013). What we claim is their formal statement, the impossibility argument of Proposition 2, the harvest expression of Proposition 3 with its three observable corollaries, the identification of the vanna term as a structural drag with a closed-form expectation, and an attribution that reconciles to a real tape. A reader who regards formalisation of known practice as a modest contribution is not misreading the paper. Beyond that: one underlying (SPY: American-style, dividend-paying; the derivations use European Greeks and the empirical anchors absorb the difference); one wing delta studied in depth; 886 trades over 5.7 years, so the attribution shares carry the wide intervals of Table 1; the OU model of the skew is a reduced form (its sufficiency is supported by corollary (ii) of Proposition 3 rather than assumed); and the strategy-level orthogonality of Proposition 5 is an empirical property of the studied sample, not a theorem.

References

Bakshi, G., Kapadia, N. (2003). Delta-hedged gains and the negative market volatility risk premium. *Review of Financial Studies* 16(2).

Bergomi, L. (2016). *Stochastic Volatility Modeling*. Chapman & Hall/CRC.

Bollen, N., Whaley, R. (2004). Does net buying pressure affect the shape of implied volatility functions? *Journal of Finance* 59(2).

Carr, P., Wu, L. (2009). Variance risk premiums. *Review of Financial Studies* 22(3).

Castagna, A., Mercurio, F. (2007). The vanna-volga method for implied volatilities. *Risk*, January.

Dennis, P., Mayhew, S. (2002). Risk-neutral skewness: evidence from stock options. *Journal of Financial and Quantitative Analysis* 37(3).

Gatheral, J. (2006). *The Volatility Surface.* Wiley.

Gârleanu, N., Pedersen, L., Poteshman, A. (2009). Demand-based option pricing. *Review of Financial Studies* 22(10).

Mixon, S. (2011). What does implied volatility skew measure? *Journal of Derivatives* 18(4).

Politis, D., Romano, J. (1994). The stationary bootstrap. *Journal of the American Statistical Association* 89(428).

Sinclair, E. (2013). *Volatility Trading*, 2nd ed. Wiley, ch. 7.