

SYNCHRONIZED STRESS RELATIVE VALUE PUBLIC-NEUTRAL OPTION BOOKS UNDER LATENT STRESS COMPRESSION

J. VIDAL LLAURADÓ 

ABSTRACT. Public option calibration can fix a finite public span while leaving synchronized-stress exposure only partly spanned. The residual compression is the listed-option relative-value object. The trading coordinate is the wedge between index-implied synchronization and an adapted realized-synchronization benchmark generated by the released Volterra–Perron stress state. The sign condition

$$\rho_{d,I,\tau}^{\text{imp}} > \widehat{\rho}_{d,I,\tau,w}^{\text{real}},$$

selects direction, not size. Candidate rows are constant-unit listed-option books. Capital enters only through the lower-bound sizing equation. When the index surface embeds more synchronized stress than admissible single-name variance carriers support, the selected book is short index convexity and long carrier-basket convexity under public-neutrality, liquidity, cost, and capital constraints.

Five maps separate the signal, realized benchmark, carrier, book formation, and sizing. The implied-synchronization map extracts the coordinate from the index-versus-single-name variance decomposition. The realized-synchronization map supplies the released stress-state benchmark. The carrier map selects single names able to absorb index stress through stress co-movement, network pressure, relative variance value, and tradability. The book map selects an element of the finite listed-option span under public-neutrality constraints and a stress-convexity hedge ratio. The sellability map converts the estimated book edge into position size through a lower bound, a book-risk denominator, and mandate caps. The main statements identify finite-dimensional public-neutral optimization, characterize the stress-convexity hedge ratio as the conditional least-squares projection of index-stress P&L onto carrier-basket P&L, prove that the lower-bound gate certifies positive expected edge under conservative reserves, and separate adapted lower-bound sizing from future P&L realization.

The empirical implementation uses listed OptionMetrics contracts, validated underlying prices, and rolling Oxford–Man Perron stress states. Over the 2001–2018 execution ledger, the balanced lower-bound specification forms 1,496 sign-positive candidate books and allocates capital to 906 rows over 805 active dates. At 1x bid–ask costs it reports total execution P&L 24.567 million, annualized Sharpe 3.958, expected-shortfall loss 23,683, and max drawdown 364,148. A mechanical leakage audit, cost-stress grid, mandate frontier, dormancy audit, and capital-realism report accompany the run. The implementation is formation-adapted under those checks, but it was developed on the available sample. The ledger is therefore an audited sample-developed implementation rather than an unused-sample trading record.

1. INTRODUCTION

Listed index and single-name options supply a finite quote-date span. They do not by themselves determine a sellable exposure to latent synchronization. The synchronization wedge compares an index-implied synchronization coordinate with an adapted realized-synchronization benchmark. The coordinate is finite and mixed-measure; it serves as a relative-value signal rather than a residual-density selector or recovery theorem. When the implied coordinate exceeds the adapted benchmark, the index leg carries excess synchronized-stress convexity relative to the admissible carrier span. The selected book is short index convexity and long carrier convexity.

Date: May 2026.

2020 Mathematics Subject Classification. 60G22, 91G20, 91G80, 62M15, 62P05.

Key words and phrases. synchronization risk, option relative value, public-neutral portfolios, Perron stress, variance dispersion, incomplete markets, expected shortfall.

DOI: 10.2139/ssrn.6731781.

The wedge inherits its state variable from the market-scale Perron state generated by directed Volterra contagion [1]. Risk-neutral compression and residual-pricing geometry supply the public quotient and residual pricing fiber of that state [2, 3]. The latent-contagion, dynamic-observability, and projected-score results provide the physical source-screened input [4, 5, 6]. The finite relative-value book is the listed-option representative of the resulting synchronization residual.

A synchronization book is admissible only if the carrier side absorbs the prescribed stress direction. Relative variance value alone does not turn a basket of single-name straddles into a synchronization carrier. The basket must load on the index stress state, lie near the network stress-transmission direction, carry admissible variance value, and remain in the executable listed-option span. The governing object is therefore a book-formation map

(public surface, latent stress state, listed option span) $\mapsto$ admissible synchronization book.

The map produces a public-neutral book with controlled synchronization-residual exposure and positive lower-bound sellability, leaving completion of the latent market outside its scope.

The execution problem is finite-dimensional. Listed options supply a finite span on each quote date. Delta, vega, gamma, maturity, public-surface principal components, liquidity, bid-ask costs, and concentration constraints are finite-dimensional controls. The latent stress geometry enters through the synchronization wedge, carrier map, danger state, and residual-capital adjustment. This separates public calibration from listed-option hedgeability. The Hilbert-space compression and residual-hedging results in [2, 3] identify the residual object; the empirical implementation selects a finite listed-contract representative.

The cited literatures enter through different objects. Option pricing and martingale asset-pricing theory supply the listed payoff and no-arbitrage background [7, 8, 9]. Evidence on option returns, jump premia, and variance-risk premia supplies the empirical-premium context [10, 11, 12, 13]. Volatility-surface dynamics, unspanned volatility, and rough volatility locate the public-surface and span boundaries [14, 15, 16]. Incomplete-market and mean-variance hedging supply the residual-hedge language [17, 18]. Good-deal restrictions and expected-shortfall control enter only at the capital boundary [19, 20].

The admitted specification sells excess synchronization only when the carrier side can absorb realized stress under cost and capital constraints. The formal sections define the implied synchronization coordinate, the carrier absorption map, the public-neutral book constraint, the stress-convexity hedge ratio, the sellability lower bound, the lower-bound sizing projection, and the capital-risk ledger.

The empirical section implements the sellability map rather than testing full latent recovery. Candidate books are formed from the sign of the synchronization wedge. Each source row is a constant-unit listed-option book, and capital enters only through the lower-bound sizing equation. The reserve calculation uses entry-known spread, carrier carry, directed-support defects, danger-state reserves, model-risk reserves, and a pre-formation capital-at-risk denominator. The empirical object is a book-level sellability test. A synchronization wedge is tradable only after observable cost, carrier, danger, and capital-risk charges are applied before realized P&L is observed.

Contributions. The main object is the lower-bound synchronization book. Six components define it.

- (1) a synchronization wedge that converts implied-versus-realized correlation information into a variance-dollar relative-value coordinate;
- (2) a carrier absorption map that selects the single-name convexity used to protect a short index-synchronization book;
- (3) a finite public-neutral option-book optimization with stress-exposure constraints;
- (4) a stress-convexity hedge ratio that replaces ordinary vega matching by conditional stress-response matching;
- (5) a sellability lower bound and lower-bound sizing theorem that separates raw premia from tradable premia; and

- (6) an adapted empirical ledger with listed-option execution, adaptedness audit, cost-stress results, mandate frontier, dormancy audit, and capital-realism accounting.

Organization. Section 2 fixes the filtered listed-option market and the public and latent information fields. Section 3 defines the synchronization coordinate and the variance-dollar wedge. Section 4 constructs carrier weights. Section 5 defines public-neutral books and proves the finite-dimensional projection and optimization statements. Section 6 introduces the lower-bound gate and sizing map. Section 7 records the capital and expected-shortfall ledger. Section 8 reports the listed-option implementation and empirical audit checks. Section 9 fixes claim boundaries and finite-span extensions. Section 10 returns to the residual sellability boundary. Appendix A collects the data-readiness and leakage-audit tables supporting the empirical run.

2. LISTED-OPTION MARKET

Trading decisions are made on a discrete quote calendar $d = 0, 1, \dots, T$. The public information available before trading on date d is $\mathcal{H}_d$. It contains current and past listed option quotes, screened public surface summaries, validated underlying prices, and previously released state estimates. The latent stress information field is $\mathcal{K}_d$, with $\mathcal{H}_d \subseteq \mathcal{K}_d$. The latent Perron stress state is a scalar or vector

$$P_{d,\text{lat}}^{(w)} \in L^2(\mathbb{P}, \mathcal{K}_d),$$

constructed on rolling windows $w \in \{63, 126, 252, 504\}$. The trading rule uses delayed, latest-available released state estimates

$$\hat{P}_d^{(w)} \in L^2(\mathbb{P}, \mathcal{H}_d).$$

Every stress-state input used at the date- d signal-formation time is therefore public at that formation time. The released coordinate may estimate a latent $\mathcal{K}_d$ -measurable object, but the execution rule acts only on the $\mathcal{H}_d$ -measurable estimate.

Let $\mathcal{U}_d$ be the finite set of eligible listed option units on date d . A unit $u \in \mathcal{U}_d$ is a contract or a small contract package, such as an index straddle or a single-name straddle in a fixed maturity bucket. Its next-horizon execution P&L is denoted

$$R_{d,h}(u),$$

where h is the holding horizon. The reported ledger uses weekly holding horizons; one-day and 21-day clocks serve as sensitivity checks.

After the contract basis e_u is fixed, the executable listed-option span is

$$\mathcal{S}_d := \text{span}\{e_u : u \in \mathcal{U}_d\} \cong \mathbb{R}^{n_d}, \quad n_d = |\mathcal{U}_d|.$$

A book is a coordinate vector $x \in \mathbb{R}^{n_d}$ in this span. Its P&L is

$$R_{d,h}(x) = \sum_{u \in \mathcal{U}_d} x_u R_{d,h}(u).$$

The public risk map

$$A_d : \mathbb{R}^{n_d} \rightarrow \mathbb{R}^{q_d}$$

collects the public exposures constrained by the book, such as delta, vega, selected gamma/theta proxies, maturity exposure, index/single-name gross exposure, and public option-surface principal components. The vector $A_d x$ is computed from $\mathcal{H}_d$ -measurable quantities.

Definition 2.1 (Public-neutral listed-option book). Given a tolerance vector $\epsilon_d \geq 0$, a listed-option book $x \in \mathbb{R}^{n_d}$ is ϵ_d -public-neutral if

$$|A_d x| \leq \epsilon_d$$

coordinatewise. It is exactly public-neutral if $\epsilon_d = 0$.

Remark 2.2. Public neutrality is a public-span condition, not absence of latent exposure. The intended book is neutral to specified public surface risks while retaining controlled exposure to the synchronization residual. This is the trading version of the compression and residual-hedging distinction in [2, 3].

3. SYNCHRONIZATION WEDGE

The relative-value coordinate must be measured on the scale of the book. Correlation units determine sign; variance and premium dollars determine execution. The map therefore defines an implied synchronization coordinate and rescales it by the index-versus-single-name synchronization denominator.

Fix an index I , a maturity bucket τ , and a state window w . Let

$$c = (I, \tau, w)$$

denote the cell. Let N_I be the admitted constituent universe in the empirical implementation, and let $a_{d,i}^I$ be the date- d index weight assigned to name i . Let $\sigma_{d,I,\tau}^{2,\text{imp}}$ be the implied total variance of the index, and let $\sigma_{d,i,\tau}^{2,\text{imp}}$ be the implied total variance of name i . The index variance decomposition has the schematic form

$$\sigma_{d,I,\tau}^{2,\text{imp}} = \sum_{i \in N_I} (a_{d,i}^I)^2 \sigma_{d,i,\tau}^{2,\text{imp}} + \rho_{d,c}^{\text{imp}} \mathcal{D}_{d,c}^{\text{imp}},$$

where the synchronization denominator is

$$\mathcal{D}_{d,c}^{\text{imp}} = 2 \sum_{i < j} a_{d,i}^I a_{d,j}^I \sigma_{d,i,\tau}^{\text{imp}} \sigma_{d,j,\tau}^{\text{imp}}.$$

Only cells with $\mathcal{D}_{d,c}^{\text{imp}} > 0$ carry a nondegenerate synchronization denominator. On those cells, the implied synchronization coordinate is

$$\rho_{d,c}^{\text{imp}} = \frac{\sigma_{d,I,\tau}^{2,\text{imp}} - \sum_{i \in N_I} (a_{d,i}^I)^2 \sigma_{d,i,\tau}^{2,\text{imp}}}{\mathcal{D}_{d,c}^{\text{imp}}}.$$

Cells with $\mathcal{D}_{d,c}^{\text{imp}} \leq 0$ are outside this synchronization coordinate unless a predeclared fallback convention is supplied.

The realized synchronization benchmark is an adapted estimate

$$\hat{\rho}_{d,c}^{\text{real}} = \hat{\rho}^{\text{real}}(\hat{P}_d^{(w)}, \text{released directed state}_d, \text{realized covariance history}_d),$$

formed only from information available at signal date d . The raw synchronization wedge is

$$\pi_{d,c}^{\text{sync}} = \rho_{d,c}^{\text{imp}} - \hat{\rho}_{d,c}^{\text{real}}.$$

Positive values indicate that the option surface prices a larger synchronization component than the date- d state-based realized benchmark. The sign gate is fixed at formation.

The wedge is a mixed-measure trading coordinate. The coordinate $\rho_{d,c}^{\text{imp}}$ is public option-implied, while $\hat{\rho}_{d,c}^{\text{real}}$ is a $\mathbb{P}$ -adapted realized-synchronization benchmark. The difference is an implied-versus-forecast-realized wedge, not by itself a residual-density selector or conditional synchronization premium.

Definition 3.1 (Variance-dollar synchronization edge). The variance-dollar synchronization edge in cell c is

$$\Delta V_{d,c}^{\text{sync}} = \pi_{d,c}^{\text{sync}} \mathcal{D}_{d,c}^{\text{imp}}.$$

Let $N_{d,c}^{\text{sync}} > 0$ be the $\mathcal{H}_d$ -measurable conversion from this variance coordinate to the one-unit listed-option book scale. It includes the variance notional, contract multiplier, and book-unit convention used by the implementation. The cost-adjusted book edge is

$$\hat{\mu}_{d,c}^{\text{book}} = N_{d,c}^{\text{sync}} \Delta V_{d,c}^{\text{sync}} - \text{TC}_{d,c} - \text{Carry}_{d,c} - \text{HE}_{d,c} - \text{Tail}_{d,c},$$

where the four subtractions are, respectively, transaction cost, single-name carry drag, hedge error reserve, and tail-state reserve, all expressed on the same one-unit book scale.

Remark 3.2. The tradable object is book edge, not raw correlation. A large $\rho^{\text{imp}} - \hat{\rho}^{\text{real}}$ is insufficient when the denominator is small, the carrier is expensive, or the tail reserve is large. Identifying this finite listed-option wedge with a dynamic residual premium would require additional projection, basis, and residual selection assumptions.

4. CARRIER ABSORPTION

The long side of a short-synchronization book carries realized index stress when the short index leg is exposed. The carrier map selects absorbent convexity, not names by heuristic rank. It assigns long convexity to names that co-move with index stress, sit near the network stress direction, carry admissible variance value, and can be traded at acceptable cost.

For $i \in N_I$, define the carrier score

$$s_{d,c,i} = z_d \left(\beta_{d,c,i}^{\text{stress}} \right) + \lambda_1 z_d \left(n_{d,w,i}^{\text{net}} \right) + \lambda_2 z_d \left(v_{d,\tau,i}^{\text{value}} \right) - \lambda_3 z_d \left(\ell_{d,\tau,i}^{\text{cost}} \right),$$

where z_d denotes an adapted cross-sectional normalization. The terms are

$$\beta_{d,c,i}^{\text{stress}} = \text{train-window response of name } i \text{'s convexity to index stress,}$$

$$n_{d,w,i}^{\text{net}} = \text{network carrier pressure from the rolling latent-state construction,}$$

$$v_{d,\tau,i}^{\text{value}} = \widehat{\text{RV}}_{d,\tau,i} - \text{IVar}_{d,\tau,i}^{\text{imp}},$$

and $\ell_{d,\tau,i}^{\text{cost}}$ is a liquidity and bid-ask penalty. It is taken to be nonnegative before the adapted cross-sectional normalization is applied.

Raw carrier scores can collapse into a few names. The shrinkage rule is

$$\tilde{s}_{d,c,i} = \omega_{d,c,i} s_{d,c,i} + (1 - \omega_{d,c,i}) \left[\alpha \bar{s}_{d,c,\text{sector}(i)} + (1 - \alpha) s_i^{\text{prior}} \right],$$

where s_i^{prior} may encode a persistent technology or sector carrier prior, and $\omega_{d,c,i} \in [0, 1]$ is an adapted reliability weight.

Definition 4.1 (Carrier basket). Let $p, q > 0$, $\epsilon > 0$, and $\ell_{d,\tau,i}^{\text{cost}} \geq 0$. Assume

$$Z_{d,c}^{\text{car}} := \sum_{j \in N_I} \frac{(\tilde{s}_{d,c,j}^+)^p}{(\ell_{d,\tau,j}^{\text{cost}} + \epsilon)^q} > 0.$$

If this denominator is zero, the cell is carrier-inadmissible unless a predeclared diversified fallback basket is supplied. Otherwise the unconstrained carrier weight is

$$\bar{b}_{d,c,i} = \frac{(\tilde{s}_{d,c,i}^+)^p / (\ell_{d,\tau,i}^{\text{cost}} + \epsilon)^q}{Z_{d,c}^{\text{car}}}.$$

The carrier basket $b_{d,c}$ is the projection of $\bar{b}_{d,c}$ onto the capped simplex

$$\Delta_{d,c} = \left\{ b \in \mathbb{R}_+^{N_I} : \sum_i b_i = 1, \quad b_i \leq b_{\max}, \quad H(b) \geq H_{\min} \right\},$$

where $H(b) = -\sum_i b_i \log b_i$ is entropy, with the convention $0 \log 0 := 0$.

The carrier projection is a compact finite projection problem. The proposition identifies the existence boundary and the convexity condition under which the projection is single-valued.

Proposition 4.2 (Existence of carrier weights). *If $\Delta_{d,c}$ is nonempty, then a carrier basket $b_{d,c}$ exists. If the projection is taken in squared Euclidean distance, the basket is unique whenever $\Delta_{d,c}$ is convex.*

Proof. The feasible carrier set is a subset of the finite simplex. The nonnegativity, unit-sum, and cap constraints are closed and bounded. Entropy is continuous on the simplex, with the convention $0 \log 0 := 0$, so the constraint $H(b) \geq H_{\min}$ is closed. Hence $\Delta_{d,c}$ is compact whenever it is nonempty. The squared Euclidean distance from $\bar{b}_{d,c}$ is continuous on $\Delta_{d,c}$, and the Weierstrass theorem gives existence of a minimizer.

For uniqueness, the nonnegative capped simplex is convex. Since entropy is concave, its superlevel set is convex; therefore the displayed $\Delta_{d,c}$ is convex. The map $b \mapsto \|b - \bar{b}_{d,c}\|_2^2$ is strictly convex on the ambient Euclidean space and hence on any affine feasible slice. A strictly convex objective has at most one minimizer on a convex set. $\square$

5. PUBLIC-NEUTRAL SYNCHRONIZATION BOOKS

Embedding the carrier basket in the listed-option span creates the executable object. The primitive book has a short index-convexity leg and a long carrier-convexity leg. The optimization then imposes public-risk constraints so that the remaining exposure is the synchronization residual, not an ordinary surface factor.

Let $u_I(d, \tau)$ be the selected index convexity unit, for example an index straddle in the τ -day bucket. Let $u_i(d, \tau)$ be the corresponding single-name convexity unit. The basic short-synchronization book is

$$x_{d,c} = -q_I e_I + q_B \sum_{i \in N_I} b_{d,c,i} e_i.$$

The sign convention is that $q_I > 0$ sells index convexity and $q_B > 0$ buys carrier-basket convexity.

5.1. Stress-convexity hedge ratio. The hedge target is stress response, not vega matching. The book is a synchronization trade, and the long basket is selected to absorb realized index stress. Let

$$R_{d,h}^I = R_{d,h}(e_I), \quad R_{d,h}^B = R_{d,h} \left(\sum_i b_{d,c,i} e_i \right).$$

On a train-window stress set $\mathcal{E}_{d,c}^{\text{stress}}$, write $\mathbb{E}_S[\cdot] = \mathbb{E}_{\text{train}}[\cdot \mid \mathcal{E}_{d,c}^{\text{stress}}]$. The conditioning set is used only when $\mathbb{P}_{\text{train}}(\mathcal{E}_{d,c}^{\text{stress}}) > 0$. Define

$$\chi_{d,c} = \frac{\text{Cov}_{\text{train}}(R_{d,h}^I, R_{d,h}^B \mid \mathcal{E}_{d,c}^{\text{stress}})}{\text{Var}_{\text{train}}(R_{d,h}^B \mid \mathcal{E}_{d,c}^{\text{stress}}) + \epsilon}.$$

The book uses $q_B = \eta_{d,c} \chi_{d,c}^+ q_I$, with $0 \leq \eta_{d,c} \leq 1$, and then imposes aggregate gamma and vega bands rather than exact equality. Cells with $\chi_{d,c} \leq 0$ do not have a positive stress-convexity carrier under this rule.

The coefficient is the stress-conditioned projection of index P&L onto carrier-basket P&L. The proposition records the least-squares characterization.

Proposition 5.1 (Stress-convexity hedge ratio). *If $\epsilon = 0$ and $\text{Var}_{\text{train}}(R_{d,h}^B \mid \mathcal{E}_{d,c}^{\text{stress}}) > 0$, then $\chi_{d,c}$ is the slope coefficient in the conditional least-squares problem*

$$\min_{a, \chi \in \mathbb{R}} \mathbb{E}_S \left[\left(R_{d,h}^I - a - \chi R_{d,h}^B \right)^2 \right].$$

Equivalently, it is the no-intercept least-squares coefficient for the centered variables $R_{d,h}^I - \mathbb{E}_S[R_{d,h}^I]$ and $R_{d,h}^B - \mathbb{E}_S[R_{d,h}^B]$.

Proof. All moments are taken under the train-window conditional law on $\mathcal{E}_{d,c}^{\text{stress}}$. Let

$$m_I = \mathbb{E}_S[R_{d,h}^I], \quad m_B = \mathbb{E}_S[R_{d,h}^B].$$

The conditional loss is a finite quadratic function of (a, χ) . Differentiation in a gives

$$a = m_I - \chi m_B.$$

After substitution the criterion becomes

$$\mathbb{E}_S \left[\left((R_{d,h}^I - m_I) - \chi (R_{d,h}^B - m_B) \right)^2 \right].$$

The first-order condition in χ is

$$-2 \text{Cov}_{\text{train}}(R_{d,h}^I, R_{d,h}^B \mid \mathcal{E}_{d,c}^{\text{stress}}) + 2\chi \text{Var}_{\text{train}}(R_{d,h}^B \mid \mathcal{E}_{d,c}^{\text{stress}}) = 0.$$

The variance assumption removes the null carrier direction and makes the centered quadratic strictly convex. The displayed coefficient is therefore the unique least-squares slope. $\square$

5.2. Finite-dimensional public-neutral optimization. Let $\hat{\mu}_d \in \mathbb{R}^{n_d}$ be the train-window expected P&L estimate, let $\hat{\Sigma}_d$ be a positive semidefinite risk estimate, and let $\text{TC}_d(x)$ be a convex transaction-cost penalty. The book optimization is

$$\max_{x \in \mathbb{R}^{n_d}} \hat{\mu}_d^\top x - \lambda x^\top \hat{\Sigma}_d x - \kappa \text{TC}_d(x)$$

subject to public-neutrality, gross and net exposure caps, liquidity caps, carrier exposure constraints, and maturity, ticker, and sector concentration caps.

The theorem states the public-neutral execution problem as a finite constrained optimization. Existence is compactness. Uniqueness is strict concavity on the feasible direction span.

Theorem 5.2 (Finite public-neutral book optimization). *Assume the feasible set $\mathcal{X}_d \subset \mathbb{R}^{n_d}$ is nonempty, closed, convex, and bounded, and assume TC_d is convex and lower semicontinuous, with $\lambda \geq 0$ and $\kappa \geq 0$. Then the public-neutral book problem has a solution. If the objective is strictly concave on $\mathcal{X}_d$, for example if $\lambda > 0$ and*

$$v^\top \hat{\Sigma}_d v > 0$$

for every nonzero feasible direction $v = x - y$ with $x, y \in \mathcal{X}_d$, the solution is unique.

Proof. The feasible set is compact by the Heine–Borel theorem. The linear term and the quadratic risk term are continuous. Since TC_d is lower semicontinuous and $\kappa \geq 0$, the term $-\kappa \text{TC}_d$ is upper semicontinuous. The full objective is therefore upper semicontinuous on a compact set, and the Weierstrass theorem gives a maximizer.

For uniqueness, let $x \neq y$ in $\mathcal{X}_d$, set $v = x - y$, and define $z_\theta = \theta x + (1 - \theta)y$ for $0 < \theta < 1$. The identity

$$z_\theta^\top \hat{\Sigma}_d z_\theta = \theta x^\top \hat{\Sigma}_d x + (1 - \theta)y^\top \hat{\Sigma}_d y - \theta(1 - \theta)v^\top \hat{\Sigma}_d v$$

shows that the quadratic penalty is strictly concave along every nonconstant feasible segment when the displayed direction condition holds and $\lambda > 0$. The cost term is concave, and the return term is affine. Hence the full objective is strictly concave on $\mathcal{X}_d$. Two distinct maximizers would make their midpoint feasible and strictly better, which is impossible. $\square$

Remark 5.3. The theorem is a finite listed-option statement. Infinite-dimensional compression enters only through the priced stress loading and the residual object. The executed book is not assumed to span the full latent market.

6. SELLABILITY AND LOWER-BOUND SIZING

A synchronization wedge is not automatically sellable. The lower-bound map separates a raw coordinate wedge from a wedge that can be traded after costs, model uncertainty, carrier quality, and tail reserves. Let $D_{d,c}$ be an adapted danger score, for example a normalized combination of Perron level, Perron slope, aggregate realized-volatility acceleration, directed-contagion breadth, realized synchronization acceleration, and index gap pressure. The tail-state reserve in $\hat{\mu}_{d,c}^{\text{book}}$ is part of the one-unit expected book-edge estimate. The lower-bound reserve below is an additional conservative haircut applied when that edge is converted into an admissible capital decision.

Definition 6.1 (Sellability lower bound). For cell c , define

$$LB_{d,c} = \hat{\mu}_{d,c}^{\text{book}} - z_\alpha \hat{\sigma}_{d,c}^{\text{model}} - \text{TC}_{d,c}^{\text{round}} - \text{Tail}_{d,c}^{\text{reserve}}.$$

The cell is lower-bound admissible if $LB_{d,c} > 0$.

The lower bound is a sign certificate, not a complete pricing statement. The proposition states the conditional expected-edge claim certified by the reserve inequality.

Proposition 6.2 (Lower-bound positive-edge certificate). *Let $\mathbb{E}_d[\cdot] := \mathbb{E}[\cdot \mid \mathcal{H}_d]$. Suppose the one-period book $P\mathcal{E}L Y_{d,c}$ satisfies the adapted conservative forecast inequality*

$$\mathbb{E}_d[Y_{d,c}] \geq \hat{\mu}_{d,c}^{\text{book}} - z_\alpha \hat{\sigma}_{d,c}^{\text{model}} - \text{TC}_{d,c}^{\text{round}} - \text{Tail}_{d,c}^{\text{reserve}}.$$

If $LB_{d,c} > 0$, then $\mathbb{E}_d[Y_{d,c}] > 0$.

Proof. By definition of $LB_{d,c}$, the right-hand side of the forecast inequality is exactly $LB_{d,c}$. Hence $\mathbb{E}_d[Y_{d,c}] \geq LB_{d,c}$. If $LB_{d,c} > 0$, then $\mathbb{E}_d[Y_{d,c}] > 0$. $\square$

Definition 6.3 (Book-risk denominator). The book-risk denominator $\mathcal{R}_{d,c}$ is an $\mathcal{H}_d$ -measurable positive reserve for one unit of the candidate book. Each input is converted to the same one-unit dollar loss or reserve scale before the maximum is taken. In the reported implementation it is the maximum of a settled pre-formation model-risk reserve, a settled tail-loss reserve, the entry-date spread proxy, a unit floor, and the formation-date margin proxy:

$$\mathcal{R}_{d,c} = \max\{\widehat{\sigma}_{d,c}^{\text{hist}}, \widehat{\text{ES}}_{d,c}^{\text{hist}}, \text{TC}_{d,c}^{\text{entry}}, 1, \text{Margin}_{d,c}^{\text{proxy}}\}.$$

The lower-bound sizing rule uses this one-unit dollar reserve as the quadratic penalty coefficient for capital allocation.

Definition 6.4 (Lower-bound size). Fix a mandate risk-aversion parameter $\lambda > 0$. The unconstrained lower-bound size is

$$x_{d,c}^+ = \frac{(LB_{d,c})_+}{\lambda \mathcal{R}_{d,c}}.$$

If $\bar{x}_{d,c}^{\text{margin}}$, $\bar{x}_{d,c}^{\text{tail}}$, and $\bar{x}_{d,c}^{\text{cost}}$ are the nonnegative $\mathcal{H}_d$ -measurable mandate caps generated by margin, tail-risk, and cost budgets, the executable size is

$$x_{d,c}^{\text{exe}} = \min\left\{x_{d,c}^+, \bar{x}_{d,c}^{\text{margin}}, \bar{x}_{d,c}^{\text{tail}}, \bar{x}_{d,c}^{\text{cost}}\right\}.$$

If $LB_{d,c} \leq 0$, the executable size is zero.

The sizing map projects a one-dimensional capital objective onto the nonnegative half-line and the mandate caps. The proposition identifies the clipped optimizer.

Proposition 6.5 (Lower-bound size as a capital decision). *For fixed λ and $\mathcal{R}_{d,c} > 0$, the unconstrained lower-bound size maximizes the one-dimensional quadratic capital objective*

$$x LB_{d,c} - \frac{\lambda}{2} \mathcal{R}_{d,c} x^2 \quad \text{over } x \geq 0.$$

Applying margin, tail, and cost caps gives the constrained executable size.

Proof. Let

$$f(x) = x LB_{d,c} - \frac{\lambda}{2} \mathcal{R}_{d,c} x^2, \quad x \geq 0.$$

Because $\lambda > 0$ and $\mathcal{R}_{d,c} > 0$, the function is strictly concave. Its first-order condition is

$$LB_{d,c} - \lambda \mathcal{R}_{d,c} x = 0.$$

If $LB_{d,c} > 0$, the unconstrained nonnegative maximizer is $LB_{d,c}/(\lambda \mathcal{R}_{d,c})$; if $LB_{d,c} \leq 0$, strict concavity and $f'(0) = LB_{d,c} \leq 0$ put the maximizer at zero. Thus the unconstrained solution is $(LB_{d,c})_+ / (\lambda \mathcal{R}_{d,c})$.

The mandate caps reduce the feasible set to the interval

$$0 \leq x \leq \min\{\bar{x}_{d,c}^{\text{margin}}, \bar{x}_{d,c}^{\text{tail}}, \bar{x}_{d,c}^{\text{cost}}\}.$$

A one-dimensional strictly concave quadratic increases up to its unconstrained maximizer and decreases after it. The constrained maximizer is therefore the displayed minimum of the unconstrained size and the three caps. $\square$

Adaptedness is a measurability statement. The theorem places the executable vector in the formation-date information field when the candidate book and sizing inputs are formation-date measurable.

Theorem 6.6 (No-look-ahead lower-bound sizing). *Let the candidate book, lower bound, book-risk denominator, and mandate caps be $\mathcal{H}_d$ -measurable. Then the executable lower-bound sized book is $\mathcal{H}_d$ -measurable. In particular, changing future P&L, exit prices, future hedging gains, realized exit spreads, or post-trade costs cannot change the position formed at date d .*

Proof. The positive-part map, the reciprocal map on $(0, \infty)$, scalar multiplication by $1/\lambda$, and finite minima are Borel measurable. Hence the map

$$(LB, \mathcal{R}, \text{caps}) \mapsto \min \left\{ \frac{(LB)_+}{\lambda \mathcal{R}}, \bar{x}^{\text{margin}}, \bar{x}^{\text{tail}}, \bar{x}^{\text{cost}} \right\}$$

is Borel measurable on the domain $\mathcal{R} > 0$. Composition with $\mathcal{H}_d$ -measurable inputs gives an $\mathcal{H}_d$ -measurable scalar size. The candidate book is $\mathcal{H}_d$ -measurable by assumption, and scalar multiplication is a Borel map on the finite listed-option span. The executable book is therefore $\mathcal{H}_d$ -measurable. Future P&L and future cost realizations are not arguments of the sizing map, so perturbing them cannot change the date- d position. $\square$

7. CAPITAL-RISK LEDGER

The capital ledger maps the mathematical book into quantities observable in the listed-option execution record. The reported coordinates are execution P&L, transaction costs, drawdown, expected shortfall, gross premium, and margin proxies. Let Y_d be daily net execution P&L and let

$$L_d = (-Y_d)^+$$

be the daily loss. The empirical expected-shortfall loss is

$$\widehat{\text{ES}}_{95} = \frac{1}{|\{d : L_d \geq q_{0.95}(L)\}|} \sum_{d: L_d \geq q_{0.95}(L)} L_d.$$

This is the implemented tie convention; if many observations equal the empirical quantile, the average may include slightly more than the upper five percent of losses. The max drawdown is computed from the cumulative P&L process. A drawdown budget $B > 0$ is breached if

$$\max_t \left(\max_{s \leq t} \sum_{u \leq s} Y_u - \sum_{u \leq t} Y_u \right) > B.$$

The ledger proposition separates datewise proxy compliance from ex-post drawdown breach. The residual breach is the positive part of the excess drawdown.

Proposition 7.1 (Capital-budget admissibility). *If the daily book process satisfies the gross-premium, margin, concentration, and drawdown-budget constraints date by date, then the realized capital ledger is admissible under those proxy constraints. If the drawdown budget is breached, the breach amount is exactly*

$$\max\{\text{MDD} - B, 0\}.$$

Proof. The ledger is a finite collection of proxy coordinates evaluated on the realized book process. If the gross-premium, margin, concentration, and drawdown-budget inequalities hold at each admissible date, the realized ledger remains in the corresponding proxy-admissible set. Let MDD denote the displayed maximum drawdown. The drawdown budget is breached exactly on $\{\text{MDD} > B\}$. The excess is $\text{MDD} - B$ on that event and zero otherwise, hence $\max\{\text{MDD} - B, 0\}$. $\square$

Remark 7.2. The capital account remains a proxy. It is closer to deployment than normalized P&L because it uses execution dollars, spread costs, gross premium, and margin proxies. Broker margin, portfolio margin, financing, borrow, exercise and assignment handling, and hard capacity constraints remain outside the proxy.

8. LISTED-OPTION IMPLEMENTATION

The empirical implementation is a listed-option ledger built from OptionMetrics contracts [21], validated underlying prices, and rolling Oxford–Man Perron stress states [22]. Signals are formed from information available at date d , entered on the next available quote date, and held under predeclared clocks. The executed specification is short overpriced synchronization:

$$\rho_{d,c}^{\text{imp}} > \hat{\rho}_{d,c}^{\text{real}} \Rightarrow \text{short index convexity plus long carrier-basket convexity,}$$

subject to the public-neutrality, carrier, cost, and lower-bound constraints above.

8.1. Data and admissible universe. The empirical universe is US listed options and US-traded equity underliers. The option component uses OptionMetrics contract data for index, liquid ETF, sector ETF, and admitted single-name candidates; the state component uses Oxford–Man realized-volatility indices to build rolling physical stress states. The restriction is empirical, not theoretical. The definitions above require an index I , an admissible carrier universe, a finite listed-option span generated by $\mathcal{U}_d$, and an adapted stress state. The same map can be applied to another option market once those objects, quote calendars, execution costs, and state inputs are rebuilt.

TABLE 1. Data inputs for the empirical implementation. The execution run is US-option based; the mathematical map requires an admissible option span, carrier universe, and adapted stress state.

Component	Role in the implementation	Main implementation choices
OptionMetrics listed contracts	Tradeable option span, bid–ask execution, Greeks, implied variance, and straddle/strangle payoffs	US index, ETF, sector ETF, and single-name candidates; 30, 60, and 90 day maturity buckets with liquidity, spread, quote, and maturity filters
OptionMetrics forwards and security prices	Forward alignment, underlying close series, and delta-hedge accounting	Validated ticker-date panel; no delta-hedged execution result is reported without matched underlying prices
Oxford–Man realized-volatility indices	Physical stress state, realized synchronization benchmark, and aggregate danger controls	Rolling windows 63, 126, 252, 504, with 126 and 252 as the primary trading windows
Directed-contagion state suite	Carrier pressure, directed confirmation, and danger-state diagnostics	Rolling-only state build with one-day availability and no use of future option outcomes

The data assembly keeps raw contract data, forward alignment, surface-derived coordinates, and realized-state inputs separate until the trade ledger is formed. This separation prevents the exercise from becoming an option-surface fit. The object is a listed-option execution test of whether a synchronization wedge remains sellable after carrier, cost, danger-state, and capital-risk charges. Appendix A reports the mechanical row counts, leakage checks, mandate frontier, cost stress, and period audit for the frozen run.

The reported specification is the lower-bound specification defined in Section 6. Candidate books are admitted by the sign-only condition

$$\rho_{d,I,\tau}^{\text{imp}} > \hat{\rho}_{d,I,\tau,w}^{\text{real}}.$$

The source rows are constant-unit listed-option books. The procedure computes the formation-date synchronization edge, entry-known spread proxy, carrier carry reserve, carrier support reserve, danger reserve, model-risk reserve, and capital-at-risk denominator. A row is used only when the resulting lower bound is positive, and its size is the lower-bound size after mandate margin, tail, and cost caps.

Four diagnostics bound the evidence. The dollar execution path records whether the synchronization book survives listed-option execution and bid–ask costs. The cost-stress grid tests sensitivity to spread assumptions. The mandate frontier locates the selected margin cap relative to nearby mandate choices. The adaptedness audit verifies that weights, lower bounds, and entry-cost proxies are invariant to perturbations of unsettled formation-date and future realized outcomes.

Remark 8.1 (Dormant sellability window). The lower-bound specification is inactive over part of the mid-2010s. The latent stress state need not disappear in that window. The synchronization wedge can remain present while the listed-option execution space fails to produce a clean

sellability certificate after liquidity, spread, carrier quality, and capital-risk constraints are imposed. Raw-exit diagnostics are consistent with this boundary. Relaxing the execution gate recovers more trades, but the additional books are expensive and drawdown-heavy. The inactivity is therefore a no-trade region of the public listed-option span.

The market setting of the period gives compatible but non-identifying context. The 2014–2016 window was a period of compressed US equity volatility, accommodative monetary conditions, calmer post-crisis risk appetite, and relatively benign realized index stress, although localized stress episodes still occurred. The claim is not causal. The narrower point is that such an environment can make a positive synchronization wedge less sellable after spread, carrier-quality, and capital-risk charges. The absence of trades is then part of the execution certificate rather than evidence that the latent geometry has vanished.

Figure 1 gives the finite ledger diagnostic. It combines the return path, drawdown path, annual execution P&L, execution intensity, and active capital exposure. The inactive mid-sample years are visible in the activity and exposure panels. They are a no-trade region of the sellability map, not an omitted-data patch. The rule stands down when the public listed-option span does not produce an acceptable lower-bound certificate.

TABLE 2. Lower-bound execution at 1x bid–ask costs. P&L is delta-hedged listed-option execution P&L in dollars. Candidate books come from the sign-positive synchronization wedge; capital is assigned by the lower-bound sizing equation with the formation-date capital-risk denominator.

Specification	P&L	Ann. Sh.	ES95	Max DD	Active	Used
Lower-bound	24.567mn	3.958	23,683	364,148	805	906




FIGURE 1. Execution diagnostics at 1x bid–ask costs. P&L is delta-hedged listed-option execution P&L in dollars. The panels show cumulative P&L, drawdown, annual P&L, trading activity, and active capital exposure for the lower-bound rule.

8.2. Adaptedness and uncertainty. Operational adaptedness is a filtration statement. At trade formation, the row size uses only the formation-date synchronization wedge, option and spread information, carrier geometry, danger variables, mandate caps, and settled pre-formation

TABLE 3. Cost-stress grid for the lower-bound rule. The same adapted formation weights are evaluated under higher bid–ask assumptions. At the 3x bid–ask stress, reported P&L remains positive and annualized Sharpe exceeds 2; drawdown increases with spread drag.

Cost multiplier	P&L	Ann. Sharpe	ES95 loss	Max drawdown
0.5	28.289mn	4.187	20,197	301,264
1.0	24.567mn	3.958	23,683	364,148
2.0	17.122mn	3.259	31,974	501,085
3.0	9.677mn	2.107	43,306	643,465

TABLE 4. Capital-realism readout at 1x costs. Margin quantities are implementation proxies rather than broker or exchange margin.

Quantity	Value
Total P&L	24.567mn
Annualized P&L	1.414mn
Total spread cost	7.445mn
Max margin proxy	3.493mn
Mean active margin proxy	0.468mn
Annualized P&L / mean active margin	3.02

history. It does not use future option P&L, future delta-hedge P&L, future returns, exit prices, realized exit-date spreads, future surface moves, or post-trade realized costs. The final audit perturbs unsettled formation-date and future realized P&L and realized spread columns across 96 formation dates. The maximum change in row weights, lower bounds, and entry-cost proxies is zero.

The reported Sharpe is bounded by two restrictions. First, the run is mechanically adapted under the verification tables in Appendix A. Second, the architecture was developed on the available sample. The run is therefore an audited research implementation rather than a fresh unused-sample record. With no unused data available, the evidentiary basis is the sign-only source, lower-bound sizing, leakage invariance, cost sensitivity, mandate frontier, dormancy audit, capital realism, and the explicit separation between discovery and execution.

9. LIMITATIONS

The implementation is recorded mathematically with its boundary conditions fixed. The empirical universe is US-centric by design, covering SPX/NDX, US index and sector ETFs, and admitted US single names. This keeps the execution span liquid and auditable while leaving the theory outside any US-equity-specific formulation. The public-neutral synchronization map is ticker agnostic. It can be extended to non-US index options, futures options, country or sector baskets, and other derivative markets if the listed-option span, carrier universe, public state, and execution-cost model are rebuilt with the same adaptedness discipline. A larger N -universe is the next finite-span extension. A broader carrier manifold gives the lower-bound rule more possible synchronized-stress carriers only while liquidity and execution-cost controls remain binding.

Several objects remain outside the present model. The risk-neutral pricing component is empirical rather than a full $\mathbb{Q}$ -dynamics for latent Perron stress. The directed Volterra operator enters through low-dimensional state summaries rather than a pathwise propagation model. The payoff space uses listed-option units and standard Greeks, although the trade is synchronization-based rather than Black–Scholes-based. The public-neutrality constraints are finite-dimensional approximations to the compression and residual-hedging objects in [2, 3]. The lower-bound rule is operationally adapted but sample-developed.

Open extensions are risk-neutral synchronization dynamics, stress-state Greeks, directed Volterra propagation, and fresh-data or paper-trading evaluation of the lower-bound rule. The present result is the finite execution component between latent stress compression and a tradable synchronization book.

10. CONCLUSION

The public-compression and residual-stress geometry of [2, 3] has a finite listed-option representative. That representative is a synchronization book. Its executable signal is a sign-positive wedge in which index-implied synchronization exceeds an adapted realized-synchronization benchmark. The book sells index convexity against a carrier basket of single-name convexity. The trade is admitted only after a lower-bound sellability certificate accounts for spread, carrier support, danger-state reserves, model-risk reserves, and capital-at-risk.

The empirical implementation reports 24.567 million of delta-hedged execution P&L at 1x bid-ask costs, with annualized Sharpe 3.958, ES95 loss 23,683, and max drawdown 364,148. These numbers are reported for an audited, sample-developed research implementation, not a fresh unused-sample trading record. The mechanical claim is adaptedness. Formation weights, lower bounds, and entry-cost proxies are invariant to perturbations of future realized outcomes in the verification audit.

Public option surfaces can leave synchronized stress only partly spanned. The residual wedge is not automatically tradable. It becomes tradable only when the listed-option span supplies an acceptable carrier and the lower-bound certificate is positive. Dormancy is therefore part of the rule's risk discipline. The book stands down when the public option market does not offer a sellable synchronization exposure.

APPENDIX A. MECHANICAL VERIFICATION AND ADAPTEDNESS AUDIT

This appendix records the mechanical outputs supporting the empirical section. The numbers are not additional tuning criteria. They are the audit trail for the frozen listed-option implementation. The tables record the admitted data, constructed windows, lower-bound sizing specification, and invariance of formation weights to future-outcome perturbations.

TABLE A.1. Data-readiness summary for the long OptionMetrics/Oxford–Man implementation. The table records the admitted contract, underlying, and state inputs before the execution ledger is evaluated.

Object	Check	Value
Configuration window	First eligible option date	2000-01-10
Configuration window	Last eligible option date	2018-06-27
Core option universe	Option-ready tickers	32
Core underlying panel	Underlying-ready tickers	32
Core underlying panel	Absent underlying rows after validation	0
Option parser	Raw option rows scanned	74,404,809
Option parser	Parsed quote rows	17,004,261
Option parser	Eligible written quotes	4,951,756
Perron state suite	First early-balanced state date	2000-05-01
Perron state suite	State rows across suites	21,989
Directed-contagion suite	State rows across suites	864,272
State windows	Rolling windows	63, 126, 252, 504

TABLE A.2. Lower-bound execution summary. Candidate rows are sign-positive synchronization books; used rows are rows whose lower-bound size remains positive after reserves and mandate caps.

Quantity	Value
Ledger observations	1,879
Sign-positive candidate books	1,496
Used lower-bound rows	906
Positive lower-bound rows	906
Active dates	805
Total P&L	24.567mn
Annualized Sharpe	3.958
ES95 loss	23,683
Max drawdown	364,148
Total spread cost	7.445mn
Max margin proxy	3.493mn

TABLE A.3. Adaptedness audit for lower-bound sizing. The audit perturbs unsettled formation-date and future realized P&L and realized spread columns, then recomputes row sizes on audited formation dates. The displayed differences are the invariance checks.

Audit quantity	Value
Audited formation dates	96
Maximum weight difference after perturbation	0
Maximum lower-bound difference after perturbation	0
Maximum entry-cost-proxy difference after perturbation	0
Same-day or future realized P&L used for weight	false
Realized exit spread used for weight	false

TABLE A.4. Mandate frontier for the positive-sync capital-risk specification. The controls are mandate inputs rather than signal thresholds. Here λ is the risk-aversion parameter and the cap is the daily overlay margin cap.

Profile	λ	Cap	P&L	Ann. Sharpe	Max DD
0.25, 250k	0.25	250k	21.016mn	3.733	363,986
0.25, 500k	0.25	500k	24.567mn	3.958	364,148
0.25, 750k	0.25	750k	25.537mn	3.978	369,585
0.25, 1m	0.25	1m	25.903mn	3.969	369,585
0.50, 750k	0.50	750k	24.708mn	3.902	353,991
0.50, 1m	0.50	1m	25.041mn	3.891	353,991
1.00, 1m	1.00	1m	23.309mn	3.727	347,618

TABLE A.5. Cost-stress grid for the balanced lower-bound profile. Spread costs are stressed after the same adapted formation weights have been fixed.

Cost multiplier	P&L	Ann. Sharpe	ES95 loss	Max DD
0.5	28.289mn	4.187	20,197	301,264
1.0	24.567mn	3.958	23,683	364,148
2.0	17.122mn	3.259	31,974	501,085
3.0	9.677mn	2.107	43,306	643,465

TABLE A.6. Dormancy and raw-exit diagnostics for 2014–2016. The lower-bound rule stands down in the eligible panel; looser raw-exit probes recover trades while producing expensive, drawdown-heavy books.

Diagnostic	Candidates	Executed	P&L	Ann. Sh.	Max DD
Eligible-panel lower-bound specification	–	0	0	–	0
Raw-exit, 10% spread gate	304	13	-1.300mn	-0.826	1.741mn
Raw-exit, 25% spread gate	304	99	1.967mn	0.366	6.048mn

REFERENCES

- [1] J. Vidal Llauradó (2026). *A Rough Theory of Markets: Latent Contagion, Market Phases, and Market-Variance Risk*. SSRN. <https://doi.org/10.2139/ssrn.6616121>.
- [2] J. Vidal Llauradó (2026). *Risk-Neutral Compression and the Geometry of Latent Market Stress*. SSRN. <https://doi.org/10.2139/ssrn.6680938>.
- [3] J. Vidal Llauradó (2026). *Perron Stress Pricing under Public Compression: Residual Price Intervals, Synchronized Stress, and Visible Hedging*. SSRN. <https://doi.org/10.2139/ssrn.6700598>.
- [4] J. Vidal Llauradó (2026). *Latent Volatility Contagion in Rough Volatility Models: Spectral Thresholds and Detectability*. SSRN. <https://doi.org/10.2139/ssrn.6520718>.
- [5] J. Vidal Llauradó (2026). *Dynamic Observability of Latent Contagion: Sequential Revelation and Observable Screening*. SSRN. <https://doi.org/10.2139/ssrn.6539199>.
- [6] J. Vidal Llauradó (2026). *Detecting Latent Volatility Contagion: A Projected Score Estimator for Rough Volatility Models*. SSRN. <https://doi.org/10.2139/ssrn.6616118>.
- [7] F. Black and M. Scholes, *The pricing of options and corporate liabilities*, Journal of Political Economy **81** (1973), no. 3, 637–654, doi:10.1086/260062.
- [8] R. C. Merton, *Theory of rational option pricing*, The Bell Journal of Economics and Management Science **4** (1973), no. 1, 141–183, doi:10.2307/3003143.
- [9] J. M. Harrison and S. R. Pliska, *Martingales and stochastic integrals in the theory of continuous trading*, Stochastic Processes and their Applications **11** (1981), no. 3, 215–260, doi:10.1016/0304-4149(81)90026-0.
- [10] J. D. Coval and T. Shumway, *Expected option returns*, Journal of Finance **56** (2001), no. 3, 983–1009, doi:10.1111/0022-1082.00352.
- [11] G. Bakshi and N. Kapadia, *Delta-hedged gains and the negative market volatility risk premium*, Review of Financial Studies **16** (2003), no. 2, 527–566, doi:10.1093/rfs/hhg002.
- [12] J. Pan, *The jump-risk premia implicit in options: Evidence from an integrated time-series study*, Journal of Financial Economics **63** (2002), no. 1, 3–50, doi:10.1016/S0304-405X(01)00088-5.
- [13] P. Carr and L. Wu, *Variance risk premiums*, Review of Financial Studies **22** (2009), no. 3, 1311–1341, doi:10.1093/rfs/hhn038.
- [14] R. Cont and J. da Fonseca, *Dynamics of implied volatility surfaces*, Quantitative Finance **2** (2002), no. 1, 45–60, doi:10.1088/1469-7688/2/1/304.
- [15] A. B. Trolle and E. S. Schwartz, *Unspanned stochastic volatility and the pricing of commodity derivatives*, Review of Financial Studies **22** (2009), no. 11, 4423–4461, doi:10.1093/rfs/hhp036.
- [16] J. Gatheral, T. Jaisson, and M. Rosenbaum, *Volatility is rough*, Quantitative Finance **18** (2018), no. 6, 933–949, doi:10.1080/14697688.2017.1393551.
- [17] H. Föllmer and M. Schweizer, *Hedging of contingent claims under incomplete information*, in M. H. A. Davis and R. J. Elliott (eds.), *Applied Stochastic Analysis*, Stochastics Monographs, vol. 5, Gordon and Breach, 1991, 389–414.
- [18] M. Schweizer, *Mean-variance hedging for general claims*, The Annals of Applied Probability **2** (1992), no. 1, 171–179, doi:10.1214/aoap/1177005776.
- [19] J. H. Cochrane and J. Saá-Requejo, *Beyond arbitrage: Good-deal asset price bounds in incomplete markets*, Journal of Political Economy **108** (2000), no. 1, 79–119, doi:10.1086/262112.
- [20] R. T. Rockafellar and S. Uryasev, *Optimization of conditional value-at-risk*, Journal of Risk **2** (2000), no. 3, 21–41, doi:10.21314/JOR.2000.038.
- [21] OptionMetrics, *IvyDB US: Historical option pricing data, volatility, and analytics*, OptionMetrics, accessed 3 May 2026, <https://optionmetrics.com/united-states/>.
- [22] A. Lunde, N. Shephard, and K. Sheppard, *Realized Library*, Oxford–Man Institute of Quantitative Finance, University of Oxford, 2010, <https://realized.oxford-man.ox.ac.uk/>.

Email address: vidal.laurado@gmail.com