

Disciplined Harvesting of the Variance Risk Premium

Maksim Aleksandrovich Andreev & Joseph Howden
Quark Quantitative Research — Working Paper — August 2026

ABSTRACT

We make an unusually narrow claim. The alpha source behind our options-income programme is *not* novel: it is the well-documented variance risk premium (VRP) — the systematic gap between option-implied and subsequently realised variance that compensates the seller of optionality. The single-name VRP (Cao & Han, 2013), the index VRP (Bakshi & Kapadia, 2003), the skew premium (Bondarenko, 2014), and the negative average return to long volatility (Coval & Shumway, 2001) are all decades-old empirical regularities. The contribution of this paper is the *execution machinery* built around that known edge, because most VRP strategies do not fail at the thesis — they fail at the discipline, paying back years of carry in a single tail event (Israelov & Tummala, 2018). We describe three components. **(1) A two-gate architecture:** an *academic-alpha gate* that refuses to fire any branch unless a documented, published edge is present (per-branch thresholds derived from the source paper's mechanism), composed with a *physics-modulated risk-cap gate* imposing net-vega / notional / gamma / theta-burn ceilings plus a volatility stand-down that refuses *new* premium-selling while the market volatility index is elevated in absolute or relative terms (the defined-risk structure and sizing caps, not the stand-down, are the protection). **(2) Wavefunction pricing and selection:** legs are designed to be priced through the pilot-wave quantum potential $Q(S)$ rather than static Black–Scholes — in the current live deployment this leg-pricing layer is *shadow-measured*, not live: live execution falls back to Black–Scholes leg pricing while the quantum-mispricing signal accrues through a dedicated shadow script under a no-live-capital discipline — strikes are read from the wavefunction's support/resistance structure, and the days-to-expiry is set by the decoherence horizon τ_{decoh} rather than a fixed 30 days. **(3) A pre-registered Harvey–Liu–Zhu (2016) validation** with frozen, Bonferroni-adjusted thresholds, and the subsequent **withdrawal** of the certification that validation produced. In August 2026 a cross-period stale-frame defect was found in the backtest engine: every test window after the first was priced from the first window's cached frame, so the underlying was pinned at a single May-2020 close and the reconstructed signals never advanced. Short-premium structures then settled at an unchanged spot, which removes most of the within-sample variance and inflates the pooled statistic the frozen bar was applied to. Re-derived on the repaired engine, none of the eight branches clears its bar, and the long-debit bull-call spread that the withdrawn table recorded as disconfirmed at $t = -19.41$ reads $t = +0.82$, so that disconfirmation is retracted together with the passes. Since the four-window design cannot reach the bar at its own trade counts even on a correct engine, we retire it as a certification instrument rather than re-tune it. Section 6 reports what survives: a twenty-year characterisation of the two income structures over 239 monthly observations, and a subperiod persistence re-read, each stated with the riders that bound it. This is a Phase-2 MVP whose backtest used proxy reconstructions of the live physics signals and synthesised option prices; live fills will differ. We claim disciplined, gated, honestly reported VRP execution — not a new alpha, and no longer a certification.

1. The Variance Risk Premium: A Documented Edge, Not a Discovery

The variance risk premium is the difference between the variance implied by option prices and the variance subsequently realised by the underlying. It is systematically positive: on average, options are priced as if future variance will be higher than it turns out to be, so the writer of optionality is paid an insurance

premium. This is one of the most robust empirical facts in asset pricing, and we lay out its lineage precisely because we are *not* claiming to have found it.

The literature we are harvesting.

- **Bakshi & Kapadia (2003)** show, via delta-hedged option gains, that the market index carries a negative volatility risk premium — index options are systematically expensive relative to realised variance.
- **Cao & Han (2013)** extend this to the *cross-section of single names*: delta-hedged option returns are decreasing in idiosyncratic volatility, and shorting the richest single-name VRP is a documented, persistent source of return.
- **Bondarenko (2014)** isolates the *skew premium* — the extra richness embedded in out-of-the-money puts beyond the at-the-money VRP — the component a put-spread seller harvests.
- **Coval & Shumway (2001)** document that *long* volatility positions (zero-beta straddles) earn significantly negative average returns: the mirror image of the seller's premium, and the reason a long-vol structure must clear a much higher bar to be justified.
- **Carr & Wu (2009)** and **Bollerslev, Tauchen & Zhou (2009)** quantify the premium's size and its time-variation, establishing that it is large, persistent, and predictably state-dependent.

The economic content is standard. Selling optionality is selling insurance; the buyer pays a premium for protection against variance, and on average that premium exceeds the cost of the claims paid out. The seller earns carry in exchange for bearing the tail. None of the strategies in this paper deviate from that mechanism. What we contribute is not a new reason the premium exists, but a disciplined apparatus for collecting it while surviving the tail it embeds.

2. Why Disciplined Execution Is the Gap

If the VRP is so well documented, why is it not trivially monetised? Because the premium is compensation for a real, severe, negatively-skewed risk, and the naive harvest of it is fragile in exactly the way insurance underwriting is fragile: long stretches of small gains punctuated by rare, large losses. The strategy that sells the premium without discipline is the one that blows up.

Israelov & Tummala (2018), "An Alternative Option to Portfolio Rebalancing," and Israelov (2017), "Pathetic Protection." The central warning we build around: a short-volatility carry strategy's apparent attractiveness is an artefact of measuring it outside its tail. The premium *is* the tail risk; the carry is funded by the rare loss. A defensible VRP harvest must therefore (i) be defined-risk where possible, (ii) size capital-at-risk small, and (iii) stand down in regimes where the tail is most likely to materialise. Equally, Israelov (2017) shows the symmetric error on the protection side: buying puts as standing insurance carries strongly negative expected return outside stressed regimes — protection is only rational when the tail is priced as live.

An earlier version of this paper illustrated the hazard with per-period profit and loss from our own Phase-2 backtest, contrasting a cash-secured put that earned in calm windows and lost more in the 2020 window against a defined-risk bull-put spread that crossed the same window with a bounded loss. Those per-period figures are withdrawn with the rest of that backtest (Section 6): the defect described there froze the underlying after the first window, so the calm-window cells and the crisis-window cell were not priced on comparable data and the contrast between them is not measurement. The structural point does not depend on our numbers and is Israelov and Tummala's own: an undefined-risk short put and a defined-risk put spread harvest the identical skew premium, and only the purchased wing bounds what the tail can take back. This is the gap the paper targets, which is not "is there a premium?" but "can it be collected through a tail-aware, gated, honestly reported machine?"

Framing of the contribution. We treat the VRP as a fixed, exogenous, citeable input and ask whether a disciplined execution layer — gating, wavefunction pricing/selection, and pre-registered validation — converts it into a survivable harvest. Every empirical claim below is a claim about the *execution* wrapper, not about the existence of the premium.

3. The Two-Gate Architecture

Every candidate option structure must pass two sequential gates before any order is placed. The first gate enforces that a documented academic edge is actually present; the second imposes physics-modulated risk caps and a volatility stand-down. Refusals from either gate are logged with a typed reason so that ex-post analysis can ask what would have happened had the refused trade fired.

3.1 Gate 1 — the academic-alpha gate

The academic-alpha gate is the structural defence against the most common failure mode of a signal-driven options system: firing a quantum-priced, mechanically-correct trade that nonetheless has no documented edge behind it. A strategy that is correctly priced but rests on no published alpha source is still, in expectation, a losing trade. The gate therefore refuses any branch unless the specific published premium that branch is supposed to harvest is measurably present, with a per-branch threshold derived from the source paper's mechanism (not in-sample tuned).

Table 1: The academic-alpha gate, by branch

BRANCH	DOCUMENTED ALPHA SOURCE	GATE THRESHOLD
Cash-secured put / covered call	Cao & Han (2013) single-name VRP	stand-down not armed AND $((IV-RV)/RV \geq 10\%$, or vol-classifier RICH/VERY RICH)
Iron condor	Bakshi & Kapadia (2003) index VRP + quantum mispricing	stand-down not armed AND quantum-overprice $\geq 1.5\%$
Bull-put spread	Bondarenko (2014) skew premium	stand-down not armed (the Israelov–Tummala defensive clause) AND $(IV-RV)/RV \geq 5\%$
Bull-call spread	Cremers & Weinbaum (2010) IV-spread + drift	signal strength $\geq 1.0\sigma$ AND stand-down not armed
Protective put	Israelov (2017) — negative carry outside stressed regimes	fires <i>only</i> when the stand-down is armed on a measured print
Bear-put spread	No standalone alpha; stressed-volatility tail hedge only	fires <i>only</i> when the stand-down is armed on a measured print
Long straddle	Coval & Shumway (2001) — long-vol negative carry	fires only inside the pre-earnings window (Patell–Wolfson 1979 IV crush is the only edge); the earnings calendar has been wired since May 2026, and the gate fails closed whenever it is unavailable

Amendment of the regime clause (2026-08-22).

Every row of Table 1 previously carried a threshold on the instanton-derived crash probability: short-premium branches required it below a floor, long-protection branches above one. Those clauses were replaced by a volatility stand-down keyed to the market volatility index, armed when the index stands at or above 28 in absolute terms or above 1.3 times its trailing twenty-session mean, and armed by default when the print is missing, implausible or stale, so that a dead feed refuses short premium rather than permitting it. The replacement was made on the forward record of the quantity it removed. Over that record nine days read a 21-day crash probability between 59% and 82%; on the four windows that have since completed, the index returned +1.96%, +0.22%, +5.06% and +3.59% over the following 21 sessions, and days reading zero did marginally worse forward than days reading between 60% and 80%. A quantity that does not order forward outcomes cannot carry a threshold, and the same reconstructed gate defaulted to zero whenever its engine was unavailable, which passed every floor written to constrain it. The crash probability is retained on every gate record as a printed readout with its standard error, so that its forward record continues to accrue, and it decides nothing. The description of these clauses as crash kill switches, which appeared in earlier versions of this paper, is withdrawn.

Read the table as a discipline, not a menu. Three of the seven rows are gated to *refuse by default*: the protective put and bear-put spread fire only while the stand-down is armed on a measured print (Israelov's insurance logic, in which protection is rational only when the tail is priced as live), and the long straddle fires *only* inside the pre-earnings window, because its sole documented edge is the pre-earnings implied-volatility crush of Patell & Wolfson (1979); the earnings calendar feeding that window has been wired since May 2026, and the gate fails closed whenever the calendar is unavailable at runtime. (The backtest rows above predate the wiring, which is why the branch recorded zero trades there.) A gate whose default action for several branches is "do not trade" is doing its job.

3.2 Gate 2 — the physics-modulated risk-cap gate

A trade that clears the academic gate next faces a quantitative risk-cap gate. This gate imposes ceilings on the book's net option exposures — net vega, net notional, net gamma, and theta-burn — with the ceilings *modulated* by the live physics state rather than fixed. The modulation is conservative: while the volatility stand-down is armed the vega and notional ceilings tighten and the gate refuses all *new* premium-selling outright. The scope of that refusal must be stated precisely, because two settled 2026-07 studies bound what any such regime gate can honestly claim. The gate refuses *new* trades only; it does not close open positions, and we do not claim prospective protective value for it. The reconstructed instanton crash gate that occupied this position until August 2026 proved contemporaneous-to-lagging rather than leading, since only 3 of its 21 fires in the tail-overlay settlement preceded a real drop and it was structurally blind to fast gaps that begin from low correlation; and closing open defined-risk income spreads on its fire was the *worst* tested de-risk policy, at a Calmar of 0.016 against 0.251 for holding to expiry, because the close realises at the bottom a loss the defined-risk structure would mostly have healed through. Those two findings are why the successor stand-down governs entry only. The protection in this architecture is the defined-risk structure itself, whose purchased wing caps the loss, together with the sizing caps of Section 6.1; the stand-down expresses the Israelov–Tummala tail discipline only in the narrow sense that new carry must not be added into the regime that funds its loss.

```

Gate 2 (schematic):
  reject if net vega(book + candidate) >
Vcap(stand-down state)
  reject if net notional >
Ncap(stand-down state)
  reject if |net gamma| >
Gcap
  reject if theta burn >
Θcap
  REFUSE all new premium-selling if the
volatility stand-down is armed
      (open defined-risk positions
are held to expiry, not closed on the gate)

  stand-down armed if VIX ≥ 28
                    or VIX > 1.3 × mean(VIX
over the trailing 20 sessions)
                    or the print is missing,
implausible or stale (fail closed)

```

The two gates compose: academic edge present AND risk caps respected AND not not standing down on volatility. The conjunction is intentionally strict. The cost of strictness is forgone trades, visible in the refusal counts recorded against every branch; the benefit is that every executed trade has both a documented edge and a bounded, regime-aware risk footprint.

4. Wavefunction Pricing and Strike / DTE Selection

The second component of the contribution is that, once a trade clears both gates, its strikes and expiry are selected from the pilot-wave wavefunction state (Bohm, 1952), and its legs are designed to be priced the same way rather than from static Black–Scholes inputs. One deployment caveat is stated up front and expanded in Section 4.1 and the caveats of Section 7: the Q(S) leg-pricing layer is currently *shadow-measured*, not live — live execution falls back to Black–Scholes leg pricing while the quantum-mispricing signal accrues through a dedicated shadow script under a no-live-capital discipline. This is the same stochastic-mechanical apparatus used elsewhere in the platform (Andreev & Howden, 2026a): under the Madelung–Nelson transform the Black–Scholes

pricing equation is the zero-potential (free-particle) limit of a Schrödinger equation — a correspondence developed at length in the quantum-finance literature (Baaquie, 2004) — and the empirically non-trivial structure lives in the quantum potential Q(S).

4.1 Leg pricing via the quantum potential Q(S)

Each leg is designed to be priced through the full pilot-wave price rather than a single flat-volatility Black–Scholes value. The quantum potential $Q(S) = -(\sigma^2/2) \cdot (\partial_S^2 \sqrt{\rho}) / \sqrt{\rho}$ encodes the curvature of the price density and shifts the fair value of a leg relative to the constant-volatility benchmark. The intended effect is a state-dependent skew correction: the price the engine uses would already reflect the support/resistance structure of the underlying's density, rather than imposing a single σ across all strikes. *Deployment status*: in the current live system no execution site passes a live implied-volatility surface into the full quantum pricer, so live leg pricing falls back to the Black–Scholes value with a zero quantum correction. The quantum-mispricing measurement instead accrues through a dedicated shadow script, under a no-live-capital discipline, until its forward sample reaches quorum. The Q(S) leg-pricing machinery described here is therefore the shadow-measured, forward-accruing design — not the currently operative live pricing layer.

4.2 Strikes from the wavefunction

Strikes are read from Q(S) directly: short calls are placed at quantum-potential *peaks* (resistance, where the density is repelled) and short puts at quantum-potential *wells* (support, where the density is attracted). This replaces the static out-of-the-money grid (e.g. fixed S·0.95 / S·1.05) with strikes located at the wavefunction's own turning points. When live wavefunction data is unavailable the engine falls back to a delta-based grid, a fallback flagged in the caveats.

4.3 Expiry from the decoherence horizon τ_{decoh}

Days-to-expiry is set by the decoherence horizon τ_{decoh} — the timescale over which the correlation eigenstructure rotates — rather than by a fixed 30-day convention. The rule is "short $\tau \rightarrow$ short DTE": when

the eigenstructure is stable (long τ_{decoh}) the engine is willing to write range-bound structures (e.g. condors) at longer horizons; when the eigenstructure is rotating fast (short τ_{decoh}) it shortens the expiry so the position is not held across a structural change it cannot price.

```
selection summary:
  strike(call) = nearest Q(S) peak above spot
(resistance)
  strike(put)  = nearest Q(S) well below spot
(support)
  DTE          = f( $\tau_{\text{decoh}}$ ): stable -> longer,
rotating -> shorter
              (income branches clamped to the
validated [21, 40]-day band;
              the ~30-day hold-to-expiry
lifecycle survived two exit-redesign
              nulls, so the modulation moves
DTE only within that band)
  size          proportional to the signal
ratio = |v| / (|v| + |u|)
              (Nelson current vs osmotic
velocity decomposition)
```

What this does and does not claim. Wavefunction pricing and selection are an *execution refinement* on top of a documented premium, not an independent alpha source. We do not claim $Q(S)$ predicts returns; the returns thesis (physics as a return predictor) settled null elsewhere in the programme. The claim here is narrower: pricing legs and locating strikes/expiry from the density structure is a more state-aware way to harvest the VRP than a flat- σ static grid. Whether that refinement adds measurable value over a Black–Scholes-priced version of the same gated strategy is itself a question for forward testing; the Phase-2 backtest below validates the gated *strategies*, not the pricing refinement in isolation.

5. Pre-Registered Harvey Validation Protocol

The third component is validation discipline. Each branch was assigned a pre-registered Harvey, Liu & Zhu (2016) threshold *before* the backtest was run, with the threshold Bonferroni-adjusted for the number of branches tested. The point of pre-registration is to

make the disconfirmations binding: a branch that fails its frozen threshold is disabled, not re-tuned until it passes.

Pre-registered protocol.

(P1) Underlying: SPY, across four stress regimes — 2020 COVID crash, 2022 rate shock, 2023 control, 2024–25 recent.

(P2) Each branch trades only when both gates of Section 3 pass; refusals are logged with typed reasons.

(P3) Per-branch P&L is pooled across the four periods; the test statistic is the pooled t -statistic on per-trade returns, computed with a Newey–West heteroskedasticity-and-autocorrelation-consistent variance (Bartlett kernel, automatic bandwidth). The legacy identically-distributed values were retained alongside the corrected ones for audit transparency.

(P4) Decisive bar: Bonferroni-adjusted Harvey threshold for the eight branches tested, $|t| > 2.91$ ($\approx$ the Harvey–Liu–Zhu $t > 3.0$ rule of thumb). A branch with $|t|$ below the bar is FAIL; a branch with a large *negative* t is DISCONFIRMED and disabled; a branch with too few trades is INSUFFICIENT (no verdict).

(P5) No in-sample tuning of gate thresholds after the verdict; a disconfirmed branch is hard-disabled and may be re-enabled only by a fresh pre-registered re-test under a reformulated alpha source.

The runner is the dedicated Harvey backtest script over the quantum options backtest engine, and the machine-readable verdicts are persisted to an archived results file in the platform's data store. The protocol above is reported as it was frozen, because the reason its output was withdrawn lies in the engine's data path rather than in the design of the test, and Section 6 cannot be read without it.

6. Results

Withdrawal of the Phase-2 certification (2026-08-04). The pooled verdict table this section previously reported, and the two figures that plotted it, are withdrawn in full. The backtest engine served every test window after the first from the first window's cached frame: the spot used for strike selection and for settlement stayed at a single May-2020 close of 304.32 across 2022, 2023, 2024 and 2025, and the reconstructed signal history stayed at its May-2020 values throughout. A short-premium structure written against an underlying that never moves settles at the strike it was written around, which is why several per-period cells returned every trade a winner at near-zero within-cell variance and why the pooled statistics were as large as they were. An earlier version of Section 7 attributed those saturated statistics to all-winning small samples; that attribution was itself wrong, and the cause was the frame defect. The defect surfaced only when real dated option chains were substituted for synthesised prices and refused the stale strikes, which did not exist on the dates the engine claimed to be trading them. It is fixed and regression-locked.

Re-derived on the repaired engine, the four-window design certifies nothing. Table 2 sets each withdrawn statistic beside its re-derivation on the same frozen bar of $|t| > 2.91$, for the branches that reached a testable sample under either engine.

Table 2: Withdrawn Phase-2 statistics and their re-derivation on the repaired engine, against the frozen Bonferroni bar $|t| > 2.91$. The verdict column of the withdrawn table is not reinstated in any form: no branch reaches the bar in either direction, and the design is retired below rather than re-run.

BRANCH	WITHDRAWN STATISTIC	REPAIRED ENGINE	STATUS AFTER WITHDRAWAL
Covered call	+13.36	−0.62	Pass withdrawn. Not certified here.
Bull-put spread	+6.67	−0.99	Pass withdrawn. Not certified here.
Cash-secured put	−0.54	−1.42	Below the bar under both engines; no verdict claimed from this design.
Bull-call spread	−19.41	+0.82	Disconfirmation retracted. See the paragraph below.
Protective put	+0.92	+1.60	Sample too small under both engines; no verdict claimed.
Iron condor, bear-put spread, long straddle	no testable sample under either engine		No verdict claimed. The condor's mispricing gate refused every reconstructed candidate, which is a statement about the reconstruction and not about the structure.

Retraction of the bull-call disconfirmation. Earlier versions of this paper presented the long-debit bull-call spread's $t = -19.41$ as the clearest demonstration that the programme disables its own losing structures. On the repaired engine that branch reads $t = +0.82$. The sign of the number the paper used as its honesty exhibit was an artefact of the same defect, and the disconfirmation is retracted without qualification. The branch remains disabled in the live configuration as a conservative default, since a long-debit structure pays the premium the income branches collect and nothing here argues for turning it on; but that disable now rests on the structure's carry profile and on the absence of any evidence either way, not on a Harvey verdict. The distinction matters because a retracted disconfirmation is not a confirmation.

The same branches were then re-run on real dated chains, priced at the quoted bid and ask rather than synthesised, which removes the pricing objection but not the sample-size one. Every branch mean turned positive, at +0.262 for the covered call, +0.121 for the bull-put spread and +0.047 for the iron condor per premium-normalised trade, and the largest statistic anywhere in that re-run was 1.57 at roughly 41 trades. Four stress windows on a single underlying cannot produce enough non-overlapping monthly entries to reach a Bonferroni-adjusted Harvey bar whatever the effect is, so the design is retired as a certification instrument rather than relaxed until it passes. What follows is the evidence that does not depend on it.

6.1 The surviving evidentiary basis: a twenty-year component characterisation

A separate twenty-year characterisation of the two income structures — the dedicated volatility-risk-premium backtest on SPY with the volatility index, 239 monthly observations spanning 2008, 2020 and 2022, priced with Black–Scholes — is untouched by the frame defect, because it never used the four-window engine. It is a component characterisation rather than a certification, and it agrees with the literature's tail warning. The covered-call core returned roughly +16.9% annualised at a Sharpe near 1.26 with a –38% maximum drawdown, ahead of buy-and-hold on all three but sensitive to cost assumptions. The bull-put spread showed an 87% win-rate together

with a –100% worst month at full capital-at-risk, which is the catastrophic tail that mandates small sizing. The productionised sleeve therefore caps bull-put capital-at-risk at 3% of equity per cycle and total capital-at-risk at 15%, with *new* put-selling refused while the volatility stand-down is armed and open defined-risk positions held to expiry, per the settled 2026-07 de-risk result in Section 3.2. The sizing discipline and the defined-risk structure are the operational form of the Israelov–Tummala tail constraint, and they, rather than any regime gate, carry the protection.

6.2 Subperiod persistence, and the two riders it must be cited with

Re-read by subperiod on 2026-08-04, the same twenty-year record gives a post-2010 covered-call statistic of 6.79 and a post-2010 bull-put statistic of 5.10. Neither figure may be cited without both of the riders below, which are part of the result rather than caveats appended to it.

Rider (a): this is a persistence reading, not a decay test. The split measures whether the component premium is still present in the later leg, and it is. It does not measure whether the premium has declined, because the early leg holds 41 observations at a ratio of minimum detectable effect to measured mean of about 3.6, so a decline of any plausible size would be invisible in it; and because a split at 2010 places essentially the whole financial crisis in the early leg, this sample cannot reach the pre-decline regime that a decay test would need to compare against. Read as evidence that the premium persists, the split is informative. Read as evidence that it has not decayed, it is not.

Rider (b): the cost objection is discharged by measurement, which is not the same as a new confirmation. Scored under a flat one-percentage-point transaction charge the covered-call leg fails the bar at 2.318, and that failure stood against the result until the charge was measured rather than assumed. On real dated chains at the sleeve's exact structures over 188 months, the post-2010 median toll is 0.0033 percentage points for the covered call against a break-even of 0.71, roughly 218 times headroom, and 0.168 against 3.93 for the bull-put spread, roughly 23 times; the ninetieth percentile of the toll distribution clears as well. Re-scored net of the measured toll the late legs give 6.773 and 5.015. The flat one-point charge was therefore a modelling artefact around 215 times the real quoted toll, and what this removes is a cost objection. It adds no new evidence for the premium, it does not affect rider (a), and slippage against the mid remains unmeasured, since no capture of option quotes at order submission exists; these figures are consequently an upper bound for a trader who crosses the spread.

7. Caveats and Limitations

This is a Phase-2 MVP. The following caveats are load-bearing.

(1) Proxy-reconstructed physics signals. The live engines compute the condensate magnitude $|\Phi|^2$, the decoherence horizon τ_{decoh} , and the instanton crash probability from real-time eigenstructure and dealer-flow data; since August 2026 the last of these is a printed readout on every gate record and governs no branch, per the amendment in Section 3.1. The backtest does *not* have those live inputs; it reconstructs them with simplified proxies from price-only history. No branch passed on the repaired engine, and a branch that fails on reconstructed signals may still work live if its edge depends on the richer live ones. The iron condor is the clearest case: it fired zero times because the reconstructed signal could not clear the quantum-mispricing gate, the engine having fallen back to Black–Scholes pricing for want of a live wavefunction, so the branch has no evidence either for or against it and the gate threshold it never cleared is untested rather than validated. The same fallback is, at

present, the *live* state as well, not merely a backtest limitation: no live execution site currently passes an implied-volatility surface into the full quantum pricer, so live leg pricing falls back to Black–Scholes with a zero quantum correction while the mispricing signal is shadow-measured under a no-live-capital discipline (Sections 4.1 and caveat 5). On the deeper limits of historical validation for wavefunction-calibrated systems, see Andreev & Howden (2026e).

(2) Synthesised option prices. Historical option prices are synthesised from the underlying via a Black–Scholes value with a skew adjustment, plus a 5% bid–ask spread and IBKR commission rates. Real fills will differ — in both directions — and the absolute P&L figures should be read as indicative of sign and rank, not as tradeable returns.

(3) The saturated statistics were a defect, not small samples, and this caveat previously said otherwise. Several per-period cells of the withdrawn table reported a saturated $t = 99.0$. Earlier versions of this paper explained them as artefacts of all-winning small samples in which every trade returned the same near-expiry payoff, and argued that they aggregated correctly into the pooled statistic. The first half of that explanation described the symptom accurately and the second half was wrong: the cells were all-winning because the underlying was frozen, so the pooled statistic inherited the defect rather than averaging it away. We record the mis-diagnosis because it is the reason the defect survived a caveats section written specifically to catch this class of artefact. A caveat that supplies a benign explanation for an anomaly removes the pressure to find the real one.

(4) Single underlying, and the single-name transfer test — now run and disconfirmed. The backtest is on SPY, and the single-name VRP edge (Cao & Han, 2013) is a cross-sectional claim that the SPY-only test does not exercise; the index VRP (Bakshi & Kapadia, 2003) is the relevant edge for SPY itself. The multi-underlying extension flagged in an earlier draft as a deferred follow-up has since been executed as a pre-registered, adversarially-adjudicated study (2026-07-02), and it *disconfirmed* the transfer of the edge to single-name selection.

Retraction of the single-name selection claim. An earlier draft presented “shorting the richest single-name VRP” (Cao & Han, 2013) as a harvestable edge for this system. Our own pre-registered three-arm study — top-VRP-selected names vs a size-matched random-name control vs the index, defined-risk bull-put spreads on real 2018–2025 daily OHLC with per-name implied volatility — found that single-name *selection* does *not* transfer the SPY-validated edge. The names the screen selected sat at the 9.9th percentile of uniform-random name-picking; the names it *rejected* earned more than the names it selected; and at matched realised volatility the VRP rank was roughly neutral, so the observed signal was dominantly a high-volatility tilt rather than a premium. Consequently the live income book was redirected from single-name selection to a defined-risk index short-volatility carry (SPY / QQQ / IWM) with a volatility stand-down. That redirection stands independently of the Phase-2 withdrawal, since it rests on its own three-arm study rather than on the four-window engine; the single-name *selection* claim is withdrawn.

(5) Pricing refinement not separately validated. As noted in Section 4, the Harvey test validates the gated *strategies*, not the wavefunction pricing/selection refinement in isolation. We do not claim that Q(S) pricing beats a Black–Scholes-priced version of the same gated strategy; that A/B is a separate, unrun test.

8. Conclusion

We have been deliberately unambitious about the edge and deliberately specific about the machinery. The variance risk premium is not our discovery; it is a decades-old, well-cited regularity (Bakshi & Kapadia, 2003; Cao & Han, 2013; Bondarenko, 2014; Coval & Shumway, 2001). The contribution of this paper is the disciplined execution layer that converts that known premium into a survivable harvest: a two-gate architecture that refuses any trade lacking a documented edge or violating physics-modulated risk caps and the volatility stand-down; wavefunction pricing and strike and expiry selection that locate the trade from the density structure rather than a flat- σ static grid; and a pre-registered validation protocol

which, when its data path was found defective, was used to withdraw its own output rather than to defend it.

What we are willing to defend after the withdrawal is narrower than what an earlier version of this paper claimed, and the narrowing is the paper's most useful result. The variance risk premium is present in a twenty-year component characterisation of the two income structures, it persists in the post-2010 leg of that record, and the transaction toll that was the standing objection to it has been measured on real chains and is two orders of magnitude below the break-even. None of that is a certification of a traded strategy, because the instrument built to certify one was retired when its data path was found defective, and because a persistence reading cannot substitute for the forward record that a certification would need. The alpha is borrowed from the literature. The contribution is an execution layer that gates on documented edges and bounded risk, together with a validation practice under which a verdict table survives only as long as the data path beneath it does.

References

- [1] G. Bakshi & N. Kapadia. Delta-hedged gains and the negative market volatility risk premium. *Review of Financial Studies*, 16(2):527–566, 2003.
- [2] C. Cao & B. Han. Cross section of option returns and idiosyncratic stock volatility. *Journal of Financial Economics*, 108(1):231–249, 2013.
- [3] O. Bondarenko. Why are put options so expensive? *Quarterly Journal of Finance*, 4(3):1450015, 2014.
- [4] J. D. Coval & T. Shumway. Expected option returns. *Journal of Finance*, 56(3):983–1009, 2001.
- [5] R. Israelov & H. Tummala. An alternative option to portfolio rebalancing. *The Journal of Derivatives*, 2018. (See also R. Israelov, "Pathetic protection," AQR working paper, 2017.)
- [6] P. Carr & L. Wu. Variance risk premiums. *Review of Financial Studies*, 22(3):1311–1341, 2009.

- [7] T. Bollerslev, G. Tauchen & H. Zhou. Expected stock returns and variance risk premia. *Review of Financial Studies*, 22(11):4463–4492, 2009.
- [8] M. Cremers & D. Weinbaum. Deviations from put-call parity and stock return predictability. *Journal of Financial and Quantitative Analysis*, 45(2):335–367, 2010.
- [9] J. M. Patell & M. A. Wolfson. Anticipated information releases reflected in call option prices. *Journal of Accounting and Economics*, 1(2):117–140, 1979.
- [10] F. Black & M. Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654, 1973.
- [11] C. R. Harvey, Y. Liu & H. Zhu. ...and the cross-section of expected returns. *Review of Financial Studies*, 29(1):5–68, 2016.
- [12] E. Madelung. Quantentheorie in hydrodynamischer Form. *Zeitschrift für Physik*, 40(3–4):322–326, 1927.
- [13] E. Nelson. Derivation of the Schrödinger equation from Newtonian mechanics. *Physical Review*, 150(4):1079–1085, 1966.
- [14] D. Bohm. A suggested interpretation of the quantum theory in terms of “hidden” variables. *Physical Review*, 85(2):166–193, 1952.
- [15] B. E. Baaquie. *Quantum Finance: Path Integrals and Hamiltonians for Options and Interest Rates*. Cambridge University Press, 2004.
- [16] M. A. Andreev & J. Howden. Stochastic mechanical methods for quantitative portfolio management. *Quark Quantitative Research Working Paper*, 2026a.
- [17] M. A. Andreev & J. Howden. Ontological limits of historical validation for wavefunction-calibrated financial systems. *Quark Quantitative Research Working Paper*, 2026e.
-
- © 2026 Maksim Aleksandrovich Andreev & Joseph Howden.
All rights reserved.
 Quark Quantitative Research Platform — Working Paper
 Generated 2026-08-25 22:08 UTC
This paper claims disciplined, gated, honestly reported harvesting of the documented variance risk premium, and not a novel alpha source. The Phase-2 verdict table is withdrawn in full following the discovery of a stale-frame defect in the backtest engine on 2026-08-04, and the bull-call spread's disconfirmation is retracted with it. The surviving evidence is the twenty-year component characterisation of Section 6.1 and the subperiod persistence re-read of Section 6.2, each of which must be cited with the riders stated there.