Components vs. the Literature

Tech Charts, Trade Plan, Conviction Engine & Expiry Optimizer (v3.95.0) vs Published Quantitative Research

Prepared: August 16, 2026

Scope: Cross-reference of COMPONENT_SPEC.md (v3.95.0) against academic and rigorous practitioner research on chart-pattern analysis, option expiration selection and price forecasting by maturity, probability-of-profit accuracy, and volatility-regime-conditioned strategy selection. Approximately 90 primary sources reviewed in four targeted sweeps; consolidated bibliography in Section 8. Companion to the earlier engine-wide review (*Engine vs. the Literature*, 2026-08-16); overlapping findings from that report are referenced, not repeated.

1. Executive Summary

The four components hold up well where it matters most. The **architecture decisions** — patterns as veto/bias rather than entries, a 21–60 DTE window with a hard floor above the 0DTE zone, POP that only downgrades verdicts, technical filters that can only *tighten* an implied-vol-based margin, cone charts that disclose their own downside bias — are each the configuration the evidence best supports, in several cases matching published best practice more closely than typical retail tooling.

Against that, this review found **4 divergences**, **6 material gaps**, and **10 refinements**. The highest-priority items:

1. **Measured-move targets are systematically too ambitious.** Chang & Osler's mirror ratios fall short of 1.0 in nearly every market tested [2], and Bulkowski's own large-sample hit rates for full-height targets run 46–71% by pattern (H&S top: 51%; flags: 46%) [15]. Tech Charts projects full pattern height. Best practice: scale targets to the pattern-specific hit rate (~50–70% of height) or present them as probability zones — this also propagates into the tape gate's measured-move veto distances.
2. **The Hurst regime input is statistically uninformative at the engine's sample sizes.** Finite-sample R/S estimation of a true random walk averages $\hat{H} \approx 0.56$ at N=500 — biased *toward* "trending" — and meaningful tests need ~5,000+ observations (Couillard & Davison [22]). One of the range-regime's three votes is mostly noise with a directional bias.
3. **The 1.25× expected-move earnings buffer is thinner than the jump distribution implies.** Under Dubinsky-Johannes near-normality of scaled announcement moves, 1.25σ contains only ≈79% of events — roughly one in five earnings threatens the widened strike. 1.5–2.0× (≈87–95%) is the literature-consistent buffer for the human path (the agent's default earnings blackout already avoids the issue) [47][48].
4. **The 14-day straddle-entry gate is ~3× wider than the documented edge.** Gao, Xing & Zhang's pre-earnings straddle profit is a 3–5-day phenomenon; entries 14 days out spend most of their life paying theta in the window the same paper shows to be flat-to-negative [51].
5. **The conviction skew bonus rewards selling into skew the literature reads as directional information.** Kozhan, Neuberger & Schneider show no independently harvestable skew premium (it is ~0.9-correlated with the variance premium the IV/HV gate already captures) [56], while

Xing-Zhang-Zhao show a steep put smirk *predicts underperformance* [58] — i.e., the surface shape is a directional signal, not a separate premium, and the bonus's sign is questionable in exactly the cases it rewards most.

6. **POP misses the Breeden-Litzenberger skew-slope term.** Using the strike's own smile IV captures the smile's *level* but omits the vega × ∂σ/∂K correction — worth roughly +2–5 probability points of P(ITM) at 10–30-delta index put strikes (the short put is riskier than displayed). One cheap finite-difference closes most of the gap [40][41].

The best news in the research: several component features that lacked justification in the earlier engine-wide review turn out to be *well* supported at component level — the hard 21-DTE floor (0DTE retail losses are now directly documented [30]), the GEX *regime* caps (gamma-fragility and hedging-demand momentum are peer-reviewed [62][63]), and the veto-only tape-gate architecture (the strongest defensible use of pattern information per the entire H&S literature [1]–[4]).

Component scorecard

Component	Verdict	Headline
Tech Charts — veto/bias architecture	**Aligned**	The evidence-supported configuration; patterns are 40–60-day relative bias, not entries
Tech Charts — measured-move targets	**Deviation**	Full-height targets hit only 46–71% of the time; scale to hit rates
Tech Charts — lifecycle/whipsaw	**Aligned (partial)**	Path-honesty is sound; ~50% throwback base rate argues for a busted-pattern *bias flip* state
Tech Charts — Hurst gate	**Divergent**	Uninformative and trend-biased at 250-bar samples
Tech Charts — ADX/BB-width regime	**Partial**	Unvalidated heuristics; the reserved HMM seam is the supported upgrade
Expiry Optimizer — 21–60 window, 21 floor	**Aligned**	0DTE/weekly evidence directly supports refusing sub-21-DTE entries
Expiry Optimizer — Gaussian 37-DTE peak	**Aligned (framing)**	Defensible *risk-adjusted* optimum; premium per unit risk is actually monotonic toward shorter DTE — don't present 37 as premium-optimal
Expiry Optimizer — gamma band / theta-vega	**Aligned**	Matches analytic gamma escalation and large backtests
Expiry Optimizer — term-structure scoring	**Aligned**	Vasquez/Johnson support both the seller and buyer variants
Expiry Optimizer — earnings buffer	**Deviation**	1.25× EM ≈ 79% containment; use 1.5–2× or exclude
Conviction — IV/HV conditioning	**Aligned (partial)**	Right sign (Goyal-Saretto); post-look-ahead-bias magnitudes are far smaller; thresholds → percentiles
Conviction — skew bonus	**Divergent (sign)**	Skew is directional information, not separable premium
Conviction — GEX regime caps vs wall nudges	**Split**	Regime use peer-reviewed; strike-local walls have a pre-registered null
Conviction — VIX bias layer	**Deviation**	Low VIX ≠ cheap options; high VIX ≠ good selling; slope/VRP dominates level (annotation-only design mitigates)
Trade Plan — cone + disclosure	**Aligned**	Symmetry disclosure matches the leverage-effect literature; expect ~72–80% realized ±1σ coverage, not 68%
Trade Plan / POP — N(d2) at BE	**Aligned (partial)**	Right object; add the skew-slope term; expect put/call calibration asymmetry

Component	Verdict	Headline
POP — managed-exit calibration	**Gap**	Realized win rates will beat POP by construction; track counterfactual expiry outcomes or a first-passage-to-target probability
Calibration methodology	**Partial**	Quintile report is the right artifact; power math says ±5 pts at POP 0.7 needs ~600 obs/bucket

---

2. Methodology

Reference point: COMPONENT_SPEC.md (v3.95.0, full-source component spec including frontend wiring and test-pinned invariants). Four parallel research sweeps: (i) chart patterns and regime detection, (ii) expiration selection and option-return term structure, (iii) POP accuracy and option-implied probability calibration, (iv) volatility-regime-conditioned strategy selection. Peer-reviewed journals first; named practitioner research (Bulkowski, CBOE benchmark studies, tastylive/spintwig/DTR large-sample backtests) is used where academia is silent and is labeled as practitioner-grade throughout. Magnitudes computed by the reviewers from standard formulas (rather than quoted) are flagged. Verdicts: **Aligned** / **Partial** / **Deviation** / **Divergent** / **Gap** as in the prior review.

---

3. Tech Charts

3.1 The architecture is the finding — ALIGNED

The entire H&S literature arc — Osler & Chang finding profits in only 2 of 6 currencies and dominated by simpler rules [1][2]; Osler finding H&S *unprofitable* in equities and its traders classifiable as noise traders [3]; Savin, Weller & Zvingelis rehabilitating the pattern only as a *risk-adjusted, market-relative* signal worth ≈5–7%/yr of alpha over 60 days with negative beta, and only after neckline-break confirmation [4] — converges on one defensible use: a **confirmed pattern is a 40–60-day directional bias with modest effect size**. That is precisely the component's design: patterns veto short-strike placement and hand a bias to Trade Plan; they never generate entries. The tape gate's "tighten-only" behavior (a technical veto can shrink but never extend an implied-vol-based safety margin) matches the review conclusion that no study shows technical levels beating option-implied probabilities [27].

Two design details also line up: confirmation gating (every positive H&S result conditions on the neckline break — the component's `confirmed_break`/lifecycle contract enforces exactly this), and the 10th–90th-percentile congestion box (roughly how the S/R literature defines trading ranges).

3.2 Measured-move targets — DEVIATION

Chang & Osler's measured-move "mirror ratios" fall short of 1.0 for every currency except JPY [2]. Bulkowski's large-sample hit rates for the full height-projected target [15]: H&S top **51%**, H&S bottom 71%, double top 64%, double bottom 69%, ascending triangle 70%, descending triangle 50%, **flags 46%** — his own prescription is to multiply the height by the hit percentage. The component projects full height for every pattern, and those targets propagate into (a) the pattern-conflict veto in the tape gate (an opposing pattern whose *measured move* breaches the short strike scores zero) and (b) the Tech Charts → Trade Plan handoff banner.

Recommendation: scale each pattern's target by its published hit rate (or render targets as a 50–70%-of-height zone). Note the interaction: a scaled-down opposing-pattern move *narrows* the tape-gate veto region — since the veto should be conservative, keep full height for the *veto* and scale for the *displayed target*, or use the zone's near edge for display and far edge for veto.

3.3 Lifecycle: throwbacks, whipsaw, and busted patterns — ALIGNED (PARTIAL)

The path-aware lifecycle (completed-is-permanent, failed, whipsaw demotion at ≥ 4 neckline recrossings) has no academic analog and its incident-driven design is sound. The literature adds two base rates worth encoding: throwbacks/pullbacks to the breakout level occur in roughly **half** of all breakouts [16][17] — one re-test is *normal* pattern behavior, which the ≥4-recross threshold already tolerates generously — and **busted patterns** (a close back through the pattern beyond the neckline) tend to produce stronger moves in the *opposite* direction than the original signal [16]. The component's `failed` status voids the bias; Bulkowski's data suggest it should *flip* it. A cheap upgrade: emit `failed` with an inverted-bias annotation the tape gate and handoff banner can consume.

One negative result worth keeping in mind: Bulkowski's volume study found heavy breakout volume added ~2 points of performance while roughly **doubling** failure rates and tripling throwback frequency [17] — the component's pattern detectors correctly use no volume confirmation (the Launcher scanner's separate 1.5× breakout-volume requirement inherits this caveat).

3.4 Range regime: Hurst — DIVERGENT; ADX/BB-width — PARTIAL

Couillard & Davison [22]: finite-sample R/S estimation of a *true Brownian motion* yields mean $\hat{H} \approx 0.56$ at N=500 and ≈ 0.54 at N=5,000; the popular error formula badly understates uncertainty; hypothesis tests need ~5,000–10,000 observations; with proper tests, H=0.5 cannot be rejected for the S&P 500 at any frequency. The component computes Hurst on ≤ ~250 daily bars — the estimate is noise with a bias *toward* reading "trend." ADX(14)<20 has no peer-reviewed validation (Wilder practitioner lineage), and the Bollinger-squeeze literature shows the directional edge collapsed post-publication (0.454%/day → 0.002%/day across 14 indices [14]), though low BB-width does legitimately forecast low near-term volatility.

Mitigating context: the regime flag feeds a *tradability heuristic* (tape gate placement scoring), not an alpha signal, and the spec already reserves an HMM seam. The supported upgrade path (Shu-Yu-Mulvey jump models beating HMMs and indicator thresholds on Sharpe/drawdown with explicit regime-persistence penalties [24]; ECB regime-duration estimates of ~50–80 weeks [65]) is: keep ADX/BB-width as *features*, replace the three-way AND with a persistent statistical model or at minimum add dwell-time hysteresis, and either drop Hurst or replace it with a Lo-MacKinlay variance-ratio test, which has known finite-sample distributions at these sample sizes.

3.5 Two enrichments the literature offers

- **Round numbers matter.** Osler: published S/R levels interrupt trends 60.8% vs 56.2% for arbitrary levels; 96% of levels end in 0 or 5; take-profit orders cluster *at* round numbers, stops *just beyond* them — producing bounces at the level and cascades through it [5][6]. Implication: snap box edges and displayed levels toward round numbers, and prefer pattern stops placed *beyond* the nearest round number rather than mechanically at extreme ±1%.
- **Detector decay is the norm.** The LMW kernel-detection lineage lost significance on post-1995 data in replication [11], and the strongest modern evidence (CNN chart-image models, Jiang-Kelly-Xiu [12]) finds the predictive information lives in *normalized range position*, not classical geometry. Implication: re-validate detector value periodically on the selection-ledger's own history, and treat the 2% pivot-prominence setting as a bandwidth hyperparameter worth sensitivity-testing (LMW's results were bandwidth-sensitive).

---

4. Expiry Optimizer

4.1 The window and the floor — ALIGNED

The literature is now unambiguous that the variance risk premium concentrates at the front of the curve: essentially all of it is earned in the shortest maturities (1-month claims ≈ |Sharpe| 1.3; nothing significant beyond ~3 months — Dew-Becker et al. [29]; Andries et al. confirm monotonic decline with maturity [31]). At the same time, the ultra-short end is now documented as hostile to retail: >75% of retail S&P option trades are 0DTE and lose systematically [30]; the 0DTE seller premium is "small and difficult to monetize after realistic frictions" [32]; weeklies beat monthlies on raw return only, losing on Sharpe and drawdown with 5–7× the transaction-cost load [34]. A 21–60 window with a hard floor sits exactly on the steep, harvestable-but-survivable part of the curve.

4.2 The Gaussian 37-DTE peak — ALIGNED, WITH A FRAMING CORRECTION

There is **no documented premium peak at 37 DTE** — per-unit-of-risk compensation rises monotonically as DTE falls [29][31], and Israelov-Tummala's surface shows the most compensated cells per unit of *stress loss* are front-month [28]. What justifies 37 is the *retail risk-adjusted* case: gamma ≈ 2× at 21 DTE and 5× at 7 DTE vs 45 [35]; unhedged-seller backtests favor 37–80 DTE on Sharpe and drawdown [33][34]; OTM theta is front-loaded, so most of a 30–45-DTE short's P&L accrues before the final two weeks [36]. The spec/UI should say this: the DTE score encodes a gamma-and-frictions optimum, not a premium optimum. A defensible enhancement: let advanced users shift the Gaussian center shorter with explicit tail warnings — the trade-off is a documented risk-preference axis, not a correctness question (Israelov-Tummala's "the optimal cell depends on the risk metric" [28]).

The gamma dimension itself — $m = \Gamma \cdot S \cdot \sigma / \sqrt{365}$ with full credit in [0.050, 0.085] — is a legitimate normalization matching the analytic $1/\sqrt{T}$ escalation, and the 21-DTE window floor doubles as the practitioner-standard management point (gamma 3–5× higher in the final 21 days; worst per-trade losses cluster there [35]).

4.3 Term structure and IV level — ALIGNED

Vasquez: the IV term-structure slope predicts straddle returns (high slope → buy vol; decile spreads economically significant [37]); Johnson: the slope prices variance risk across instruments [59]. The optimizer's seller-backwardation / buyer-trough scoring is the correct pair of signs, and the buyer/seller symmetric IV-level scores (pinned to sum to 20) implement the Goyal-Saretto sign correctly at expiry granularity. The most evidence-backed *addition* would be an option-momentum input — past 6–36-month straddle returns predict next-month option returns with decile spreads ≈ 3.6%/month (Heston et al., JF 2023 [38]).

4.4 Strike selection and the earnings buffer

Delta choice: no academic 16Δ-vs-30Δ study exists; the surface evidence says higher delta = more premium per stress loss, lower delta = smoother equity curve, no free lunch [28]; the largest practitioner grid found best per-trade normalized P&L at 16–20Δ with 73–80 DTE [33]. The 0.30Δ default is the premium-capture end of a documented preference axis — worth a user toggle, not a fix.

Earnings buffer — DEVIATION: implied earnings-jump vols average ~8.5–10%; implied exceeds realized by ~20–30% on average, and *scaled* announcement moves are near-normal [47] — so a 1.25× expected-move push is ≈1.25σ ≈ **79% one-event containment**, and raw moves are fat-tailed on top (≈2.5% of mass in each extreme tail [48]). For the human path that trades through earnings, 1.5–2.0× EM (≈87–95%) is the supported buffer; the agent's default blackout (and whitelist-with-conditions) is

the cleaner posture the literature prefers for unattended selling.

---

5. Conviction Engine

5.1 IV/HV conditioning — ALIGNED (PARTIAL)

Goyal & Saretto remains the direct support for the rich/cheap axis (sell rich-vs-HV, buy cheap-vs-HV; economically large decile spreads [49]). Two modern corrections: **Duarte et al. (RFS, forthcoming) show much of the published option-anomaly alpha — the Goyal-Saretto family included — was inflated by look-ahead quote filters**; the sign survives, the retail-realizable magnitude is far smaller [50]. And the effect is a *deciles* phenomenon — strongest at extremes — while the engine draws lines at 1.20/0.85. Refinements: percentile-rank IV/HV within the scanned universe rather than absolute cutoffs; note the tenor mismatch (G&S used 12-month HV vs 1-month IV; a 20-day HV vs 45-DTE IV is noisier); and carry post-Duarte humility in any displayed edge expectations. (The prior review's recommendation to promote NVRP to the primary gate stands and is strengthened by this.)

5.2 The skew bonus — DIVERGENT (sign risk)

The conviction engine awards up to +5 for "aligned" skew (BUY & inverted call skew; SELL & steep-normal call skew), treating surface shape as harvestable edge. The literature says two things against this: (1) **there is no independently harvestable skew premium** — the SPX skew premium is ~0.9-correlated with the variance premium, and a variance-hedged skew trade earns statistically zero [56][57]; whatever the bonus is capturing, the IV/HV component already prices it. (2) **Skew shape is directional information**: steep put smirks predict ~10.9%/yr *underperformance* [58]; rich/inverted call wings often reflect informed call buying [58b]. Selling put spreads into a steep put skew is selling insurance exactly where informed flow says downside risk is elevated; selling call spreads into an inverted call skew shorts the tail the market is actively pricing.

Recommendation: repurpose skew from a *bonus* to a *directional modifier/veto* — steep put skew downgrades bullish credit conviction, inverted call skew adds caution to bear-call structures — and let the premium-richness question remain with IV/HV-NVRP. (The skew_analyzer's edge-detection for *strike selection within a chosen structure* is a narrower, more defensible use than the conviction-level bonus.)

5.3 GEX: regime yes, walls no — SPLIT

The engine's two GEX uses now have opposite evidential status. **Regime caps** (negative aggregate gamma → trendy/fragile, positive → range/vol-suppressed) are supported by peer-reviewed work: dealer-gamma imbalance predicts intraday momentum when negative and reversal when positive [62]; hedging-demand momentum replicates across 60+ markets [63]. **Strike-local wall nudges** now have a direct null: a pre-registered, control-matched SPY test found the top GEX "wall" statistically indistinguishable from rank-2–5 strikes across break-resistance, pinning, and vol-dampening tests, under multiple GEX definitions [64]. The one strike-local effect that *is* documented is expiration-day pinning to high-OI strikes [20] — an expiry-week, cents-scale effect. Recommendation: keep the regime caps and the index-only sign trust (both already correct), demote wall-based ±10 placement nudges to cosmetic/informational, and if any strike-local logic is kept, tie it to expiry-week pinning rather than walls.

5.4 The VIX bias layer — DEVIATION (mitigated by design)

Three results cut against level-based structure preference: **low VIX does not mean cheap options** — the vol premium persists in calm markets, so buying protection because VIX < 15 is still a losing

trade on average (Israelov & Nielsen, "Still Not Cheap" [60]); **high VIX does not mean good selling** — the ex-ante short-vol premium falls and often inverts exactly when risk spikes (Cheng [61]); and **slope/VRP measures dominate the level** in the one clean head-to-head design (Johnson [59]). The engine's mitigations are real: vix_bias is annotation-only and never overrides; the agent's VIX guard conditions on FEAR/PANIC *plus backwardation* (a composite the literature supports); and the term-slope sizing rule was already the best-supported VIX feature in the prior review. Remaining fixes: rewrite the COMPLACENT→DEBIT leg (its rationale is wrong even as annotation copy), and add hysteresis to regime labels — vol regimes persist for weeks-to-months (expected durations ~50–80 weeks for low/medium regimes [65]), so a label that can flip on a 0.1-point VIX print is noise; enter/exit bands (e.g., 14/16 instead of a hard 15) cost nothing.

5.5 Structure expectations — condors and straddles

- **Iron condors:** the CNDR benchmark earned ~9.8% CAGR 1988–2010, then **−0.9%/yr for 2010–2019** [54][55] — condors are a volatility-dampened exposure with a small, regime-dependent edge and decade-long flat stretches. The HOLD+rich mapping is consistent with the VRP literature, but user-facing expectation-setting should reflect the modern subsample.
- **Long straddles — DEVIATION:** the catalyst gate (≤ 14 days) is ~3× wider than the documented effect. GXZ's profitable window is entry ~3 trading days pre-announcement (avg +3.34% into the event; deteriorating sharply with distance and with DTE beyond ~10 days) [51][52]. Tighten the gate to ≤ 5 trading days, or stage the recommendation ("straddle candidate — enter T−3"). Vasquez's slope signal [37] is a second, independently supported straddle gate the engine already has the data for (term_slope) but does not currently apply to the straddle decision.

---

6. Trade Plan & POP Accuracy

6.1 The probability object — ALIGNED (PARTIAL), one quantified gap

N(d2) at the break-even — not delta — is the correct Black-Scholes object, and the delta-proxy fallback ordering is right. The literature adds three calibrated expectations:

1. **The missing skew-slope term.** The exact Breeden-Litzenberger digital probability is N(−d2) at the strike's own IV **plus vega × ∂σ/∂K** [40][41]. The component uses the strike's smile IV (level) but omits the slope term. Magnitude at a ~25-delta 30-day index put with typical skew: ≈ **+3.5–4.5 probability points of P(ITM)** (~20% relative); at 10-delta, smaller absolute but larger relative error. Direction: the short put is riskier than displayed; the short call less so. A finite-differenced smile slope closes most of the gap for two lines of code.
2. **Asymmetric calibration by side.** Risk-neutral probabilities embed the equity premium and VRP: realized win rates for short index puts run *above* delta/N(d2)-implied rates while short calls run at-or-below (large practitioner samples [44]; RND calibration tests fail without a risk-premium tilt of CRRA ≈ 2–4 — Bliss & Panigirtzoglou [39]). The calibration report should expect and test by direction, not pooled.
3. **Favorite-longshot confirmation.** Deep-OTM index options exhibit the longshot bias — all puts overpriced, deep-OTM calls longshot-priced [45] — favorable to a seller *on frequency* and silent on severity; the existing max-loss/CVaR displays are the correct complement to POP-based verdict caps.

6.2 The cone — ALIGNED, with a coverage expectation

The symmetric-cone disclosure ("downside understated") matches the leverage-effect literature, and the drawn ±1σ bounds will empirically *over-cover*: with the VRP, index outcomes land inside the implied ±1σ move materially more often than 68% (fair-value straddle math alone puts inside-the-breakevens at ~58% for ±0.8σ [43]; VRP pushes realized containment into the 70s). Calibration guidance: validate cone coverage against *expected-with-VRP* rates, not nominal 68/95 — a cone that covers exactly 68% would actually be evidence the σ input was too *narrow*. GBM touch probabilities remain lower bounds for downside barriers under equity skew (stochastic-vol barrier pricing differs from GBM by up to ~10% of payoff [42]); the existing disclosure covers this.

6.3 Managed exits and the calibration loop — GAP (upgrade path)

With 50%-profit/21-DTE management, realized win rates exceed hold-to-expiry POP *by construction* (P(touch profit target before expiry) > P(profit at expiry); practitioner grids show managed win rates in the high-90s vs entry POPs in the 70s [46][33]). The component's calibration report carries the right caveat; the literature suggests two concrete upgrades: (a) track the **counterfactual expiry outcome** per trade (what the position would have done held to expiry) and calibrate POP against that — the cleanest apples-to-apples; or (b) publish a separate Monte-Carlo **first-passage-to-target probability** ("P(hit 50% profit before 21 DTE)") as the managed-trade forecast, since no closed form exists for a barrier in option-price space.

Methodology (power math, computed): detecting that a predicted 0.70 is really 0.75 at 80% power needs ≈ **620 observations per bucket**; 50–100 trades per quintile detects only ±10-point errors. The report should display Wilson/consistency bands (fleet metrics already does), state the minimum detectable effect at current n, use CORP/PAV-style reliability curves as n grows, and recalibrate with Platt (not isotonic) below ~1,000 labeled trades — consistent with the prior review's ML findings.

---

7. Consolidated Findings

Divergences:

#	Component	Finding	Evidence	Action
D1	Tech Charts	Hurst gate uninformative and trend-biased at engine sample sizes	[22][23]	Drop or replace with variance-ratio test; least weight meanwhile
D2	Conviction	Skew bonus rewards selling into directional skew information; no separable skew premium	[56][58]	Repurpose skew as directional modifier/veto
D3	Conviction	Straddle catalyst gate 14 d vs documented 3–5-day edge window	[51][52]	Tighten to ≤5 trading days or stage entry at T−3
D4	Conviction	GEX wall-based strike nudges have a pre-registered null	[64]	Demote to informational; keep regime caps

Gaps:

#	Component	Finding	Evidence	Action
G1	POP	Missing BL skew-slope term: short-put P(ITM) understated ~2–5 pts at 10–30Δ	[40][41]	Add finite-differenced vega×slope correction
G2	POP	No counterfactual-expiry or first-passage-to-target calibration object for managed trades	[46]	Track counterfactual outcomes or MC P(hit target)
G3	Expiry Optimizer	Earnings buffer 1.25× EM ≈ 79% containment	[47][48]	1.5–2.0× EM or exclude earnings expiries on the human path
G4	Tech Charts	Busted patterns void bias; evidence says they should flip it	[16]	`failed` status carries inverted-bias annotation
G5	Expiry Optimizer	No option-momentum input (strongest recent predictor, 3.6%/mo decile spread)	[38]	Candidate scoring input from straddle-return history
G6	Conviction/VIX	No hysteresis on regime labels vs weeks-to-months regime persistence	[65]	Enter/exit bands (e.g., 14/16) or 5–10-day median VIX

Refinements: measured-move targets scaled to hit rates, veto kept conservative (R1, [2][15]); round-number snapping for box edges and stops (R2, [5][6]); ADX/BB-width as features under an HMM/jump regime layer — the reserved seam (R3, [24]); periodic detector re-validation on ledger history + pivot-prominence sensitivity (R4, [11][12]); DTE peak framed as risk-adjusted with an advanced shift-shorter toggle (R5, [28][29]); delta-preference toggle 16Δ↔30Δ mapped to the risk-metric axis (R6, [28][33]); IV/HV percentile-ranking + tenor-matched HV (R7, [49][50]); condor expectations set off the post-2010 CNDR record (R8, [54][55]); by-side POP calibration with VRP-adjusted cone coverage targets (R9, [39][44]); COMPLACENT→DEBIT annotation copy rewritten (R10, [60]).

Where the components are ahead of common practice: veto-only, tighten-only pattern architecture (§3.1); the hard 21-DTE floor against a now-documented 0DTE retail loss pool; POP-only-downgrades verdicts; the cone's self-disclosure; index-only GEX sign trust; path-honest pattern lifecycle with permanent completion; the existing calibration report with its exit-policy caveat; fail-closed execution tickets.

---

8. References

Numbering continuous; grouped by research area. URLs verified August 2026. Practitioner sources labeled (P).

Chart patterns & regime detection

[1] Osler, C.L. & Chang, P.H.K. (1995). "Head and Shoulders: Not Just a Flaky Pattern." FRBNY Staff Report 4.

[2] Chang, P.H.K. & Osler, C.L. (1999). "Methodical Madness: Technical Analysis and the Irrationality of Exchange-Rate Forecasts." *Economic Journal* 109(458), 636–661.

[3] Osler, C.L. (1998). "Identifying Noise Traders: The Head-and-Shoulders Pattern in U.S. Equities." FRBNY Staff Report 42.

[4] Savin, G., Weller, P. & Zvingelis, J. (2007). "The Predictive Power of 'Head-and-Shoulders' Price Patterns in the U.S. Stock Market." *Journal of Financial Econometrics* 5(2), 243–265.

[5] Osler, C.L. (2000). "Support for Resistance: Technical Analysis and Intraday Exchange Rates." *FRBNY Economic Policy Review* 6(2), 53–68.

[6] Osler, C.L. (2003). "Currency Orders and Exchange Rate Dynamics." *Journal of Finance* 58(5), 1791–1819.

[8] Donaldson, R.G. & Kim, H.Y. (1993). "Price Barriers in the Dow Jones Industrial Average." *JFQA* 28(3).

[9] De Ceuster, M., Dhaene, G. & Vennet, T. (1998). "On the Hypothesis of Psychological Barriers in Stock Markets and Benford's Law." *Journal of Empirical Finance* 5(3).

[10] Lo, A., Mamaysky, H. & Wang, J. (2000). "Foundations of Technical Analysis." *Journal of Finance* 55(4), 1705–1765.

[11] Nekrasov, V. (2018). Replication of LMW on 1995–2010 data: most patterns lose significance. (P)

[12] Jiang, J., Kelly, B. & Xiu, D. (2023). "(Re-)Imag(in)ing Price Trends." *Journal of Finance* 78(6).

[14] Fang, J., Jacobsen, B. & Qin, Y. (2014). "Popularity versus Profitability: Evidence from Bollinger Bands." *Journal of Portfolio Management*.

[15] Bulkowski, T. *The Measure Rule* — target hit rates by pattern (*Encyclopedia of Chart Patterns*, 2nd ed., Wiley 2005). (P)

[16] Bulkowski, T. Failure-rate study, ~14,000 patterns 1991–2008. (P)

[17] Bulkowski, T. Volume study — heavy breakout volume: +1.8 pp performance, ~2× failure, ~3× throwbacks. (P)

[19] Sullivan, R., Timmermann, A. & White, H. (1999). "Data-Snooping, Technical Trading Rule Performance, and the Bootstrap." *Journal of Finance* 54(5).

[20] Ni, S.X., Pearson, N.D. & Poteshman, A.M. (2005). "Stock Price Clustering on Option Expiration Dates." *JFE* 78(1), 49–87.

[22] Couillard, M. & Davison, M. (2005). "A Comment on Measuring the Hurst Exponent of Financial Time Series." *Physica A* 348, 404–418.

[23] Weron, R. (2002). "Rescaled Range Analysis and Detrended Fluctuation Analysis: Finite Sample Properties and Confidence Intervals."

[24] Shu, Y., Yu, C. & Mulvey, J. (2024). "Downside Risk Reduction Using Regime-Switching Signals: A Statistical Jump Model Approach." *Journal of Asset Management* 25.

[27] (Gap noted in text: no published work links technical levels to option short-strike safety; nearest anchors [5][20].)

Expiration selection & option-return term structure

[28] Israelov, R. & Tummala, H. (2017). "Which Index Options Should You Sell?" SSRN 2990542.

[29] Dew-Becker, I., Giglio, S., Le, A. & Rodriguez, M. (2017). "The Price of Variance Risk." *JFE* 123(2).

[30] Beckmeyer, H., Branger, N. & Gayda, L. (2023). "Retail Traders Love 0DTE Options... But Should They?" SSRN 4404704.

[31] Andries, M., Eisenbach, T., Kahn, R.J. & Schmalz, M. "The Term Structure of the Price of Variance Risk." FRBNY Staff Report 736.

[32] Vilkov, G. "0DTE Trading Rules." SSRN 4641356; Dim, C., Eraker, B. & Vilkov, G. "0DTEs: Trading, Gamma Risk and Volatility Propagation." SSRN 4692190.

[33] DTR Trading. SPX strangle/condor backtest grids (56k+ variants). (P)

[34] OptionsPilot. "Weekly vs Monthly Options for Selling Premium: 10 Years of Data." (P)

[35] Days to Expiry. "The 21 DTE Rule Explained" (summarizing tastylive ~200k-trade research; gamma 3–5× in final 21 days). (P)

[36] projectoption / Days to Expiry. OTM vs ATM theta-decay curve treatments. (P)

[37] Vasquez, A. (2017). "Equity Volatility Term Structures and the Cross-Section of Option Returns." *JFQA* 52(6).

[38] Heston, S., Jones, C., Khorram, M., Li, S. & Mo, H. (2023). "Option Momentum." *Journal of Finance* 78(6).

POP accuracy & probability calibration

[39] Bliss, R.R. & Panigirtzoglou, N. (2004). "Option-Implied Risk Aversion Estimates." *Journal of Finance* 59(1) (RNDs rejected raw; pass with CRRA ≈ 2–4 tilt); Berkowitz, J. (2001). "Testing Density Forecasts." *JBES* 19(4).

[40] Breeden, D. & Litzenberger, R. (1978). "Prices of State-Contingent Claims Implicit in Option Prices." *Journal of Business* 51(4); digital-option skew-sensitivity treatments (Quant Next; Bahra, Bank of England WP 1997).

[41] Gatheral, J. (2006). *The Volatility Surface* (digital pricing with smile slope). Magnitudes in text computed from standard formulas.

[42] Thul, M.; Haugh, M. Barrier-option model risk under local vs stochastic vol; Wystup, U. (2006). *FX Options and Structured Products* (one-touch deviations ≈ ±10% of payoff).

[43] Brenner, M. & Subrahmanyam, M. (1988). ATM straddle ≈ 0.8·S·σ√T; Moontower analysis of straddle win rates. (P)

[44] spintwig aggregate short-put/short-call win-rate-vs-delta findings; Option Alpha "Moment in Time" calibration study (>1M trades, Brier ≈ 0.16–0.20). (P)

[45] Hodges, S., Tompkins, R. & Ziemba, W. (2003). "The Favorite/Long-shot Bias in S&P 500 and FTSE 100 Index Futures Options." SSRN 424421; Snowberg, E. & Wolfers, J. (2010). NBER w15923; Bondarenko, O. (2014). "Why Are Put Options So Expensive?"

[46] DTR Trading profit-target grids; spintwig managed-vs-held backtests; Dimitriadis, T., Gneiting, T. & Jordan, A. (2021). "Stable Reliability Diagrams." *PNAS*; Niculescu-Mizil & Caruana (2005), ICML. Power computations in text from binomial formulas.

[47] Dubinsky, A. & Johannes, M. (2006/2019). "Earnings Announcements and Equity Options" / "Option Pricing of Earnings Announcement Risks." *Management Science* (implied jump vols 8.5–10.4%; P/Q ≈ 0.74–0.82; scaled moves near-normal).

[48] Lipkin, M. & Tatevossian, L. "Earnings Moves and Pre-Earnings Implied Volatility." *Journal of Risk* (fat raw tails, ~2.5% per extreme tail); ORATS 343-stock earnings study. (P)

Regime-conditioned strategy selection

[49] Goyal, A. & Saretto, A. (2009). "Cross-Section of Option Returns and Volatility." *JFE* 94(2).

[50] Duarte, J., Jones, C., Khorram, M., Mo, H. & Wang, J. (forthcoming). "Too Good to Be True: Look-ahead Bias in Empirical Options Research." *RFS*.

[51] Gao, C., Xing, Y. & Zhang, X. (2018). "Anticipating Uncertainty: Straddles Around Earnings Announcements." *JFQA* 53(6) (3–5-day pre-announcement window; deteriorates with distance and DTE > ~10).

[52] CXO Advisory / Alpha Architect summaries of [51] with window-by-window magnitudes. (P)

[54] Cboe CNDR Index methodology; Black, K. & Szado, E. (2016). CBOE options-selling benchmark study (CNDR stdev 7.23%); Wilshire (2019) benchmark study 1986–2018.

[55] CNDR sub-period decomposition: ~9.8% CAGR 1988–2010, −0.9%/yr 2010–2019. (P)

[56] Kozhan, R., Neuberger, A. & Schneider, P. (2013). "The Skew Risk Premium in the Equity Index Market." *RFS* 26(9).

[57] Cao, J. & Han, B. (2013). "Cross-Section of Option Returns and Idiosyncratic Stock Volatility." *JFE*.

[58] Xing, Y., Zhang, X. & Zhao, R. (2010). "What Does the Individual Option Volatility Smirk Tell Us About Future Equity Returns?" *JFQA* 45(3). [58b] Cremers, M. & Weinbaum, D. (2010). *JFQA* 45(2).

[59] Johnson, T.L. (2017). "Risk Premia and the VIX Term Structure." *JFQA* 52(6).

[60] Israelov, R. & Nielsen, L. (2015). "Still Not Cheap: Portfolio Protection in Calm Markets." *Journal of Portfolio Management* 41(4).

[61] Cheng, I.-H. (2019). "The VIX Premium." *RFS* 32(1).

[62] Barbon, A. & Buraschi, A. (2021). "Gamma Fragility." SSRN 3725454.

[63] Baltussen, G., Da, Z., Lammers, S. & Martens, M. (2021). "Hedging Demand and Market Intraday Momentum." *JFE* 142(1).

[64] Popovici, R. (2025/26). "Do Gamma-Exposure Walls Exhibit Strike-Local Effects? A Pre-Registered, Control-Matched Test in SPY." SSRN 7082418.

[65] ECB Financial Stability Review (May 2018). "Regime Shifts in Stock Market Volatility" (regime durations ~50–80 weeks); regime-switching SV literature (self-transition > 0.97 daily).

Prepared from COMPONENT_SPEC.md v3.95.0 and ~90 primary sources across four targeted research sweeps. Practitioner-grade evidence labeled (P) throughout; reviewer-computed magnitudes flagged in text. Research commentary, not investment advice.