

RISK-NEUTRAL COMPRESSION AND THE GEOMETRY OF LATENT MARKET STRESS

J. VIDAL LLAURADÓ 

ABSTRACT. Public option surfaces can price and quote local surface coefficients without determining the latent market-stress loading that moves non-public claims. This paper formulates that boundary as risk-neutral compression. For a square-integrable claim or local option-surface summary X , the $\mathbb{Q}$ -priced structural loading is $L_X^{\mathbb{Q}} = C_{\mathcal{H}\varpi}^{\mathbb{Q}}(X)$. A public information set $\mathcal{H} \subseteq \mathcal{G}_T$ fixes only $C_{\mathcal{H}}^{\mathbb{Q}}(L_X^{\mathbb{Q}})$. The main theorem identifies this conditional projection as the common source of public pricing summaries, option-surface transfers, matched signed $\mathbb{P}/\mathbb{Q}$ coefficient comparisons, nonlinear variance-compression gaps, and visible-hedging residuals. Calibration of a surface therefore leaves latent-state recovery and tradeable hedgeability as separate conditions. Residual components survive as orthogonal projection errors unless the specified public information or hedge span completes the priced structural loading. The SPX/NDX OptionMetrics exercise, rebuilt from raw Oxford–Man realized-volatility data and external option quotes, estimates the finite public projection: implied-volatility level, total-variance level, and signed downside total-variance skew align with the rebuilt Perron stress direction in full-sample and rolling-window designs. Companion A supplies the Gaussian source-family estimates used by the option-surface specialization.

1. INTRODUCTION

Public option surfaces are public price functionals. They can determine a price, a skew coefficient, or a local implied-volatility derivative without determining the latent loading that prices market-variance risk. The paper treats option surfaces, premium comparisons, variance claims, and visible hedge classes as public images of one object, the $\mathbb{Q}$ -priced structural loading.

The latent input is the synchronized market-variance state $(\varpi, \mathcal{H}_{\varpi}, \Delta_{Q_M})$. The market-scale Volterra–Perron construction derives this state from directed contagion and Perron synchronization [1]. The preceding smoothing, dynamic, and projected-score papers separate pricing visibility, path-space detectability, observable screening, and feasible $\mathbb{P}$ -side inference [2, 3, 4]. The present paper changes the measure and changes the question. It asks which part of the resulting priced loading is selected by public option and hedge information under $\mathbb{Q}$.

The first boundary is measure-theoretic. The directed Volterra kernels, rough exponents, signed edge support, and Perron mode of the normalized synchronization matrix are fixed by the primitive geometry. Equivalent pricing reweights conditional expectations. It does not make public screened edge selections, visibility weights, compressed loadings, or claim-specific projections measure invariant unless a separate projection-invariance assumption is imposed. The Perron-generated reduction from [1] supplies the sigma-field and factor on which risk-neutral compression acts. Write

$$(\varpi, \mathcal{H}_{\varpi}, \Delta_{Q_M}) \longmapsto L_X^{\mathbb{Q}} = C_{\mathcal{H}\varpi}^{\mathbb{Q}}(X) \longmapsto C_{\mathcal{H}}^{\mathbb{Q}}(L_X^{\mathbb{Q}}), \quad \mathcal{H} \subseteq \mathcal{G}_T.$$

The market-scale construction of [1] derives the synchronized aggregate state. This paper prices its public images.

The pricing layer is fixed before the empirical layer. Prices are represented by a state-price density; visible pricing summaries are $\mathbb{Q}$ -measurable public summaries; visible hedge

Date: April 2026.

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

Key words and phrases. risk-neutral compression, latent market stress, rough volatility, incomplete markets, option-surface geometry, state-price densities, market-variance risk, visible hedging.

DOI: 10.2139/ssrn.6680938.

classes are closed subspaces of $L^2(\mathbb{Q}, \mathcal{G}_T)$. The martingale-pricing background is the state-price-density and equivalent-martingale-measure tradition [5, 6, 7, 8, 9]. The projection step uses the Hilbert-space geometry common to local-risk-minimization, mean-variance hedging, and incomplete-information hedging [10, 11, 12, 13, 14]. Conditional expectation is the standard L^2 projection geometry [15, 16]. Partial-information investment and information-based pricing provide boundary references for the visible/latent split, not inputs to the theorem [17, 18].

The governing operator is $C_{\mathcal{H}}^{\mathbb{Q}}$. A claim's priced structural loading is visible, screened while still priced through $\mathcal{H}_{\varpi}$, or absent according to $\|C_{\mathcal{H}}^{\mathbb{Q}}(L_X^{\mathbb{Q}})\|_{L^2(\mathbb{Q})}$ and $\|L_X^{\mathbb{Q}}\|_{L^2(\mathbb{Q})}$. A centered aggregate-factor component is then separated inside $L_X^{\mathbb{Q}}$ relative to the market-variance benchmark selected by $\mathcal{H}_{\varpi}$. Public summaries classify the full priced loading. The centered factor enters only when nonlinear compression gaps or visible-hedging residuals expose it.

The same loading then moves through the market objects studied below. Equivalent pricing preserves the mode system and the derived sigma-field. Compression and projection theorems make public pricing and hedging objects factor through $L_X^{\mathbb{Q}}$. Option sections turn that projection into surface coefficients, with signed downside total-variance skew as the tail-facing empirical coordinate. The premium section compares matched signed $\mathbb{P}/\mathbb{Q}$ projected coefficients on a common basis. The variance and hedging sections record the convexity gaps and projection residuals left by the same operator.

Main results. Fix the synchronized Perron state. A claim or local option summary has under $\mathbb{Q}$ the priced structural loading

$$L_X^{\mathbb{Q}} = C_{\mathcal{H}_{\varpi}}^{\mathbb{Q}}(X),$$

and public information $\mathcal{H} \subseteq \mathcal{G}_T$ determines only

$$C_{\mathcal{H}}^{\mathbb{Q}}(L_X^{\mathbb{Q}}).$$

This is the paper's common object. The first results prove that, once the Perron-generated reduction and the equivalent pricing measure are fixed, every treated public pricing or hedging summary is a compression, projection, or measurable transformation of $L_X^{\mathbb{Q}}$. The option-transfer block then states the additional quote-level conditions under which this compressed loading appears as a skew, price-surface, or implied-volatility coefficient. Companion A supplies the primitive Volterra estimates that reduce the Gaussian option-facing loading to the scalar seam used by the local option-factor statement.

The other branches are not parallel theorem copies. The premium branch compares matched signed $\mathbb{P}/\mathbb{Q}$ projected coefficients on a common source-screened basis and uses the comparison for support classification. The variance branch identifies the convexity gaps generated by the aggregate factor and by recentering defects. The hedging branch expresses visible incompleteness as the distance from $L_X^{\mathbb{Q}}$ to a visible hedge space.

The governing results are Theorems 3.1, 4.12, 6.3, 6.7, 9.7, 10.16, and 11.1. The appendices supply the endpoint calculations: option-transfer variants in Appendix A, premium estimator and support details in Appendix B, endpoint obstructions in Appendix C, and Gaussian source-family estimates from Companion A, summarized in Propositions D.53, D.58, and D.59.

The Perron-market construction supplies the latent stress coordinate and visible channel geometry. The present analysis identifies its risk-neutral public projections.

Organization. Sections 2–5 define the Perron-generated reduction, prove invariance of the primitive geometry under equivalent pricing, define $L_X^{\mathbb{Q}}$, and classify public visibility by the norm of the compressed loading. Theorem 4.12 is the common projection identity.

Section 6 gives the option-facing pricing setup, the abstract first-variation statement, and the cumulant-class transfer from compressed loading to quoted surface coefficient. Section 7 records the option consequence used empirically. The Hermite, primitive non-Gaussian, rare-jump, Gaussian-Volterra, Bachelier, and Black calculations are in Appendices A, D, and E; Companion A supplies the Gaussian source-family verification [19].

Section 8 reports the SPX/NDX finite-projection evidence. Section 9 compares harmonized signed $\mathbb{P}/\mathbb{Q}$ projected coefficients on a cleaned matched derivative surface. Sections 10 and 11

treat nonlinear variance compression and visible hedge residuals. Section 12 states the endpoint scope, and Appendix C records the detailed obstructions. Section 13 closes the argument.

2. LATENT MARKET GEOMETRY UNDER $\mathbb{P}$ AND $\mathbb{Q}$

Let $N \geq 2$ be fixed and let $W = (W^1, \dots, W^N)$ be an N -dimensional Brownian motion with independent coordinates on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$. The Volterra kernel specification used below follows the rough-volatility and affine-Volterra lineage [20, 21, 22, 23], with broader survey and calibration-context references in [24, 25]. The screened Perron construction is taken from [1]; the Perron–Frobenius terminology follows the standard nonnegative-matrix background [26, 27].

Assumption 2.1 (Directed Volterra market under $\mathbb{P}$). For each $1 \leq i, j \leq N$, there exist exponent $H_{ij} \in (0, 1)$, amplitude $\beta_{ij} \in \mathbb{R}$, and continuous modulation $a_{ij}: [0, T]^2 \rightarrow \mathbb{R}$ such that for $0 \leq s < t \leq T$,

$$K_{ij}(t, s) = \beta_{ij} a_{ij}(t, s) (t - s)^{H_{ij}-1/2}.$$

Moreover $a_{ij}(t, t) > 0$ whenever $\beta_{ij} \neq 0$, and the drift $t \mapsto \mu_i^{\mathbb{P}}(t)$ is deterministic and continuous. The latent volatility channel of asset i is

$$V_i(t) = \mu_i^{\mathbb{P}}(t) + \sum_{j=1}^N \int_0^t K_{ij}(t, s) dW_s^j.$$

Fix deterministic weights $w \in (0, \infty)^N$ and define the aggregate market-volatility process

$$M(t) := \sum_{i=1}^N w_i V_i(t).$$

The latent contagion geometry is encoded by the signed edge matrix and the screened synchronization matrix of [1]. Fix a reference time $t_0 \in (0, T]$, screening weights $\pi_{ij} \in [0, 1]$, and a market-dominant active set $\mathcal{E}_M \subseteq \{1, \dots, N\}^2$. Define the structural signed edge matrix

$$B_{ij} := \beta_{ij} a_{ij}(t_0, t_0) \mathbf{1}_{\{(i,j) \in \mathcal{E}_M\}}, \quad 1 \leq i, j \leq N.$$

The nonnegative screened synchronization matrix is

$$A_{ij} := |B_{ij}| \pi_{ij}.$$

The weights π_{ij} are normalized visibility weights for the local stress-envelope experiment. They are part of the observational projection, not primitive edge signs. Throughout, work on the irreducible aggregate regime inherited from [1]. In particular, A is assumed irreducible; let $v \in (0, \infty)^N$ be its Perron right eigenvector and let $\text{spr}(A)$ denote its spectral radius.

This separates structural and observational edge objects. The structural edge set

$$E^{\text{str}} := \{(i, j) : \beta_{ij} a_{ij}(t_0, t_0) \neq 0\}$$

is part of the primitive Volterra geometry. Physical screened selections $E^{\text{vis}, \mathbb{P}}$ and weights $\pi_{ij}^{\mathbb{P}}$, and risk-neutral or option-surface projected selections $E^{\text{vis}, \mathbb{Q}}$ and weights $\pi_{ij}^{\mathbb{Q}}$, are observational objects. Equivalent pricing preserves the primitive directed Volterra geometry under the stated representation assumptions. It does not, without an additional projection-invariance assumption, preserve finite-sample screened edge selections or empirical visibility weights.

Assumption 2.2 (Equivalent square-integrable pricing measure). There exists a predictable $\mathbb{R}^N$ -valued process θ such that

$$\int_0^T \|\theta_s\|^2 ds < \infty \quad \mathbb{P}\text{-a.s.},$$

and the stochastic exponential

$$Z_t := \exp\left(-\int_0^t \theta_s^\top dW_s - \frac{1}{2} \int_0^t \|\theta_s\|^2 ds\right)$$

is a strictly positive martingale. Define $\mathbb{Q}$ by

$$\frac{d\mathbb{Q}}{d\mathbb{P}} = Z_T.$$

All claims considered below are assumed to belong to $L^2(\mathbb{Q})$. When the notation is used for discounted no-arbitrage prices of traded assets, those discounted traded assets are $\mathbb{Q}$ -local martingales.

Condition 2.3 (Optional bounded-density L^2 -comparability). The state-price density satisfies

$$0 < \underline{z} \leq Z_T \leq \bar{z} < \infty \quad \mathbb{P}\text{-a.s.}$$

for some deterministic constants $\underline{z}, \bar{z}$. This optional condition is invoked only in comparability statements that identify $L^2(\mathbb{P})$ and $L^2(\mathbb{Q})$ setwise. The main risk-neutral compression results are stated natively on $L^2(\mathbb{Q})$ and do not require this bounded-density condition.

Assumption 2.4 (Aggregate Perron sufficient-state reduction). For the aggregate claim families treated here, the market-scale construction of [1] supplies a Perron-projected amplitude ϖ , a visible sigma-field $\mathcal{G}_T$, a derived structural sigma-field $\mathcal{H}_\varpi$, and, for integrated market variance, a centered structural increment Δ_{Q_M} . These objects satisfy the aggregate Perron sufficient-state reduction

$$(\varpi, \mathcal{H}_\varpi, \Delta_{Q_M}) \text{ is sufficient for the aggregate claim families considered below.}$$

The analysis below studies the $\mathbb{Q}$ -priced public compressions generated by this reduction.

Definition 2.5 (Aggregate Perron structural sigma-field). Assume the aggregate claim families treated here satisfy the Perron-projected sufficient-state reduction of [1]. Let $u, v \in (0, \infty)^N$ be the left and right Perron eigenvectors of the irreducible screened synchronization matrix, normalized so that

$$u^\top v = 1.$$

Let x denote the aggregate latent envelope from the market-scale construction of [1] and define the canonically projected aggregate amplitude by

$$\varpi(t) := u^\top x(t), \quad 0 \leq t \leq T.$$

Assume

$$\int_0^T \varpi(t)^2 dt \in L^2(\mathbb{P}, \mathcal{F}_T) \cap L^2(\mathbb{Q}, \mathcal{F}_T).$$

Set

$$\mathcal{L}_T := \sigma(\varpi(t) : 0 \leq t \leq T), \quad \mathcal{H}_\varpi := \mathcal{G}_T \vee \mathcal{L}_T.$$

Thus $\mathcal{H}_\varpi$ is the derived aggregate structural sigma-field inherited from [1] for the aggregate claim classes treated here; for integrated market variance, the same reduction selects the centered structural increment

$$\Delta_{Q_M} = C_{\mathcal{H}_\varpi}(Q_M) - C_{\mathcal{G}_T}(Q_M).$$

[1] derives these objects from primitive aggregate coupling, Perron projection, and off-Perron aggregate screening, and illustrates the reduction empirically on the Oxford-Man realized-measure panel [28, 29, 30].

Remark 2.6 (Aggregate reduction, not a fresh existence claim). Definition 2.5 imports the aggregate-first reduction established in [1]: the Perron-projected amplitude $\varpi = u^\top x$, the derived sigma-field $\mathcal{H}_\varpi$, the market-variance increment Δ_{Q_M} , and the claim families for which that state is sufficient. The present argument starts after that reduction and identifies the corresponding $\mathbb{Q}$ -priced loadings and their public compressions, with no new Perron-factor existence assumption inside this manuscript.

Definition 2.7 (Pricing, visible summaries, and risk-neutral compression). Let $\mathcal{G}_t \subseteq \mathcal{F}_t$ be the visible market filtration. For a discounted claim $X \in L^2(\mathbb{Q}, \mathcal{F}_T)$, define the pricing operator

$$\pi_t(X) := \mathbb{E}_{\mathbb{Q}}[X \mid \mathcal{F}_t].$$

For a date- t public information set $\mathcal{H}_t \subseteq \mathcal{G}_t$, define the public date- t compression by

$$C_{t,\mathcal{H}}^{\mathbb{Q}}(X) := \mathbb{E}_{\mathbb{Q}}[X \mid \mathcal{H}_t].$$

For every sub- σ -field $\mathcal{H} \subseteq \mathcal{G}_T$, define the risk-neutral compression of X onto $\mathcal{H}$ by

$$C_{\mathcal{H}}^{\mathbb{Q}}(X) := \mathbb{E}_{\mathbb{Q}}[X \mid \mathcal{H}].$$

A *visible pricing summary* is any $\mathfrak{S} \in L^2(\mathbb{Q}, \mathcal{G}_T)$, and a *visible hedge class* is any closed linear subspace $\mathcal{S} \subseteq L^2(\mathbb{Q}, \mathcal{G}_T)$. The time-zero option calculations below use the terminal notation $C_{\mathcal{H}}^{\mathbb{Q}}$ after suppressing the observation date. The empirical panel is the repeated-date analogue of the same local conditional-pricing operation, with each option surface measurable with respect to the public information available on its observation date.

Definition 2.8 (Derived sufficient structural sigma-field and structural claim spaces). Under Definition 2.5, define the $\mathbb{Q}$ -side structural claim space

$$\mathcal{M}_{\mathbb{Q}} := L^2(\mathbb{Q}, \mathcal{H}_{\varpi}),$$

and the corresponding $\mathbb{P}$ -side structural claim space

$$\mathcal{M}_{\mathbb{P}} := L^2(\mathbb{P}, \mathcal{H}_{\varpi}).$$

Remark 2.9 (Scope of the pricing layer). The state-price density gives an arbitrage-free pricing operator on discounted claims, but the analysis remains incomplete-market in scope. It does not seek a representative-agent foundation or a market-completion theorem. This is the same martingale-valuation separation between pricing measures and market completeness that underlies [5, 8, 9]. The pricing layer identifies which parts of latent market stress remain publicly visible under $\mathbb{Q}$ and which survive only through compression gaps and residual hedge risk.

3. STRUCTURAL GEOMETRY AND MEASURE CHANGE

Equivalent pricing preserves the primitive Volterra–Perron geometry but not the conditional projection operators. The first theorem fixes that separation before any visibility or hedging claim is made.

Theorem 3.1 (Geometry invariance under equivalent pricing). *Assume Assumptions 2.1 and 2.2, together with Definition 2.5. Then there exists an N -dimensional Brownian motion $W^{\mathbb{Q}}$ on $(\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbb{Q})$ such that for every asset i ,*

$$V_i(t) = \mu_i^{\mathbb{Q}}(t) + \sum_{j=1}^N \int_0^t K_{ij}(t, s) dW_s^{\mathbb{Q},j},$$

where the possibly random Volterra drift shift is

$$\mu_i^{\mathbb{Q}}(t) := \mu_i^{\mathbb{P}}(t) - \sum_{j=1}^N \int_0^t K_{ij}(t, s) \theta_s^j ds.$$

It is deterministic, or of finite variation as a function of t , only under additional regularity restrictions on the Girsanov kernel and the Volterra kernels; no subsequent structural statement uses such restrictions. In particular:

- (i) *the kernels K_{ij} , the exponents H_{ij} , the structural active set $\mathcal{E}_M$, the signed edge matrix B , the normalized synchronization matrix A , its spectral radius $\text{spr}(A)$, and its Perron mode v are the same under $\mathbb{P}$ and $\mathbb{Q}$, subject to the declared representation and normalization assumptions;*
- (ii) *the structural mode sigma-field $\mathcal{H}_{\varpi} = \mathcal{G}_T \vee \mathcal{L}_T$ is the same under $\mathbb{P}$ and $\mathbb{Q}$; consequently the mode classes $\mathcal{M}_{\mathbb{P}}$ and $\mathcal{M}_{\mathbb{Q}}$ are built from the same underlying information, although they need not coincide as sets without additional integrability or comparability conditions;*
- (iii) *the structural projections $C_{\mathcal{H}_{\varpi}}^{\mathbb{P}}$ and $C_{\mathcal{H}_{\varpi}}^{\mathbb{Q}}$ act on the same mode-generating sigma-field, but they remain measure-dependent. Thus equivalent pricing preserves the latent geometry while potentially changing its public projection.*

Proof. By Girsanov's theorem [16],

$$W_t^{\mathbb{Q}} := W_t + \int_0^t \theta_s \, ds$$

is an N -dimensional Brownian motion under $\mathbb{Q}$. Substituting $dW_s = dW_s^{\mathbb{Q}} - \theta_s \, ds$ into the Volterra representation gives

$$V_i(t) = \mu_i^{\mathbb{P}}(t) + \sum_{j=1}^N \int_0^t K_{ij}(t, s) \, dW_s^{\mathbb{Q},j} - \sum_{j=1}^N \int_0^t K_{ij}(t, s) \theta_s^j \, ds,$$

which is the claimed representation. Since the kernels themselves are unchanged, so are the exponents, the active set, the screened operator, and the Perron system built from them.

Because the aggregate Perron amplitude $\varpi = u^\top x$ is fixed by the underlying measurable structure rather than by either probability measure, its path sigma-field $\mathcal{L}_T = \sigma(\varpi(t) : 0 \leq t \leq T)$ is unchanged by the equivalent measure change. Hence the sigma-field $\mathcal{H}_\varpi = \mathcal{G}_T \vee \mathcal{L}_T$ is the same measurable object under $\mathbb{P}$ and $\mathbb{Q}$, and the spaces $\mathcal{M}_{\mathbb{P}}$ and $\mathcal{M}_{\mathbb{Q}}$ are simply the respective L^2 -classes on that common sigma-field. The last statement follows because conditional expectation is determined by both the sub- σ -field and the measure. The sub- σ -field is common, but the $L^2(\mathbb{P})$ - and $L^2(\mathbb{Q})$ -inner products, and therefore the corresponding orthogonal projections, may differ. $\square$

Proposition 3.2 (Optional L^2 -comparability under bounded density). *Under Assumption 2.2 and Condition 2.3, the spaces $L^2(\mathbb{P})$ and $L^2(\mathbb{Q})$ coincide as sets and have equivalent norms. Consequently $\mathcal{M}_{\mathbb{P}}$ and $\mathcal{M}_{\mathbb{Q}}$ also coincide as sets and have equivalent norms.*

Proof. Let X be measurable. If $X \in L^2(\mathbb{P})$, then Condition 2.3 gives

$$\mathbb{E}_{\mathbb{Q}}[X^2] = \mathbb{E}_{\mathbb{P}}[Z_T X^2] \leq \bar{z} \mathbb{E}_{\mathbb{P}}[X^2] < \infty.$$

Conversely, if $X \in L^2(\mathbb{Q})$, then

$$\mathbb{E}_{\mathbb{P}}[X^2] = \mathbb{E}_{\mathbb{Q}}\left[\frac{X^2}{Z_T}\right] \leq \underline{z}^{-1} \mathbb{E}_{\mathbb{Q}}[X^2] < \infty.$$

So $L^2(\mathbb{P}) = L^2(\mathbb{Q})$ as sets. The two norms are equivalent because

$$\underline{z} \mathbb{E}_{\mathbb{P}}[X^2] \leq \mathbb{E}_{\mathbb{Q}}[X^2] \leq \bar{z} \mathbb{E}_{\mathbb{P}}[X^2]$$

whenever $X \in L^2(\mathbb{P}) = L^2(\mathbb{Q})$. Restricting to $(\mathcal{G}_T \vee \mathcal{L}_T)$ -measurable random variables gives the same conclusion for $\mathcal{M}_{\mathbb{P}}$ and $\mathcal{M}_{\mathbb{Q}}$. $\square$

Corollary 3.3 (Same geometry, different visible projections). *Equivalent pricing preserves the structural mode system and the mode-generating sigma-field $\mathcal{H}_\varpi$, but not its public visibility. The core measure-invariant objects are the kernels, the screened Perron construction, and $\mathcal{H}_\varpi$. The measure-dependent objects are the visible projections $C_{\mathcal{G}_T}^{\mathbb{P}}$ and $C_{\mathcal{G}_T}^{\mathbb{Q}}$ together with the structural projections $C_{\mathcal{H}_\varpi}^{\mathbb{P}}$ and $C_{\mathcal{H}_\varpi}^{\mathbb{Q}}$. If Condition 2.3 also holds, one may additionally identify $\mathcal{M}_{\mathbb{P}}$ and $\mathcal{M}_{\mathbb{Q}}$ setwise by Proposition 3.2.*

Proof. Theorem 3.1 fixes the measurable sigma-field $\mathcal{H}_\varpi$ and the kernel-Perron objects under the equivalent change of measure. It does not identify the conditional-expectation operator. For any sub- σ -field $\mathcal{A} \subseteq \mathcal{H}_\varpi$, Bayes' formula gives, whenever the expression is defined,

$$C_{\mathcal{A}}^{\mathbb{Q}}(Y) = \frac{\mathbb{E}_{\mathbb{P}}[Z_T Y \mid \mathcal{A}]}{\mathbb{E}_{\mathbb{P}}[Z_T \mid \mathcal{A}]},$$

whereas $C_{\mathcal{A}}^{\mathbb{P}}(Y) = \mathbb{E}_{\mathbb{P}}[Y \mid \mathcal{A}]$. The range sigma-field is the same, but the projection may change with the state-price density. Proposition 3.2 gives the stronger optional setwise identification of the corresponding L^2 -classes under bounded-density comparability. $\square$

Remark 3.4 (Same latent geometry). The phrase means the same kernel system, the same screened Perron construction, and the same mode-generating sigma-field. Under the optional Condition 2.3, the $\mathbb{P}$ - and $\mathbb{Q}$ -side mode classes may also be identified setwise. It does not mean that the $\mathbb{P}$ - and $\mathbb{Q}$ -conditional projections of a fixed claim agree. The distinction is exactly that they need not.

4. RISK-NEUTRAL COMPRESSION AND PROJECTION

The compression problem begins after the structural sigma-field has been fixed. The map $C_{\mathcal{H}_\varpi}^{\mathbb{Q}}$ extracts the priced structural loading, and the smaller map $C_{\mathcal{H}}^{\mathbb{Q}}$ records the part admitted by public information.

Proposition 4.1 (Optional projection stability under bounded-density comparability). *Under Assumption 2.2 and Condition 2.3, fix a sub- σ -field $\mathcal{H} \subseteq \mathcal{H}_\varpi$. Then $L^2(\mathbb{P}, \mathcal{H}) = L^2(\mathbb{Q}, \mathcal{H})$ as sets, and the operators*

$$C_{\mathcal{H}}^{\mathbb{P}}: \mathcal{M}_{\mathbb{P}} = \mathcal{M}_{\mathbb{Q}} \rightarrow L^2(\mathbb{P}, \mathcal{H}), \quad C_{\mathcal{H}}^{\mathbb{Q}}: \mathcal{M}_{\mathbb{P}} = \mathcal{M}_{\mathbb{Q}} \rightarrow L^2(\mathbb{Q}, \mathcal{H})$$

are bounded linear projections on the same ambient claim class. For every $X \in \mathcal{M}_{\mathbb{P}} = \mathcal{M}_{\mathbb{Q}}$,

$$\|C_{\mathcal{H}}^{\mathbb{P}}X\|_{L^2(\mathbb{Q})} \leq \left(\frac{\bar{z}}{z}\right)^{1/2} \|X\|_{L^2(\mathbb{Q})},$$

$$\|C_{\mathcal{H}}^{\mathbb{Q}}X\|_{L^2(\mathbb{P})} \leq \left(\frac{\bar{z}}{z}\right)^{1/2} \|X\|_{L^2(\mathbb{P})}.$$

Hence the $\mathbb{Q}$ -side visible projection is a continuous reweighting of the same structural span, not a new geometric object. This is an optional comparability statement; the compression theorems use the native $L^2(\mathbb{Q})$ Hilbert space and do not rely on this bounded-density identification.

Proof. By Proposition 3.2, $L^2(\mathbb{P}) = L^2(\mathbb{Q})$ as sets with equivalent norms, and the same holds after restricting to $\mathcal{H}$ -measurable random variables. Since conditional expectations are orthogonal projections in their native Hilbert norms,

$$\|C_{\mathcal{H}}^{\mathbb{P}}X\|_{L^2(\mathbb{P})} \leq \|X\|_{L^2(\mathbb{P})}, \quad \|C_{\mathcal{H}}^{\mathbb{Q}}X\|_{L^2(\mathbb{Q})} \leq \|X\|_{L^2(\mathbb{Q})}.$$

Applying the norm equivalence from Proposition 3.2 to the projected terms gives the displayed bounds. The final statement is the interpretation of the fact that the range $L^2(\mathcal{H})$ is fixed while the projection operator changes continuously with the weighted norm. $\square$

Proposition 4.2 (Supporting structural projection of claims). *Under Assumptions 2.1 and 2.2, together with Definition 2.8, every claim $X \in L^2(\mathbb{Q}, \mathcal{F}_T)$ admits a unique orthogonal decomposition*

$$X = L_X^{\mathbb{Q}} + U_X^{\mathbb{Q}},$$

where

$$L_X^{\mathbb{Q}} := C_{\mathcal{H}_\varpi}^{\mathbb{Q}}(X) \in \mathcal{M}_{\mathbb{Q}}, \quad U_X^{\mathbb{Q}} := X - C_{\mathcal{H}_\varpi}^{\mathbb{Q}}(X) \perp_{L^2(\mathbb{Q})} \mathcal{M}_{\mathbb{Q}}.$$

Equivalently,

$$\mathbb{E}_{\mathbb{Q}}[U_X^{\mathbb{Q}} \mid \mathcal{H}_\varpi] = 0.$$

For every public sub- σ -field $\mathcal{H} \subseteq \mathcal{G}_T$,

$$C_{\mathcal{H}}^{\mathbb{Q}}(X) = C_{\mathcal{H}}^{\mathbb{Q}}(L_X^{\mathbb{Q}}).$$

If $L_X^{\mathbb{Q}} \neq 0$, define

$$\alpha_X^{\mathbb{Q}} := \|L_X^{\mathbb{Q}}\|_{L^2(\mathbb{Q})}, \quad \Lambda_X := \frac{L_X^{\mathbb{Q}}}{\alpha_X^{\mathbb{Q}}};$$

if $L_X^{\mathbb{Q}} = 0$, set $\alpha_X^{\mathbb{Q}} := 0$ and $\Lambda_X := 0$. The term $L_X^{\mathbb{Q}}$ is the claim's priced structural loading under $\mathbb{Q}$.

The full loading and the residual loadings have different jobs. For every public sub- σ -field $\mathcal{H} \subseteq \mathcal{G}_T$ and every closed tradeable span $\mathcal{S} \subseteq L^2(\mathbb{Q}, \mathcal{H})$, write

$$R_{X,\mathcal{H}}^{\mathbb{Q}} := L_X^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}}(L_X^{\mathbb{Q}}), \quad R_{X,\mathcal{S}}^{\mathbb{Q}} := L_X^{\mathbb{Q}} - C_{\mathcal{S}}^{\mathbb{Q}}(L_X^{\mathbb{Q}}).$$

Public compression concerns $C_{\mathcal{H}}^{\mathbb{Q}}(L_X^{\mathbb{Q}})$. Price intervals and convex residual charges concern $R_{X,\mathcal{H}}^{\mathbb{Q}}$. Visible hedge errors concern $R_{X,\mathcal{S}}^{\mathbb{Q}}$.

Proof. Because $\mathcal{M}_{\mathbb{Q}} = L^2(\mathbb{Q}, \mathcal{H}_{\varpi})$ is a closed subspace of $L^2(\mathbb{Q}, \mathcal{F}_T)$, the Hilbert projection theorem yields the unique orthogonal projection $P_{\mathcal{M}_{\mathbb{Q}}}^{\mathbb{Q}}(X)$. Conditional expectation onto $\mathcal{H}_{\varpi}$ is exactly this projection, so

$$L_X^{\mathbb{Q}} = C_{\mathcal{H}_{\varpi}}^{\mathbb{Q}}(X)$$

and $U_X^{\mathbb{Q}} = X - L_X^{\mathbb{Q}}$ is orthogonal to $\mathcal{M}_{\mathbb{Q}}$, equivalently $\mathbb{E}_{\mathbb{Q}}[U_X^{\mathbb{Q}} \mid \mathcal{H}_{\varpi}] = 0$. Since every public sigma-field $\mathcal{H} \subseteq \mathcal{G}_T \subseteq \mathcal{H}_{\varpi}$, the tower property gives

$$C_{\mathcal{H}}^{\mathbb{Q}}(X) = C_{\mathcal{H}}^{\mathbb{Q}}(C_{\mathcal{H}_{\varpi}}^{\mathbb{Q}}(X)) = C_{\mathcal{H}}^{\mathbb{Q}}(L_X^{\mathbb{Q}}).$$

The remaining statements are notation. $\square$

Proposition 4.3 (Approximate sufficiency of the Perron-generated aggregate structural sigma-field). *Let $X \in L^2(\mathbb{Q}, \mathcal{F}_T)$, and suppose that*

$$X = X_{\varpi} + R_X,$$

where $X_{\varpi} \in L^2(\mathbb{Q}, \mathcal{H}_{\varpi})$. Let

$$L_X^{\mathbb{Q}} := C_{\mathcal{H}_{\varpi}}^{\mathbb{Q}}(X)$$

be the priced structural loading from Proposition 4.2. Then

$$L_X^{\mathbb{Q}} = X_{\varpi} + C_{\mathcal{H}_{\varpi}}^{\mathbb{Q}}(R_X).$$

Consequently,

$$\|L_X^{\mathbb{Q}} - X_{\varpi}\|_{L^2(\mathbb{Q})} \leq \|R_X\|_{L^2(\mathbb{Q})}.$$

Moreover, for every public sub- σ -field $\mathcal{H} \subseteq \mathcal{G}_T$,

$$\|C_{\mathcal{H}}^{\mathbb{Q}}(L_X^{\mathbb{Q}}) - C_{\mathcal{H}}^{\mathbb{Q}}(X_{\varpi})\|_{L^2(\mathbb{Q})} \leq \|R_X\|_{L^2(\mathbb{Q})},$$

and for every visible hedge class $\mathcal{S} \subseteq L^2(\mathbb{Q}, \mathcal{G}_T)$,

$$\|P_{\mathcal{S}}^{\mathbb{Q}}(L_X^{\mathbb{Q}}) - P_{\mathcal{S}}^{\mathbb{Q}}(X_{\varpi})\|_{L^2(\mathbb{Q})} \leq \|R_X\|_{L^2(\mathbb{Q})}.$$

Thus, if $\|R_X\|_{L^2(\mathbb{Q})} \leq \delta_X$, replacing fuller aggregate structural content by its $\mathcal{H}_{\varpi}$ -measurable Perron-generated component perturbs structural projections, public compressions, and visible projections by at most δ_X .

Proof. Because X_{ϖ} is $\mathcal{H}_{\varpi}$ -measurable,

$$L_X^{\mathbb{Q}} = C_{\mathcal{H}_{\varpi}}^{\mathbb{Q}}(X_{\varpi} + R_X) = X_{\varpi} + C_{\mathcal{H}_{\varpi}}^{\mathbb{Q}}(R_X),$$

which gives the first display. Since conditional expectation is a contraction on $L^2(\mathbb{Q})$,

$$\|L_X^{\mathbb{Q}} - X_{\varpi}\|_{L^2(\mathbb{Q})} = \|C_{\mathcal{H}_{\varpi}}^{\mathbb{Q}}(R_X)\|_{L^2(\mathbb{Q})} \leq \|R_X\|_{L^2(\mathbb{Q})}.$$

Applying the contractions $C_{\mathcal{H}}^{\mathbb{Q}}$ and $P_{\mathcal{S}}^{\mathbb{Q}}$ to $L_X^{\mathbb{Q}} - X_{\varpi}$ yields the remaining bounds. $\square$

Remark 4.4 (Economic reading of $\mathcal{H}_{\varpi}$ for aggregate claim families). Proposition 4.3 is deterministic. If the off-Perron aggregate remainder is small in $L^2(\mathbb{Q})$, then the Perron-generated sigma-field $\mathcal{H}_{\varpi}$ is sufficient for the structural reduction and approximately sufficient in operator norm for the aggregate claim families inherited from [1]. Read together with the Oxford-Man appendix of [1], where the aggregate off-Perron share is empirically reported to be small for the market-stress proxy treated there [28, 29, 30], the sigma-field $\mathcal{H}_{\varpi}$ has a specific economic interpretation. For the aggregate claims studied in the market-scale analysis of [1] and in the present analysis, it is the empirically dominant enlargement in that reported design.

Definition 4.5 (Canonical centered aggregate factor direction). For the aggregate claim families inherited from [1], let

$$\Lambda_M^* \in L^2(\mathbb{Q}, \mathcal{H}_\varpi)$$

denote the centered aggregate factor direction obtained by normalizing the canonical centered structural increment of integrated market variance under the present pricing measure. Thus

$$\mathbb{E}_\mathbb{Q}[\Lambda_M^* \mid \mathcal{G}_T] = 0, \quad \|\Lambda_M^*\|_{L^2(\mathbb{Q})} = 1$$

whenever the centered increment is nonzero; if that centered increment vanishes, set $\Lambda_M^* := 0$. Define the centered structural subspace

$$\mathcal{M}_\mathbb{Q}^\circ := \{Y \in L^2(\mathbb{Q}, \mathcal{H}_\varpi) : \mathbb{E}_\mathbb{Q}[Y \mid \mathcal{G}_T] = 0\}.$$

Then

$$\mathcal{M}_\mathbb{Q}^\circ = \mathcal{M}_\mathbb{Q} \ominus L^2(\mathbb{Q}, \mathcal{G}_T),$$

and $\Lambda_M^* \in \mathcal{M}_\mathbb{Q}^\circ$.

Proposition 4.6 (Canonical aggregate factor decomposition inside the priced structural loading). *Under Proposition 4.2 and Definition 4.5, every claim $X \in L^2(\mathbb{Q}, \mathcal{F}_T)$ admits a unique decomposition*

$$L_X^\mathbb{Q} = V_X^\mathbb{Q} + \alpha_X^{\mathbb{Q},*} \Lambda_M^* + R_X^{\mathbb{Q},*},$$

where

$$V_X^\mathbb{Q} := C_{\mathcal{G}_T}^\mathbb{Q}(L_X^\mathbb{Q}) \in L^2(\mathbb{Q}, \mathcal{G}_T), \\ \alpha_X^{\mathbb{Q},*} \in \mathbb{R}, \quad R_X^{\mathbb{Q},*} \in \mathcal{M}_\mathbb{Q}^\circ, \quad R_X^{\mathbb{Q},*} \perp_{L^2(\mathbb{Q})} \Lambda_M^*.$$

Equivalently, if

$$\tilde{L}_X^\mathbb{Q} := L_X^\mathbb{Q} - C_{\mathcal{G}_T}^\mathbb{Q}(L_X^\mathbb{Q}) \in \mathcal{M}_\mathbb{Q}^\circ,$$

then

$$\tilde{L}_X^\mathbb{Q} = \alpha_X^{\mathbb{Q},*} \Lambda_M^* + R_X^{\mathbb{Q},*}, \quad \alpha_X^{\mathbb{Q},*} = \langle \tilde{L}_X^\mathbb{Q}, \Lambda_M^* \rangle_{L^2(\mathbb{Q})}.$$

If $\Lambda_M^* = 0$, then $\alpha_X^{\mathbb{Q},*} := 0$ and $R_X^{\mathbb{Q},*} := \tilde{L}_X^\mathbb{Q}$. A claim is called aggregate-factor aligned when

$$R_X^{\mathbb{Q},*} = 0.$$

Proof. The space $\mathcal{M}_\mathbb{Q}^\circ$ is a closed subspace of $L^2(\mathbb{Q}, \mathcal{F}_T)$. Since $L_X^\mathbb{Q} \in \mathcal{M}_\mathbb{Q}$ and $\tilde{L}_X^\mathbb{Q} = L_X^\mathbb{Q} - \mathbb{E}_\mathbb{Q}[L_X^\mathbb{Q}]$ has zero $\mathbb{Q}$ -mean, $\tilde{L}_X^\mathbb{Q} \in \mathcal{M}_\mathbb{Q}^\circ$. If $\Lambda_M^* = 0$, set $\alpha_X^{\mathbb{Q},*} := 0$ and $R_X^{\mathbb{Q},*} := \tilde{L}_X^\mathbb{Q}$. Otherwise, $\text{span}\{\Lambda_M^*\}$ is a one-dimensional closed subspace of $\mathcal{M}_\mathbb{Q}^\circ$, so the Hilbert projection theorem yields the unique orthogonal decomposition

$$\tilde{L}_X^\mathbb{Q} = \langle \tilde{L}_X^\mathbb{Q}, \Lambda_M^* \rangle_{L^2(\mathbb{Q})} \Lambda_M^* + R_X^{\mathbb{Q},*}$$

with $R_X^{\mathbb{Q},*} \perp \Lambda_M^*$. Adding back $V_X^\mathbb{Q} = C_{\mathcal{G}_T}^\mathbb{Q}(L_X^\mathbb{Q})$ gives the stated claim. $\square$

Corollary 4.7 (Linear public screening of the centered aggregate factor). *Under Definition 4.5, for every public sub- σ -field $\mathcal{H} \subseteq \mathcal{G}_T$,*

$$C_{\mathcal{H}}^\mathbb{Q}(\Lambda_M^*) = 0.$$

Consequently, if X is aggregate-factor aligned so that

$$L_X^\mathbb{Q} = V_X^\mathbb{Q} + \alpha_X^{\mathbb{Q},*} \Lambda_M^*,$$

then

$$C_{\mathcal{H}}^\mathbb{Q}(L_X^\mathbb{Q}) = C_{\mathcal{H}}^\mathbb{Q}(V_X^\mathbb{Q}).$$

Thus the canonical centered aggregate factor is always screened from linear public compression, even when its coefficient $\alpha_X^{\mathbb{Q},*}$ is nonzero inside the priced structural loading.

Proof. Because $\Lambda_M^* \in \mathcal{M}_\mathbb{Q}^\circ$, one has $\mathbb{E}_\mathbb{Q}[\Lambda_M^* \mid \mathcal{G}_T] = 0$. Since $\mathcal{H} \subseteq \mathcal{G}_T$, the tower property gives

$$C_\mathcal{H}^\mathbb{Q}(\Lambda_M^*) = \mathbb{E}_\mathbb{Q}[\mathbb{E}_\mathbb{Q}[\Lambda_M^* \mid \mathcal{G}_T] \mid \mathcal{H}] = 0.$$

Linearity of $C_\mathcal{H}^\mathbb{Q}$ gives

$$C_\mathcal{H}^\mathbb{Q}(L_X^\mathbb{Q}) = C_\mathcal{H}^\mathbb{Q}(V_X^\mathbb{Q}) + \alpha_X^{\mathbb{Q},*} C_\mathcal{H}^\mathbb{Q}(\Lambda_M^*) = C_\mathcal{H}^\mathbb{Q}(V_X^\mathbb{Q}),$$

which proves the aligned-claim identity. $\square$

Remark 4.8 (Structural objects versus priced claim objects). The measure-invariant objects are the kernel system, the screened Perron construction, and the derived sufficient structural sigma-field $\mathcal{H}_\varpi$. By contrast, the priced structural loading $L_X^\mathbb{Q}$ is claim-specific and measure-dependent. Its scalar weight $\alpha_X^\mathbb{Q}$ is a pricing object, and its normalized direction Λ_X is the normalized direction of that priced loading for the chosen claim under $\mathbb{Q}$. Thus Λ_X is not measure-invariant claim by claim. The invariant content is the structural geometry in which those claim-specific loadings live. The aggregate factor benchmark Λ_M^* is the exception isolated here. Its role is tied to the canonical centered market-variance factor construction of [1] rather than to an arbitrary claim. Even there, the $\mathbb{Q}$ -side realization remains measure-dependent through the conditional projection defining that centered structural increment.

Remark 4.9 (What the factor decomposition does and does not classify). Proposition 4.6 is a factor-level decomposition inside the centered structural loading while leaving the loading itself in place. The claim-specific object classified by public visibility remains $L_X^\mathbb{Q} = C_{\mathcal{H}_\varpi}^\mathbb{Q}(X)$, because visible summaries act on the full priced structural loading. The canonical centered aggregate factor Λ_M^* has a separate role. It isolates the aggregate synchronized benchmark direction inside that loading and contributes an exact orthogonal term to compression and hedging residuals. Corollary 4.7 identifies why Λ_M^* is not directly linearly visible through the option-surface sigma-field.

Corollary 4.10 (Compression principle for priced structural loading). *Let $X \in L^2(\mathbb{Q}, \mathcal{F}_T)$ with canonical decomposition*

$$X = L_X^\mathbb{Q} + U_X^\mathbb{Q}.$$

Then:

(i) *for every sub- σ -field $\mathcal{H} \subseteq \mathcal{G}_T$,*

$$C_\mathcal{H}^\mathbb{Q}(X) = C_\mathcal{H}^\mathbb{Q}(L_X^\mathbb{Q});$$

(ii) *for every closed visible hedge class $\mathcal{S} \subseteq L^2(\mathbb{Q}, \mathcal{G}_T)$,*

$$P_\mathcal{S}^\mathbb{Q}(X) = P_\mathcal{S}^\mathbb{Q}(L_X^\mathbb{Q}),$$

where $P_\mathcal{S}^\mathbb{Q}$ is the $L^2(\mathbb{Q})$ -orthogonal projection onto $\mathcal{S}$.

(iii) *if $L_X^\mathbb{Q} \neq 0$, so that $L_X^\mathbb{Q} = \alpha_X^\mathbb{Q} \Lambda_X$, then*

$$C_\mathcal{H}^\mathbb{Q}(X) = \alpha_X^\mathbb{Q} C_\mathcal{H}^\mathbb{Q}(\Lambda_X), \quad P_\mathcal{S}^\mathbb{Q}(X) = \alpha_X^\mathbb{Q} P_\mathcal{S}^\mathbb{Q}(\Lambda_X),$$

for every public $\mathcal{H} \subseteq \mathcal{G}_T$ and every visible hedge class $\mathcal{S} \subseteq L^2(\mathbb{Q}, \mathcal{G}_T)$.

(iv) *more generally, if $\Psi_\mathcal{H}: L^2(\mathbb{Q}, \mathcal{H}) \rightarrow \mathbb{R}^m$ is any Borel measurable functional on a public claim class $L^2(\mathbb{Q}, \mathcal{H})$, then*

$$\Psi_\mathcal{H}(C_\mathcal{H}^\mathbb{Q}(X)) = \Psi_\mathcal{H}(C_\mathcal{H}^\mathbb{Q}(L_X^\mathbb{Q})),$$

and if $L_X^\mathbb{Q} \neq 0$,

$$\Psi_\mathcal{H}(C_\mathcal{H}^\mathbb{Q}(X)) = \Psi_\mathcal{H}(\alpha_X^\mathbb{Q} C_\mathcal{H}^\mathbb{Q}(\Lambda_X)).$$

In particular, every public pricing object treated here is obtained by compressing, projecting, or measurably transforming the same priced structural loading $L_X^\mathbb{Q} = \alpha_X^\mathbb{Q} \Lambda_X$.

Proof. Because $U_X^\mathbb{Q} \perp \mathcal{M}_\mathbb{Q}$ and $L^2(\mathbb{Q}, \mathcal{G}_T) \subseteq \mathcal{M}_\mathbb{Q}$, one has $C_\mathcal{H}^\mathbb{Q}(U_X^\mathbb{Q}) = 0$ for every $\mathcal{H} \subseteq \mathcal{G}_T$, which proves part (i). Likewise $U_X^\mathbb{Q} \perp \mathcal{S}$ for every $\mathcal{S} \subseteq L^2(\mathbb{Q}, \mathcal{G}_T)$, which proves part (ii). Part (iii) is the identity $L_X^\mathbb{Q} = \alpha_X^\mathbb{Q} \Lambda_X$. Part (iv) follows by applying the measurable map $\Psi_\mathcal{H}$ to part (i), and then using part (iii) when $L_X^\mathbb{Q} \neq 0$. $\square$

Corollary 4.11 (Optional same-span reading under bounded-density comparability). *Under Condition 2.3, the geometric span is the same under $\mathbb{P}$ and $\mathbb{Q}$, but public visibility changes because $C_\mathcal{H}^\mathbb{Q}$ and $C_\mathcal{H}^\mathbb{P}$ are different bounded projections on that common span. Thus screening under $\mathbb{Q}$ is a reweighting of the same priced structural directions, not a loss of structural support.*

Proof. Proposition 4.1 identifies the structural span as a common vector space under $\mathbb{P}$ and $\mathbb{Q}$ when the density is bounded above and away from zero. Corollary 4.10 then says that the public pricing object is obtained by applying the risk-neutral projection $C_\mathcal{H}^\mathbb{Q}$ to the priced loading in that span. Thus the support directions are common, while the visible component may change because the conditional-expectation operator is measure-dependent. $\square$

The risk-neutral compression theorem identifies the Hilbert-space source of public pricing, visible projection, and residual hedge risk.

Theorem 4.12 (Main risk-neutral compression theorem). *Assume Assumptions 2.1, 2.2, and 2.4, together with Definition 2.5. Let*

$$\mathcal{H} \subseteq \mathcal{G}_T \subseteq \mathcal{H}_\varpi \subseteq \mathcal{F}_T.$$

For every $X \in L^2(\mathbb{Q}, \mathcal{F}_T)$, set

$$L_X^\mathbb{Q} = C_{\mathcal{H}_\varpi}^\mathbb{Q}(X), \quad U_X^\mathbb{Q} = X - L_X^\mathbb{Q}.$$

Then

$$X = L_X^\mathbb{Q} + U_X^\mathbb{Q}, \quad U_X^\mathbb{Q} \perp_{L^2(\mathbb{Q})} L^2(\mathbb{Q}, \mathcal{H}_\varpi).$$

Every public compression factors through the priced structural loading:

$$C_\mathcal{H}^\mathbb{Q}(X) = C_\mathcal{H}^\mathbb{Q}(L_X^\mathbb{Q}).$$

For every closed visible hedge class $\mathcal{S} \subseteq L^2(\mathbb{Q}, \mathcal{G}_T)$,

$$P_\mathcal{S}^\mathbb{Q}(X) = P_\mathcal{S}^\mathbb{Q}(L_X^\mathbb{Q})$$

and the residual risk decomposes as

$$\inf_{H \in \mathcal{S}} \mathbb{E}_\mathbb{Q}[(X - H)^2] = \text{dist}_{L^2(\mathbb{Q})}(L_X^\mathbb{Q}, \mathcal{S})^2 + \|U_X^\mathbb{Q}\|_{L^2(\mathbb{Q})}^2.$$

If $L_X^\mathbb{Q} \neq 0$, then the normalized loading

$$\Lambda_X = \frac{L_X^\mathbb{Q}}{\|L_X^\mathbb{Q}\|_{L^2(\mathbb{Q})}}$$

determines every public compression and visible projection up to the scalar $\alpha_X^\mathbb{Q} = \|L_X^\mathbb{Q}\|_{L^2(\mathbb{Q})}$.

After centering against $L^2(\mathbb{Q}, \mathcal{G}_T)$, the same loading has the unique aggregate-factor split

$$L_X^\mathbb{Q} = V_X^\mathbb{Q} + \alpha_X^{\mathbb{Q},*} \Lambda_M^* + R_X^{\mathbb{Q},*}, \quad R_X^{\mathbb{Q},*} \perp_{L^2(\mathbb{Q})} \Lambda_M^*.$$

Finally, whenever $X = X_\varpi + R_X$ with X_ϖ $\mathcal{H}_\varpi$ -measurable, replacing the full aggregate structural content by the Perron-generated component X_ϖ perturbs structural projection, public compression, and visible projection by at most $\|R_X\|_{L^2(\mathbb{Q})}$.

Proof. Conditional expectation onto $\mathcal{H}_\varpi$ is the $L^2(\mathbb{Q})$ -orthogonal projection onto the closed subspace $L^2(\mathbb{Q}, \mathcal{H}_\varpi)$. Hence $U_X^\mathbb{Q}$ is orthogonal to $L^2(\mathbb{Q}, \mathcal{H}_\varpi)$. Since $\mathcal{H} \subseteq \mathcal{G}_T \subseteq \mathcal{H}_\varpi$, the tower property gives

$$C_\mathcal{H}^\mathbb{Q}(X) = C_\mathcal{H}^\mathbb{Q}(C_{\mathcal{H}_\varpi}^\mathbb{Q}(X)) = C_\mathcal{H}^\mathbb{Q}(L_X^\mathbb{Q}).$$

Also $\mathcal{S} \subseteq L^2(\mathbb{Q}, \mathcal{G}_T) \subseteq L^2(\mathbb{Q}, \mathcal{H}_\varpi)$, so $U_X^\mathbb{Q} \perp \mathcal{S}$. Therefore the closest element of $\mathcal{S}$ to $X = L_X^\mathbb{Q} + U_X^\mathbb{Q}$ is the closest element of $\mathcal{S}$ to $L_X^\mathbb{Q}$, and Pythagoras gives

$$\|X - H\|_{L^2(\mathbb{Q})}^2 = \|L_X^\mathbb{Q} - H\|_{L^2(\mathbb{Q})}^2 + \|U_X^\mathbb{Q}\|_{L^2(\mathbb{Q})}^2, \quad H \in \mathcal{S}.$$

Taking the infimum over $H \in \mathcal{S}$ proves the hedge residual identity. The normalized-loading statement follows from $L_X^\mathbb{Q} = \alpha_X^\mathbb{Q} \Lambda_X$.

The aggregate-factor split is the orthogonal decomposition of $L_X^\mathbb{Q} - C_{\mathcal{G}_T}^\mathbb{Q}(L_X^\mathbb{Q})$ along the one-dimensional span of Λ_M^* , as in Proposition 4.6. If $X = X_\varpi + R_X$ with X_ϖ $\mathcal{H}_\varpi$ -measurable, then

$$L_X^\mathbb{Q} - X_\varpi = C_{\mathcal{H}_\varpi}^\mathbb{Q}(R_X),$$

so the final perturbative bound follows from the $L^2(\mathbb{Q})$ -contraction property of conditional expectation and orthogonal projection. $\square$

Remark 4.13 (Relation to partial-information hedging). The Hilbert-space step in Theorem 4.12 is the familiar projection step from reduced-information pricing and hedging. Conditional expectation onto a smaller information set is an L^2 -projection, and closed visible hedge spaces enter by orthogonal projection. This is the same geometric layer underlying local risk minimization, mean-variance hedging, the Föllmer–Schweizer decomposition, and its relationship with the Galtchouk–Kunita–Watanabe decomposition [10, 11, 12, 13, 14]. The theorem leaves the Galtchouk–Kunita–Watanabe decomposition and the Föllmer–Schweizer partial-information hedging construction in place. Its role is to identify the object projected here: the $\mathbb{Q}$ -priced structural loading $L_X^\mathbb{Q} = C_{\mathcal{H}_\varpi}^\mathbb{Q}(X)$ generated by the Perron-sufficient latent market state. That object is then the common source of the option, premium, variance, and hedging endpoints below.

5. RISK-NEUTRAL VISIBILITY AND SCREENING

Visibility is a norm statement on the compressed loading. A public summary sees a claim only through $C_{\mathcal{H}}^\mathbb{Q}(L_X^\mathbb{Q})$; the uncompressed loading may still be nonzero.

Definition 5.1 (Compression-based $\mathbb{Q}$ -visibility of priced structural loading). Fix a claim $X \in L^2(\mathbb{Q}, \mathcal{F}_T)$ with priced structural loading $L_X^\mathbb{Q}$, and fix a public sub- σ -field $\mathcal{H} \subseteq \mathcal{G}_T$. The claim's structural loading is called

(i) $\mathbb{Q}$ -visible for X at public scale $\mathcal{H}$ if

$$\|C_{\mathcal{H}}^\mathbb{Q}(L_X^\mathbb{Q})\|_{L^2(\mathbb{Q})} > 0;$$

(ii) $\mathbb{Q}$ -screened for X at public scale $\mathcal{H}$ if

$$\|L_X^\mathbb{Q}\|_{L^2(\mathbb{Q})} > 0 \quad \text{and} \quad \|C_{\mathcal{H}}^\mathbb{Q}(L_X^\mathbb{Q})\|_{L^2(\mathbb{Q})} = 0;$$

(iii) structurally absent for X if

$$\|L_X^\mathbb{Q}\|_{L^2(\mathbb{Q})} = 0.$$

Proposition 5.2 (Operational summary criterion). Let $\{X_h\}_{h>0} \subset L^2(\mathbb{Q}, \mathcal{F}_T)$ be a claim family with priced structural loadings $L_{X_h}^\mathbb{Q}$, let $\mathcal{H}_h \subseteq \mathcal{G}_T$ be public sigma-fields, and let $\{\mathfrak{S}_h(\varepsilon)\}_{h>0, |\varepsilon|<\varepsilon_0} \subset L^2(\mathbb{Q}, \mathcal{H}_h)$ be visible summary families. Suppose that for a positive deterministic normalization $a_h > 0$ there are remainders $r_h(\varepsilon) \in L^2(\mathbb{Q})$ such that

$$\frac{\mathfrak{S}_h(\varepsilon) - \mathfrak{S}_h(0)}{\varepsilon} = C_{\mathcal{H}_h}^\mathbb{Q}(L_{X_h}^\mathbb{Q}) + r_h(\varepsilon),$$

and

$$\lim_{h \downarrow 0, \varepsilon \rightarrow 0} a_h^{-1} \|r_h(\varepsilon)\|_{L^2(\mathbb{Q})} = 0.$$

Assume moreover that the normalized compressed priced structural loadings converge:

$$a_h^{-1} C_{\mathcal{H}_h}^\mathbb{Q}(L_{X_h}^\mathbb{Q}) \longrightarrow Y \quad \text{in } L^2(\mathbb{Q}).$$

Then the normalized perturbation has the joint $L^2(\mathbb{Q})$ -limit

$$D_{\text{load}}^{\mathbb{Q}} \mathfrak{S} := \lim_{h \downarrow 0, \varepsilon \rightarrow 0} \frac{\mathfrak{S}_h(\varepsilon) - \mathfrak{S}_h(0)}{\varepsilon a_h} = Y,$$

and consequently

$$D_{\text{load}}^{\mathbb{Q}} \mathfrak{S} \neq 0 \iff \lim_{h \downarrow 0} a_h^{-1} \|C_{\mathcal{H}_h}^{\mathbb{Q}}(L_{X_h}^{\mathbb{Q}})\|_{L^2(\mathbb{Q})} = \|Y\|_{L^2(\mathbb{Q})} > 0.$$

Thus the derivative-based criterion detects the limiting normalized compressed loading. The normalization a_h need not vanish; this convention covers both decaying and exploding short-maturity quote scales. Without convergence of the normalized compressed loadings it should be read only as a scale-wise sufficient test.

Proof. After division by a_h , the difference between the normalized summary perturbation and $a_h^{-1} C_{\mathcal{H}_h}^{\mathbb{Q}}(L_{X_h}^{\mathbb{Q}})$ is $a_h^{-1} r_h(\varepsilon)$, which converges to zero in the stated joint limit. The normalized compressed loading converges to Y in $L^2(\mathbb{Q})$. The normalized summary perturbation therefore has the same limit. Norm convergence follows from the reverse triangle inequality, and the final equivalence is exactly $\|Y\|_{L^2(\mathbb{Q})} > 0$. $\square$

Proposition 5.3 (Risk-neutral visibility classification for priced structural loadings). *Fix a claim $X \in L^2(\mathbb{Q}, \mathcal{F}_T)$, let*

$$X = L_X^{\mathbb{Q}} + U_X^{\mathbb{Q}}$$

be the canonical decomposition of Proposition 4.2, and fix a public sub- σ -field $\mathcal{H} \subseteq \mathcal{G}_T$. Then exactly one of the following three cases holds:

(i) visible and structurally priced:

$$\|C_{\mathcal{H}}^{\mathbb{Q}}(L_X^{\mathbb{Q}})\|_{L^2(\mathbb{Q})} > 0;$$

(ii) screened but structurally priced:

$$\|L_X^{\mathbb{Q}}\|_{L^2(\mathbb{Q})} > 0 \quad \text{and} \quad \|C_{\mathcal{H}}^{\mathbb{Q}}(L_X^{\mathbb{Q}})\|_{L^2(\mathbb{Q})} = 0;$$

(iii) structurally absent for the claim:

$$\|L_X^{\mathbb{Q}}\|_{L^2(\mathbb{Q})} = 0.$$

These conditions are necessary and sufficient.

$$\text{If } \alpha_X^{\mathbb{Q}} > 0, \quad \|C_{\mathcal{H}}^{\mathbb{Q}}(L_X^{\mathbb{Q}})\|_{L^2(\mathbb{Q})} = \alpha_X^{\mathbb{Q}} \|C_{\mathcal{H}}^{\mathbb{Q}}(\Lambda_X)\|_{L^2(\mathbb{Q})}.$$

Proof. Definition 5.1 partitions the claim by the two norms $\|L_X^{\mathbb{Q}}\|_{L^2(\mathbb{Q})}$ and $\|C_{\mathcal{H}}^{\mathbb{Q}}(L_X^{\mathbb{Q}})\|_{L^2(\mathbb{Q})}$. If the first norm is zero, the loading is absent. If it is positive, the second norm is either positive or zero, giving visibility or screening and no fourth case. When $\alpha_X^{\mathbb{Q}} > 0$, the scalar factorization $L_X^{\mathbb{Q}} = \alpha_X^{\mathbb{Q}} \Lambda_X$ and linearity of $C_{\mathcal{H}}^{\mathbb{Q}}$ give the displayed identity. $\square$

Remark 5.4 (Why this is the classification theorem). The classification does not test existence of the structural geometry; Theorem 3.1 has already fixed that geometry. It tests whether a chosen public compression selects a nonzero component of the claim's priced structural loading. Proposition 5.3 characterizes the distinction by the norm of $C_{\mathcal{H}}^{\mathbb{Q}}(L_X^{\mathbb{Q}})$.

6. MARKET-INDEX PRICE DYNAMICS AND OPTION SETUP UNDER $\mathbb{Q}$

The option block separates two maps. Risk-neutral compression produces the option-facing priced loading; a quote functional decides which part of that loading appears as an observed surface coefficient. The abstract first-variation theorem is Gaussian-free. The cumulant-class condition then turns the compressed scalar signal into a short-maturity skew law, and the Gaussian-Volterra/Bachelier specialization gives kernel-level coefficients. The quote conventions and option-surface background come from the Bachelier, Black-Scholes/Black, rational option-pricing, state-price-density, local-volatility, and skew/moment literatures [31, 32, 33, 7, 34, 35, 36, 37, 38].

Assumption 6.1 (Regular public option-summary first variation). For every sufficiently small maturity $\tau > 0$, there exist:

- (i) a public option sigma-field $\mathcal{H}_{M,\tau}^{\text{opt}} \subseteq \mathcal{G}_T$;
- (ii) an option first-variation pricing object $\Xi_{M,\tau} \in L^2(\mathbb{Q}, \mathcal{F}_T)$;
- (iii) a continuous linear functional

$$\mathfrak{D}_{M,\tau}: L^2(\mathbb{Q}, \mathcal{H}_{M,\tau}^{\text{opt}}) \rightarrow \mathbb{R};$$

- (iv) a deterministic scale $a_{M,\tau} > 0$ and a remainder $r_{M,\tau} \in \mathbb{R}$
- such that the quoted option summary family $\mathfrak{S}_M(\tau; \varepsilon)$ satisfies

$$\partial_\varepsilon \mathfrak{S}_M(\tau; \varepsilon)|_{\varepsilon=0} = \mathfrak{D}_{M,\tau}(C_{\mathcal{H}_{M,\tau}^{\text{opt}}}^{\mathbb{Q}}(\Xi_{M,\tau})) + r_{M,\tau},$$

with

$$a_{M,\tau}^{-1} r_{M,\tau} \rightarrow 0 \quad \text{as } \tau \downarrow 0.$$

Definition 6.2 (Option first-variation pricing object and option-facing priced structural loading). Under Assumption 6.1, define

$$L_{M,\tau}^{\mathbb{Q}} := C_{\mathcal{H}_\omega}^{\mathbb{Q}}(\Xi_{M,\tau}), \quad U_{M,\tau}^{\mathbb{Q}} := \Xi_{M,\tau} - C_{\mathcal{H}_\omega}^{\mathbb{Q}}(\Xi_{M,\tau}).$$

Under the standing hypotheses of Proposition 4.2, this is the canonical priced-structural decomposition of the option first-variation claim $\Xi_{M,\tau}$. If $L_{M,\tau}^{\mathbb{Q}} \neq 0$, set

$$\alpha_{M,\tau}^{\mathbb{Q}} := \|L_{M,\tau}^{\mathbb{Q}}\|_{L^2(\mathbb{Q})}, \quad \Lambda_{M,\tau} := \frac{L_{M,\tau}^{\mathbb{Q}}}{\alpha_{M,\tau}^{\mathbb{Q}}};$$

otherwise set $\alpha_{M,\tau}^{\mathbb{Q}} := 0$ and $\Lambda_{M,\tau} := 0$. The family $\{L_{M,\tau}^{\mathbb{Q}}\}_{\tau>0}$ is the option-facing priced structural loading.

The abstract option theorem separates the compressed loading from the scalar quote functional that may annihilate part of it.

Theorem 6.3 (Abstract option-summary compression of priced structural loading). *Under Assumptions 2.1, 2.2, and 6.1, together with Definitions 2.8 and 6.2,*

$$\partial_\varepsilon \mathfrak{S}_M(\tau; \varepsilon)|_{\varepsilon=0} = \mathfrak{D}_{M,\tau}(C_{\mathcal{H}_{M,\tau}^{\text{opt}}}^{\mathbb{Q}}(L_{M,\tau}^{\mathbb{Q}})) + r_{M,\tau}.$$

If $\alpha_{M,\tau}^{\mathbb{Q}} > 0$, then

$$\partial_\varepsilon \mathfrak{S}_M(\tau; \varepsilon)|_{\varepsilon=0} = \mathfrak{D}_{M,\tau}(\alpha_{M,\tau}^{\mathbb{Q}} C_{\mathcal{H}_{M,\tau}^{\text{opt}}}^{\mathbb{Q}}(\Lambda_{M,\tau})) + r_{M,\tau}.$$

In particular, the first variation of every regular public option summary treated here is a continuous linear functional of the compressed option-facing priced structural loading. Hence the norm of the compressed option-facing priced structural loading governs structural visibility and screening at the chosen option scale:

$$\|C_{\mathcal{H}_{M,\tau}^{\text{opt}}}^{\mathbb{Q}}(L_{M,\tau}^{\mathbb{Q}})\|_{L^2(\mathbb{Q})}.$$

The particular scalar option summary detects that visible loading only through the additional linear functional $\mathfrak{D}_{M,\tau}$; a nonzero compressed loading may be invisible to this scalar summary if it lies in $\ker \mathfrak{D}_{M,\tau}$.

Proof. By Corollary 4.10,

$$C_{\mathcal{H}_{M,\tau}^{\text{opt}}}^{\mathbb{Q}}(\Xi_{M,\tau}) = C_{\mathcal{H}_{M,\tau}^{\text{opt}}}^{\mathbb{Q}}(L_{M,\tau}^{\mathbb{Q}}).$$

Substituting this identity into Assumption 6.1 gives the first displayed formula. The second formula follows from the scalar factorization $L_{M,\tau}^{\mathbb{Q}} = \alpha_{M,\tau}^{\mathbb{Q}} \Lambda_{M,\tau}$ when $\alpha_{M,\tau}^{\mathbb{Q}} > 0$. The final statement separates two kernels. Visibility is governed by the kernel of the projection $C_{\mathcal{H}_{M,\tau}^{\text{opt}}}^{\mathbb{Q}}$ on the option-facing loading. Detection by the chosen scalar quote is governed, after that projection, by $\ker \mathfrak{D}_{M,\tau}$. $\square$

Remark 6.4 (Scope of the abstract option theorem). Theorem 6.3 identifies the option-side compression map. The first variation of a regular public option summary is generated by compressing the option-facing priced structural loading and then applying a continuous linear summary functional. No Gaussian assumption enters that result. The explicit non-Gaussian cumulant-class law and the later Gaussian-Volterra/Bachelier asymptotics are specializations of the same abstract compression statement.

Assumption 6.5 (Non-Gaussian cumulant-class public skew condition). For every sufficiently small maturity $\tau > 0$, let

$$v_M(\tau) := \text{Var}_{\mathbb{Q}}(R_\tau^0)$$

denote the baseline option-facing return variance at maturity τ , and let

$$\kappa_{3,M}^{\mathbb{Q}}(\tau; \varepsilon)$$

denote the third cumulant of the perturbed option-facing return at that maturity. Assume:

- (i) there exists $\vartheta_M > 0$ such that

$$v_M(\tau) = \vartheta_M \tau + o(\tau), \quad \tau \downarrow 0;$$

- (ii) there exists a finite set of exponents $\mathcal{B}_M \subset (0, \infty)$ and coefficients $\{A_M(\beta)\}_{\beta \in \mathcal{B}_M} \subset \mathbb{R}$ such that

$$\text{supp}_M^{\text{comp}} := \{\beta \in \mathcal{B}_M : A_M(\beta) \neq 0\}$$

and, whenever $\text{supp}_M^{\text{comp}} \neq \emptyset$,

$$\beta_M^{\text{comp}} := \inf \text{supp}_M^{\text{comp}},$$

with the convention $\beta_M^{\text{comp}} := +\infty$ when $\text{supp}_M^{\text{comp}} = \emptyset$, and such that:

- (a) if $\text{supp}_M^{\text{comp}} \neq \emptyset$, then

$$\partial_\varepsilon \kappa_{3,M}^{\mathbb{Q}}(\tau; \varepsilon)|_{\varepsilon=0} = 6\vartheta_M^{3/2} \sum_{\beta \in \mathcal{B}_M} A_M(\beta) \tau^{\beta+3/2} + o\left(\tau^{\beta_M^{\text{comp}}+3/2}\right), \quad \tau \downarrow 0,$$

- (b) if $\text{supp}_M^{\text{comp}} = \emptyset$, then

$$\partial_\varepsilon \kappa_{3,M}^{\mathbb{Q}}(\tau; \varepsilon)|_{\varepsilon=0} = 0 \quad \text{for every sufficiently small } \tau > 0.$$

- (iii) there exists a deterministic function $G_M(\tau) \rightarrow G_M(0) > 0$ such that

- (a) if $\text{supp}_M^{\text{comp}} \neq \emptyset$, then

$$\partial_\varepsilon \mathfrak{S}_M(\tau; \varepsilon)|_{\varepsilon=0} = \frac{G_M(\tau)}{6v_M(\tau)^{3/2}\tau^{1/2}} \partial_\varepsilon \kappa_{3,M}^{\mathbb{Q}}(\tau; \varepsilon)|_{\varepsilon=0} + o\left(\tau^{\beta_M^{\text{comp}}-1/2}\right).$$

- (b) if $\text{supp}_M^{\text{comp}} = \emptyset$, then

$$\partial_\varepsilon \mathfrak{S}_M(\tau; \varepsilon)|_{\varepsilon=0} = 0 \quad \text{for every sufficiently small } \tau > 0.$$

Definition 6.6 (Cumulant-selected compressed exponent support of the scalar public signal). Under Assumption 6.5, define the cumulant-selected compressed exponent support by

$$\text{supp}^{\text{comp}}(\Gamma_{M,\tau}^{\mathbb{Q}}) := \{\beta \in \mathcal{B}_M : A_M(\beta) \neq 0\},$$

where

$$\Gamma_{M,\tau}^{\mathbb{Q}} := \mathfrak{D}_{M,\tau}(C_{\mathcal{H}_{M,\tau}^{\text{opt}}}^{\mathbb{Q}}(L_{M,\tau}^{\mathbb{Q}}))$$

is the scalar public first-variation signal generated from the compressed option-facing priced structural loading. The effective public leading exponent is

$$\beta_M^{\text{comp}} := \inf \text{supp}^{\text{comp}}(\Gamma_{M,\tau}^{\mathbb{Q}}),$$

with the convention $+\infty$ when the compressed support is empty. The exponent support is selected by the cumulant-class condition for the scalar public signal generated by the compressed loading. The full L^2 -valued compressed loading has no such scalar support notion.

The cumulant-class theorem characterizes when the scalar public skew coefficient is the leading power selected by the compressed option-facing loading.

Theorem 6.7 (Non-Gaussian cumulant-class market skew as compressed priced structural loading). *Assume Assumptions 2.1, 2.2, 6.1, and 6.5, together with Definitions 2.8 and 6.2. Let*

$$\Gamma_{M,\tau}^{\mathbb{Q}} := \mathfrak{D}_{M,\tau}(C_{\mathcal{H}_{M,\tau}^{\text{opt}}}^{\mathbb{Q}}(L_{M,\tau}^{\mathbb{Q}})).$$

If $\text{supp}^{\text{comp}}(\Gamma_{M,\tau}^{\mathbb{Q}}) = \emptyset$, then the cumulant-class coefficient system does not select a nontrivial leading public power at the modeled scales. Assume henceforth that

$$\text{supp}^{\text{comp}}(\Gamma_{M,\tau}^{\mathbb{Q}}) \neq \emptyset, \quad \text{equivalently } \beta_M^{\text{comp}} < \infty.$$

Then, as $\tau \downarrow 0$, the quoted public first variation satisfies

$$\partial_{\varepsilon} \mathfrak{S}_M(\tau; \varepsilon)|_{\varepsilon=0} = G_M(\tau) \sum_{\beta \in \text{supp}^{\text{comp}}(\Gamma_{M,\tau}^{\mathbb{Q}})} A_M(\beta) \tau^{\beta-1/2} + o(\tau^{\beta_M^{\text{comp}}-1/2}).$$

Moreover Theorem 6.3 gives

$$\partial_{\varepsilon} \mathfrak{S}_M(\tau; \varepsilon)|_{\varepsilon=0} = \Gamma_{M,\tau}^{\mathbb{Q}} + r_{M,\tau}, \quad a_{M,\tau}^{-1} r_{M,\tau} \rightarrow 0.$$

If the abstract option-summary remainder is negligible at the cumulant leading scale,

$$r_{M,\tau} = o(\tau^{\beta_M^{\text{comp}}-1/2}),$$

then the scalar public compressed-loading signal itself has the same expansion:

$$\Gamma_{M,\tau}^{\mathbb{Q}} = G_M(\tau) \sum_{\beta \in \text{supp}^{\text{comp}}(\Gamma_{M,\tau}^{\mathbb{Q}})} A_M(\beta) \tau^{\beta-1/2} + o(\tau^{\beta_M^{\text{comp}}-1/2}).$$

Equivalently, under this scale-compatible remainder condition and when $\alpha_{M,\tau}^{\mathbb{Q}} > 0$,

$$\mathfrak{D}_{M,\tau}(\alpha_{M,\tau}^{\mathbb{Q}} C_{\mathcal{H}_{M,\tau}^{\text{opt}}}^{\mathbb{Q}}(\Lambda_{M,\tau})) = G_M(\tau) \sum_{\beta \in \text{supp}^{\text{comp}}(\Gamma_{M,\tau}^{\mathbb{Q}})} A_M(\beta) \tau^{\beta-1/2} + o(\tau^{\beta_M^{\text{comp}}-1/2}).$$

Thus the leading public skew derivative is generated by the first-order public compression of the option-facing priced structural loading. The expansion for the scalar compressed-loading signal requires the displayed scale compatibility; otherwise the cumulant-class expansion is an expansion of the quoted derivative rather than of $\Gamma_{M,\tau}^{\mathbb{Q}}$ itself.

Proof. Assume $\text{supp}^{\text{comp}}(\Gamma_{M,\tau}^{\mathbb{Q}}) \neq \emptyset$. The case of empty compressed support is exactly the qualitative first sentence of the theorem. Assumption 6.5(iii) expresses the public first variation through the normalized third-cumulant derivative. Substituting the asymptotic expansion from Assumption 6.5(ii) and using

$$v_M(\tau)^{3/2} = \vartheta_M^{3/2} \tau^{3/2} (1 + o(1))$$

from Assumption 6.5(i) yields

$$\partial_{\varepsilon} \mathfrak{S}_M(\tau; \varepsilon)|_{\varepsilon=0} = G_M(\tau) \sum_{\beta \in \mathcal{B}_M} A_M(\beta) \tau^{\beta-1/2} + o(\tau^{\beta_M^{\text{comp}}-1/2}).$$

Removing the zero coefficients is exactly the passage from $\mathcal{B}_M$ to $\text{supp}^{\text{comp}}(\Gamma_{M,\tau}^{\mathbb{Q}})$. The identity with $\Gamma_{M,\tau}^{\mathbb{Q}} + r_{M,\tau}$ is Theorem 6.3, written in scalar notation. If $r_{M,\tau} = o(\tau^{\beta_M^{\text{comp}}-1/2})$, subtracting this remainder from the derivative expansion gives the displayed expansion for $\Gamma_{M,\tau}^{\mathbb{Q}}$. The factorized form follows from Theorem 6.3 when $\alpha_{M,\tau}^{\mathbb{Q}} > 0$. $\square$

The Hermite obstruction, primitive non-Gaussian Volterra estimates, and rare-jump exact price calculations delimit the option-transfer endpoint. They are collected in Appendix A. The article uses the compact consequence in Theorem 6.7. When a quoted surface coefficient satisfies the stated transfer condition, it is a public compression of the option-facing priced structural loading.

Assumption 6.8 (Gaussian-Volterra market-index dynamics with synchronized perturbation). Fix deterministic index weights $b = (b_1, \dots, b_N) \in (0, \infty)^N$, a continuous baseline normal volatility $\sigma_0^M: [0, T] \rightarrow (0, \infty)$, and an additional $\mathbb{Q}$ -Brownian motion $B^\mathbb{Q}$ such that

$$\langle B^\mathbb{Q}, W^{\mathbb{Q},j} \rangle_t = \rho_j t, \quad \sum_{j=1}^N \rho_j^2 \leq 1, \quad 1 \leq j \leq N.$$

Equivalently, possibly after enlarging the filtered space,

$$B_t^\mathbb{Q} = \sum_{j=1}^N \rho_j W_t^{\mathbb{Q},j} + \sqrt{1 - \|\rho\|^2} W_t^{\mathbb{Q},0},$$

where $W^{\mathbb{Q},0}$ is a $\mathbb{Q}$ -Brownian motion independent of $W^\mathbb{Q}$. For a local synchronized perturbation parameter ε , let the discounted market-index price satisfy the additive dynamics

$$S_t^{M,\varepsilon} = S_0^M + \int_0^t \sigma_M^\varepsilon(u) dB_u^\mathbb{Q}, \quad \sigma_M^\varepsilon(u) := \sigma_0^M(u) + \varepsilon Y_M(u),$$

where

$$Y_M(t) := \sum_{j=1}^N \int_0^t G_j(t, s) dW_s^{\mathbb{Q},j},$$

with

$$G_j(t, s) := \sum_{i=1}^N b_i \dot{\beta}_{ij} a_{ij}(t, s) (t - s)^{H_{ij}-1/2} \mathbf{1}_{\{(i,j) \in \mathcal{E}_M\}}.$$

Here $\dot{\beta}_{ij} \in \mathbb{R}$ is the synchronized-mode perturbation loading of the market-visible channel (i, j) . The synchronized perturbation satisfies

$$\mathbb{E}_\mathbb{Q} \int_0^T Y_M(u)^2 du < \infty,$$

so the stochastic integral defining Z_τ below is well defined for $\tau \leq T$. Set

$$X_\tau := \int_0^\tau \sigma_0^M(u) dB_u^\mathbb{Q}, \quad Z_\tau := \int_0^\tau Y_M(u) dB_u^\mathbb{Q}, \quad R_\tau^\varepsilon := S_\tau^{M,\varepsilon} - S_0^M = X_\tau + \varepsilon Z_\tau.$$

Assume further that for every sufficiently small $\tau > 0$,

$$X_\tau^2 Z_\tau \in L^2(\mathbb{Q}).$$

For $\tau > 0$, let $C_M(\tau, k; \varepsilon)$ denote the local date- t price, with t suppressed, of a maturity- τ call on $S^{M,\varepsilon}$ with strike $S_t^M + k$, let $\sigma_N^M(\tau, k; \varepsilon)$ be its Bachelier implied volatility, let $\mathcal{H}_{M,\tau}^{\text{opt}} \subseteq \mathcal{G}_t$ denote the public sigma-field generated by the maturity- τ quote family at that observation date, and let

$$\mathfrak{S}_{M,N}(\tau; \varepsilon) \in \mathbb{R} \subset L^2(\mathbb{Q}, \mathcal{H}_{M,\tau}^{\text{opt}})$$

denote the associated public ATM normal skew summary. The displayed local formulas are written after conditioning on the observation-date public information. In a single-date time-zero calculation this sigma-field may be trivial and the quote is identified with a constant element of $L^2(\mathbb{Q})$; in the empirical panel, the same coefficient is observed repeatedly across dates as a $\mathcal{G}_t$ -measurable option-surface coordinate. Its local value is the quoted normal slope

$$\partial_k \sigma_N^M(\tau, k; \varepsilon)|_{k=0}.$$

Definition 6.9 (Market-visible option coefficients and exponents). Under Assumption 6.8, define the return-visible channel set

$$\mathcal{E}_M^{\text{ret}} := \{(i, j) \in \mathcal{E}_M : b_i \dot{\beta}_{ij} \rho_j \neq 0\}.$$

For each $(i, j) \in \mathcal{E}_M^{\text{ret}}$, set

$$\rho_{ij}^M := \frac{b_i \dot{\beta}_{ij} a_{ij}(0, 0) \rho_j}{H_{ij} + 1/2}, \quad \kappa_{ij}^M := \frac{6 (\sigma_0^M(0))^2 b_i \dot{\beta}_{ij} a_{ij}(0, 0) \rho_j}{(H_{ij} + 1/2)(H_{ij} + 3/2)},$$

and define

$$c_{ij}^M := \frac{\kappa_{ij}^M}{6(\sigma_0^M(0))^3} = \frac{b_i \dot{\beta}_{ij} a_{ij}(0,0) \rho_j}{\sigma_0^M(0)(H_{ij} + 1/2)(H_{ij} + 3/2)}.$$

Set the market-scale leading exponents

$$\beta_M := \min\{H_{ij} : (i, j) \in \mathcal{E}_M^{\text{ret}}\},$$

$$\beta_M^{\text{off}} := \min\{H_{ij} : i \neq j, (i, j) \in \mathcal{E}_M^{\text{ret}}\}, \quad \beta_M^{\text{diag}} := \min\{H_{ii} : (i, i) \in \mathcal{E}_M^{\text{ret}}\},$$

with the convention $+\infty$ when the corresponding set is empty; the same convention also applies to β_M . Define the aggregate leading coefficients

$$A_M^{\text{off}} := \sum_{i \neq j: H_{ij} = \beta_M^{\text{off}}} c_{ij}^M, \quad A_M^{\text{diag}} := \sum_{i: H_{ii} = \beta_M^{\text{diag}}} c_{ii}^M.$$

When $\mathcal{E}_M^{\text{ret}} = \emptyset$, all coefficients above are taken to vanish and the skew response is trivial; in particular $\beta_M = +\infty$.

Remark 6.10 (Role of the market-index setup). Assumption 6.8 supplies the market-facing input for the option calculation. It fixes the pricing measure, return Brownian channel, leverage coefficients, and synchronized perturbation map. The option summary is generated by this market-index model, not by an external single-channel pricing specification. In the Gaussian–Volterra/Bachelier specialization, the cumulant-class condition is verified in closed form.

7. OPTION-SURFACE TRANSFER: MAIN STATEMENT

The empirical option input is deliberately narrow. Section 6 supplies the abstract first-variation theorem, the cumulant-class transfer, and the local option-facing loading. Appendices D–E supply the Gaussian–Volterra, Bachelier, and Black calculations. Under the source-family hypotheses verified in Companion A, the local option-surface coefficient is a public $\mathbb{Q}$ -compression of the same priced structural loading that appears in Theorem 4.12. The SPX/NDX exercise estimates that finite projection by aligning option-surface level, total-variance level, and downside total-variance skew with the SPX/NDX restriction of the full Oxford–Man Perron direction.

8. EMPIRICAL OPTION-SURFACE EVIDENCE

The empirical exercise is a finite public-projection diagnostic. The theory treats a local pricing problem under a fixed $\mathbb{Q}$. After conditioning on observation-date public information, a single time-zero quote is a deterministic element of that local price system. The panel uses the date-indexed analogue: on date d , the fitted option-surface coordinate is the realized value of the same local public-summary coefficient map. The regressions test the panel analogue of the compression coefficient, not randomness of one time-zero quote.

The Perron state is rebuilt in a separate reproducibility archive from raw Oxford–Man realized-volatility data and then matched to SPX/NDX OptionMetrics surfaces over the common window 2000–2018. The $\mathbb{P}$ -side object is the latent Perron stress direction rebuilt on the full eight-index Oxford–Man universe

$$\{\text{SPX}, \text{NASDAQ}, \text{FTSE}, \text{DAX}, \text{CAC}, \text{EuroStoxx 50}, \text{Nikkei}, \text{Hang Seng}\}.$$

The $\mathbb{Q}$ -side object is the SPX/NDX option-surface vector projected onto the SPX/NDX restriction of the full Perron direction, not onto a Perron vector recomputed from the two-index submatrix. The test asks whether public option-surface coordinates align with the already rebuilt $\mathbb{P}$ -side direction, rather than calibrating the Perron state to options.

For feature f , maturity bucket τ , and observation date d , the panel regression is

$$q_{d,\tau}^{(f)} = \alpha_\tau^{(f)} + \beta_\tau^{(f)} p_d + \varepsilon_{d,\tau}^{(f)}.$$

Here p_d is the rebuilt $\mathbb{P}$ -side Perron coordinate and $q_{d,\tau}^{(f)}$ is the $\mathbb{Q}$ -side option-feature vector, centered ticker-by-ticker and projected onto the same restricted full-universe Perron direction. The reported slopes use the raw feature units. Correlations, sign shares, and standardized

signed compression-gap diagnostics use separately standardized p_d and $q_{d,\tau}^{(f)}$ within the selected feature/state-mode/maturity cell. Newey–West standard errors use the automatic lag rule in the empirical code,

$$\ell = \max\{1, \min(20, \text{round}(4(n/100)^{2/9}))\}.$$

Regressions are estimated separately by maturity bucket; the reported HAC errors are time-series Newey–West errors within each bucket rather than cross-sectional date-clustered errors. Rolling midpoint rows use one observation per rolling Perron window, while rolling completed-window rows match each option date to the most recent completed rolling state. For each matched date, let $Q_{d,\tau}^{(f)} \in \mathbb{R}^2$ be the centered SPX/NDX option-feature vector and write its Perron coordinate as

$$q_{d,\tau}^{(f)} = \langle \hat{u}_{\text{opt}}, Q_{d,\tau}^{(f)} \rangle, \quad \Pi_{\text{Perron}} Q_{d,\tau}^{(f)} = q_{d,\tau}^{(f)} \hat{v}_{\text{opt}},$$

where $\hat{u}_{\text{opt}}$ and $\hat{v}_{\text{opt}}$ are the SPX/NDX restrictions of the full left and right Perron vectors, normalized by $\|\hat{v}_{\text{opt}}\|_2 = 1$ and $\langle \hat{u}_{\text{opt}}, \hat{v}_{\text{opt}} \rangle = 1$. With $a_{\text{opt}} = (1/2, 1/2)$, the off-Perron aggregate residual is

$$o_{d,\tau}^{(f)} = a_{\text{opt}}^\top (Q_{d,\tau}^{(f)} - q_{d,\tau}^{(f)} \hat{v}_{\text{opt}}),$$

and the off-Perron statistic reported below is

$$\frac{1}{n_\tau} \sum_{d=1}^{n_\tau} |o_{d,\tau}^{(f)}|,$$

in the raw feature units of f . The statistic is therefore a residual magnitude diagnostic on the two-index restriction rather than a scale-free percentage. Placebo and residualized controls compare the Perron coordinate with the off-Perron coordinate, rotated-date Perron series, median random directions in the Perron/off-Perron plane, and regime/aggregate controls. For the downside total-variance skew coordinate, the residualized checks also control directly for the contemporaneous ATM total-variance level.

The surface coordinates are deliberately low-dimensional. The level coordinates are the fitted ATM implied volatility and ATM total variance at 30, 60, and 90 days. The tail-facing coordinate is the signed downside total-variance skew

$$-\partial_k \{\tau \sigma_B(\tau, k)^2\} \big|_{k=0},$$

estimated by a filtered local fit to the total-variance smile. Raw implied-volatility skew is retained as a secondary tail statistic, but the main tail coordinate is total-variance based because it is the surface jet that matches the theorem-side scaling most directly. The ingestion keeps European index-option quotes with positive bid, offer, mid, implied volatility, and forward, valid expiry dates, $7 \leq \text{DTE} \leq 365$, and strike scaled as `strike_price/1000`. The local fit uses quotes in the maturity buckets $\{30, 60, 90\}$ days with a ten-day half-width and $|k| \leq 0.35$. The certified skew fit requires both SPX and NDX quotes to satisfy local-surface quality conditions: enough quotes on both sides of ATM, near-ATM support, high direct total-variance fit quality, and a well-conditioned local quadratic design.

Table 1 is the descriptive benchmark. ATM implied-volatility level explains roughly two thirds of the Perron-aligned option movement, and ATM total-variance level explains about one half. Filtered downside total-variance skew has lower explanatory power than level, but it remains a Perron-compressed tail coordinate. Raw IV skew is weaker and less stable.

Table 2 separates the descriptive full-sample benchmark from rolling-state evidence. The midpoint column uses one observation per rolling Perron window; the completed-window column matches each option date to the most recent completed rolling state. The level effects fall, but remain statistically visible. The downside total-variance skew retains a nontrivial rolling signal, which is the empirical counterpart of the option-tail branch used below.

The SPX/NDX pair supplies a focused two-index setting for the compression diagnostic. The measured coordinates match the finite-projection pattern. The dominant public option coordinates are compressed versions of the rebuilt latent Perron state, while skew becomes stable only after the signed total-variance surface jet and the local-surface quality filters are imposed.

Feature	Days	Obs.	Corr.	R^2	NW t	Off-Perron
ATM IV level	30	3207	0.807	0.651	23.219	0.058
ATM IV level	60	2903	0.805	0.648	25.040	0.054
ATM IV level	90	1997	0.801	0.641	21.607	0.050
ATM total-variance level	30	3207	0.708	0.501	13.338	0.004
ATM total-variance level	60	2903	0.734	0.539	15.409	0.005
ATM total-variance level	90	1997	0.738	0.545	14.247	0.006
Downside total-variance skew	30	3207	0.685	0.469	10.308	0.009
Downside total-variance skew	60	2903	0.692	0.479	10.642	0.007
Downside total-variance skew	90	1997	0.668	0.446	9.283	0.010

TABLE 1. Full-sample SPX/NDX compression evidence. The $\mathbb{Q}$ -side option-surface Perron coordinate is regressed on the rebuilt $\mathbb{P}$ -side Perron coordinate. The final column reports the mean absolute $\mathbb{Q}$ -side off-Perron component on the same two-index restriction.

Feature	Days	Midpoint R^2	Midpoint t	Completed R^2	Completed t
ATM IV level	30	0.268	4.275	0.274	7.301
ATM IV level	60	0.220	4.022	0.282	7.211
ATM IV level	90	0.255	4.629	0.238	5.687
ATM total variance	30	0.218	3.082	0.232	5.312
ATM total variance	60	0.202	3.389	0.266	5.577
ATM total variance	90	0.248	3.902	0.233	4.674
Downside TV skew	30	0.305	4.116	0.267	6.660
Downside TV skew	60	0.300	4.682	0.324	7.242
Downside TV skew	90	0.330	4.664	0.279	5.948

TABLE 2. Rolling-window compression evidence. The midpoint column uses one observation per rebuilt rolling Perron window; the completed-window column matches each option date to the most recent completed rolling state.

Days	Obs.	Corr.	R^2	NW t	Mean abs. off-Perron
30	3087	0.684	0.468	10.002	0.009
60	2888	0.695	0.483	10.671	0.007
90	1973	0.674	0.455	9.331	0.010

TABLE 3. Filtered downside total-variance skew statistics from the local surface fit.

8.1. Signed premium evidence. The premium branch compares matched $\mathbb{P}$ - and $\mathbb{Q}$ -projected coefficients on a common basis. The empirical implementation is conservative. It reports the raw-unit transfer slope from the regression above and separately records the standardized signed compression gap

$$\Delta_{M,\tau}^z = z(\mathbb{Q}\text{-option Perron coordinate}) - z(\mathbb{P}\text{-Perron stress coordinate}).$$

Because the mean of this standardized signed compression gap is mechanically zero in the full sample, the signed information comes from the raw transfer slope, the share of dates on which the $\mathbb{P}$ - and $\mathbb{Q}$ -coordinates have the same standardized sign, and the difference between the mean gap in the top and bottom $\mathbb{P}$ -stress quintiles.

The signed diagnostic contains information that an absolute-gap proxy discards. The transfer slopes are positive and statistically separated from zero, and same-sign shares are high across the level, total-variance, and filtered downside-skew coordinates. The high-minus-low stress gap is uniformly negative. In high $\mathbb{P}$ -stress states the standardized option coordinate remains aligned with the Perron direction but lies below the standardized physical stress coordinate. The object is a signed compression and shrinkage diagnostic rather than a direct economic premium magnitude.

Table 5 shows that the signed premium comparison survives the same rebuilt rolling-state discipline used in the compression experiment. The rolling statistics are smaller than the

Feature	Days	Raw slope	R^2	NW t	Same sign	High-low gap
ATM IV level	30	0.023	0.651	23.219	0.828	-0.582
ATM IV level	60	0.022	0.648	25.040	0.832	-0.599
ATM IV level	90	0.019	0.641	21.607	0.834	-0.590
ATM total variance	30	0.001	0.501	13.338	0.776	-0.910
ATM total variance	60	0.002	0.539	15.409	0.787	-0.832
ATM total variance	90	0.003	0.545	14.247	0.805	-0.786
Downside TV skew	30	0.002	0.468	10.002	0.770	-1.009
Downside TV skew	60	0.003	0.483	10.671	0.776	-1.008
Downside TV skew	90	0.003	0.455	9.331	0.764	-1.006

TABLE 4. Signed common-basis compression statistics. Slope is reported in raw feature units. The high-low gap is the mean of $\Delta_{M,\tau}^z$ in the top $\mathbb{P}$ -stress quintile minus the corresponding mean in the bottom quintile.

Feature	State	R^2 range	NW t range	Same-sign range	High-low range
ATM IV level	Full sample	0.641–0.651	21.607–25.040	0.828–0.834	-0.599– -0.582
ATM IV level	Rolling daily	0.238–0.282	5.687–7.301	0.600–0.624	-1.546– -1.427
ATM IV level	Rolling midpoint	0.220–0.268	4.022–4.629	0.610–0.657	-1.543– -1.320
ATM total variance	Full sample	0.501–0.545	13.338–15.409	0.776–0.805	-0.910– -0.786
ATM total variance	Rolling daily	0.232–0.266	4.674–5.577	0.604–0.622	-1.654– -1.514
ATM total variance	Rolling midpoint	0.202–0.248	3.082–3.902	0.598–0.635	-1.554– -1.469
Downside TV skew	Full sample	0.455–0.483	9.331–10.671	0.764–0.776	-1.009– -1.006
Downside TV skew	Rolling daily	0.257–0.324	5.952–7.243	0.656–0.686	-1.560– -1.344
Downside TV skew	Rolling midpoint	0.244–0.330	4.225–4.682	0.621–0.697	-1.563– -1.461

TABLE 5. Rolling signed common-basis premium statistics. Ranges are taken over the 30-, 60-, and 90-day maturity buckets. Rolling daily uses the completed-window convention; rolling midpoint uses one observation per rolling Perron window.

full-sample benchmark, but transfer R^2 and same-sign shares remain positive, while the high-minus-low $\mathbb{P}$ -stress gap remains negative across the three option-surface coordinates. The comparison is a targeted two-index signed comparison aligned with the rolling implementation used for the core compression evidence.

9. PAIRWISE LATENT-CONTAGION PREMIUM ON A MATCHED $\mathbb{P}/\mathbb{Q}$ SURFACE

The premium section compares the visible pricing geometry under $\mathbb{Q}$ with the projected physical-measure object supplied by the projected-score construction of [4]. The estimand is a signed coefficient difference. Both sides are projected onto the same source-screened basis, and the $\mathbb{Q}$ -side comes from a cleaned option-implied derivative surface observed on dates matched to the realized-volatility sample. Option-implied surface information is used in the state-price-density and local-surface sense [34, 35]; the realized side follows the Oxford-Man/realized-kernel provenance [28, 29, 30]. The object is the matched $\mathbb{P}/\mathbb{Q}$ projected coefficient, not a standalone roughness comparison.

Definition 9.1 (Harmonized pairwise projected basis and premiums). Fix an ordered pair (i, j) . Let φ_{ij} denote the common unit-norm source-screened projected basis direction supplied by the projected-score geometry of [4] for that pair, normalized identically under $\mathbb{P}$ and $\mathbb{Q}$. Let $D_{ij}^{\mathbb{P}}$ and $D_{ij}^{\mathbb{Q}}$ denote the corresponding local covariance-derivative objects under $\mathbb{P}$ and $\mathbb{Q}$, represented on the same reduced experiment. Let $L_{ij}^{\mathbb{P}}$ and $L_{ij}^{\mathbb{Q}}$ denote the pairwise priced structural loadings of those reduced experiments on the harmonized basis, so that $D_{ij}^{\mathbb{P}}$ and $D_{ij}^{\mathbb{Q}}$ are their public projected representatives. Define the projected pairwise structural-loading coefficients

$$\gamma_{ij}^{\mathbb{P}} := \langle D_{ij}^{\mathbb{P}}, \varphi_{ij} \rangle_{\mathbb{P},ij}, \quad \gamma_{ij}^{\mathbb{Q}} := \langle D_{ij}^{\mathbb{Q}}, \varphi_{ij} \rangle_{\mathbb{Q},ij},$$

the pairwise latent-contagion premium

$$\Pi_{ij} := \gamma_{ij}^{\mathbb{Q}} - \gamma_{ij}^{\mathbb{P}},$$

and the companion reduced-form roughness gap

$$\Delta \tilde{H}_{ij} := \tilde{H}_{ij}^{\mathbb{Q}} - \tilde{H}_{ij}^{\mathbb{P}},$$

where $\tilde{H}_{ij}^{\mathbb{P}}$ and $\tilde{H}_{ij}^{\mathbb{Q}}$ are pairwise reduced-form roughness exponents estimated from the harmonized physical and option-implied experiments, not the structural Volterra exponents of Section 2.

Remark 9.2 (Why the premium is coefficient-based). The primary premium object is Π_{ij} , with $\Delta \tilde{H}_{ij}$ retained only as a secondary roughness summary. A coefficient gap is closer to the projected structural-loading geometry, uses the same source-screened normalization on both sides, and maps directly to the projected-score machinery of [4].

Assumption 9.3 (Matched cleaned derivative-surface experiment). For each ordered pair (i, j) , there exist matched observation dates

$$d = 1, \dots, n,$$

a maturity grid $\mathcal{T}_n \subset (0, \bar{\tau}]$, and a strike or log-moneyness grid $\mathcal{M}_n \subset [-\bar{k}, \bar{k}]$ such that:

- (i) on every matched date d , the $\mathbb{P}$ -side reduced experiment is represented by a projected-score object $\hat{D}_{ij,d}^{\mathbb{P}}$ built from the realized-volatility geometry of [4];
- (ii) on the same matched date d , the $\mathbb{Q}$ -side reduced experiment is represented by a cleaned, arbitrage-regularized option surface

$$\hat{C}_{i,d}(\tau, k) \quad \text{or equivalently} \quad \hat{\sigma}_{i,d}^{\text{imp}}(\tau, k), \quad (\tau, k) \in \mathcal{T}_n \times \mathcal{M}_n,$$

from which a local derivative object $\hat{D}_{ij,d}^{\mathbb{Q}}$ is extracted by a linear surface-functional compatible with the same source-screened nuisance projection and the same pairwise normalization as on the $\mathbb{P}$ -side;

- (iii) both sides are projected onto the same harmonized source-screened basis φ_{ij} , yielding paired coefficient observations

$$\hat{g}_{ij,d}^{\mathbb{P}} := \langle \hat{D}_{ij,d}^{\mathbb{P}}, \varphi_{ij} \rangle_{\mathbb{P},ij}, \quad \hat{g}_{ij,d}^{\mathbb{Q}} := \langle \hat{D}_{ij,d}^{\mathbb{Q}}, \varphi_{ij} \rangle_{\mathbb{Q},ij},$$

and sample averages

$$\hat{\gamma}_{ij,n}^{\mathbb{P}} := \frac{1}{n} \sum_{d=1}^n \hat{g}_{ij,d}^{\mathbb{P}}, \quad \hat{\gamma}_{ij,n}^{\mathbb{Q}} := \frac{1}{n} \sum_{d=1}^n \hat{g}_{ij,d}^{\mathbb{Q}};$$

- (iv) the paired estimator vector satisfies a joint central limit theorem:

$$\sqrt{n} \begin{pmatrix} \hat{\gamma}_{ij,n}^{\mathbb{P}} - \gamma_{ij}^{\mathbb{P}} \\ \hat{\gamma}_{ij,n}^{\mathbb{Q}} - \gamma_{ij}^{\mathbb{Q}} \end{pmatrix} \Rightarrow \mathcal{N} \left(0, \begin{pmatrix} \Sigma_{ij}^{\mathbb{P}} & \Sigma_{ij}^{\mathbb{P}\mathbb{Q}} \\ \Sigma_{ij}^{\mathbb{P}\mathbb{Q}} & \Sigma_{ij}^{\mathbb{Q}} \end{pmatrix} \right).$$

Definition 9.4 (Cleaned option-calibration panel). For each matched date $d = 1, \dots, n$, asset or index i , and quote index $\ell = 1, \dots, L_{i,d}$, observe a mid quote or implied-volatility quote

$$Y_{i,d,\ell}^{\text{obs}}$$

at maturity–moneyness coordinate

$$x_{i,d,\ell} := (\tau_{i,d,\ell}, k_{i,d,\ell}) \in (0, \bar{\tau}] \times [-\bar{k}, \bar{k}].$$

A cleaned $\mathbb{Q}$ -surface is an estimator

$$\hat{C}_{i,d} \in \mathcal{X}$$

obtained from the quote cloud by weighted least squares, likelihood, or constrained smoothing over a no-arbitrage admissible class Θ_{arb} :

$$\hat{\theta}_{i,d} \in \arg \min_{\theta \in \Theta_{\text{arb}}} \sum_{\ell=1}^{L_{i,d}} w_{i,d,\ell} (Y_{i,d,\ell}^{\text{obs}} - C_{\theta}(x_{i,d,\ell}))^2 + \text{pen}_{i,d}(\theta), \quad \hat{C}_{i,d} := C_{\hat{\theta}_{i,d}}.$$

Here $\mathcal{X}$ is a Banach space strong enough to control the local surface derivatives used below, for example a weighted C^r -norm on the maturity–moneyness window.

Definition 9.5 (Option-surface coefficient extraction). Fix an ordered pair (i, j) . Let

$$\mathcal{A}_{ij} : \mathcal{X} \rightarrow \mathcal{D}_{ij}$$

be the local linear derivative map that extracts the $\mathbb{Q}$ -side pairwise derivative object from a cleaned option surface. Examples include local ATM-skew derivatives, local total-variance curvature derivatives, or a vector of maturity-local derivatives. Let

$$P_{ij} : \mathcal{D}_{ij} \rightarrow \mathcal{D}_{ij}^{\text{scr}}$$

be the source-screening nuisance projection, $N_{ij} > 0$ the pairwise normalizer, and $\varphi_{ij} \in \mathcal{D}_{ij}^{\text{scr}}$ the unit harmonized source-screened basis vector. The population $\mathbb{Q}$ -side coefficient observation on date d is

$$g_{ij,d}^{\mathbb{Q}} := \left\langle N_{ij}^{-1} P_{ij} \mathcal{A}_{ij}(\mathbb{C}_{i,d}^{\mathbb{Q}}), \varphi_{ij} \right\rangle_{ij},$$

where $\mathbb{C}_{i,d}^{\mathbb{Q}}$ is the latent arbitrage-free option surface on that date. The empirical coefficient observation is

$$\hat{g}_{ij,d}^{\mathbb{Q}} := \left\langle \hat{N}_{ij,n}^{-1} \hat{P}_{ij,n} \mathcal{A}_{ij}(\hat{\mathbb{C}}_{i,d}), \hat{\varphi}_{ij,n} \right\rangle_{ij,n}.$$

The $\mathbb{P}$ -side coefficient $\hat{g}_{ij,d}^{\mathbb{P}}$ is defined analogously from the realized-volatility projected-score experiment on the same date.

Assumption 9.6 (Derivative-stable calibration and harmonization). For each ordered pair (i, j) , the following hold.

- (i) The derivative extraction map $\mathcal{A}_{ij}$ is bounded linear from $\mathcal{X}$ to $\mathcal{D}_{ij}$, and P_{ij} is a bounded projection.
- (ii) The empirical projection, basis, inner product, and normalizer are stable in the sense that substituting

$$(\hat{P}_{ij,n}, \hat{\varphi}_{ij,n}, \hat{N}_{ij,n}, \langle \cdot, \cdot \rangle_{ij,n})$$

for

$$(P_{ij}, \varphi_{ij}, N_{ij}, \langle \cdot, \cdot \rangle_{ij})$$

changes the sample mean of the $\mathbb{Q}$ -side coefficient observations by $o_{\mathbb{P}}(n^{-1/2})$.

- (iii) The cleaned surface has the coefficient-level asymptotic linear representation

$$\hat{g}_{ij,d}^{\mathbb{Q}} = g_{ij,d}^{\mathbb{Q}} + \zeta_{ij,d}^{\mathbb{Q}} + b_{ij,n}^{\mathbb{Q}} + r_{ij,d,n}^{\mathbb{Q}},$$

where $\zeta_{ij,d}^{\mathbb{Q}}$ is the calibration influence contribution, $b_{ij,n}^{\mathbb{Q}}$ is a deterministic smoothing or regularization bias, and

$$\sqrt{n} b_{ij,n}^{\mathbb{Q}} \rightarrow b_{ij}^{\mathbb{Q}}, \quad \frac{1}{\sqrt{n}} \sum_{d=1}^n r_{ij,d,n}^{\mathbb{Q}} = o_{\mathbb{P}}(1).$$

The same representation holds on the $\mathbb{P}$ -side, with influence term $\zeta_{ij,d}^{\mathbb{P}}$ and bias limit $b_{ij}^{\mathbb{P}}$.

- (iv) The matched vector

$$Y_{ij,d} := \begin{pmatrix} g_{ij,d}^{\mathbb{P}} - \gamma_{ij}^{\mathbb{P}} + \zeta_{ij,d}^{\mathbb{P}} \\ g_{ij,d}^{\mathbb{Q}} - \gamma_{ij}^{\mathbb{Q}} + \zeta_{ij,d}^{\mathbb{Q}} \end{pmatrix}$$

is strictly stationary and satisfies a bivariate central limit theorem

$$\frac{1}{\sqrt{n}} \sum_{d=1}^n Y_{ij,d} \Rightarrow \mathcal{N}(0, \Omega_{ij}^{\text{cal}}).$$

Theorem 9.7 (Calibration-to-premium coefficient transfer). *Under Assumption 9.6, define*

$$\hat{\gamma}_{ij,n}^{\mu} := \frac{1}{n} \sum_{d=1}^n \hat{g}_{ij,d}^{\mu}, \quad \mu \in \{\mathbb{P}, \mathbb{Q}\},$$

and

$$\hat{\Pi}_{ij,n} := \hat{\gamma}_{ij,n}^{\mathbb{Q}} - \hat{\gamma}_{ij,n}^{\mathbb{P}}.$$

Then

$$\sqrt{n} \begin{pmatrix} \hat{\gamma}_{ij,n}^{\mathbb{P}} - \gamma_{ij}^{\mathbb{P}} \\ \hat{\gamma}_{ij,n}^{\mathbb{Q}} - \gamma_{ij}^{\mathbb{Q}} \end{pmatrix} \Rightarrow \mathcal{N} \left(\begin{pmatrix} b_{ij}^{\mathbb{P}} \\ b_{ij}^{\mathbb{Q}} \end{pmatrix}, \Omega_{ij}^{\text{cal}} \right).$$

Consequently,

$$\sqrt{n} (\hat{\Pi}_{ij,n} - \Pi_{ij}) \Rightarrow \mathcal{N} \left(b_{ij}^{\mathbb{Q}} - b_{ij}^{\mathbb{P}}, (-1 \quad 1) \Omega_{ij}^{\text{cal}} \begin{pmatrix} -1 \\ 1 \end{pmatrix} \right).$$

If the regularization biases are undersmoothed or bias-corrected so that $b_{ij}^{\mathbb{P}} = b_{ij}^{\mathbb{Q}} = 0$, this is exactly the matched premium CLT in Proposition 9.10, with the covariance matrix enlarged to include option-calibration and realized-volatility estimation noise.

Proof. The coefficient-level representations imply

$$\sqrt{n} \begin{pmatrix} \hat{\gamma}_{ij,n}^{\mathbb{P}} - \gamma_{ij}^{\mathbb{P}} \\ \hat{\gamma}_{ij,n}^{\mathbb{Q}} - \gamma_{ij}^{\mathbb{Q}} \end{pmatrix} = \frac{1}{\sqrt{n}} \sum_{d=1}^n Y_{ij,d} + \begin{pmatrix} \sqrt{n} b_{ij,n}^{\mathbb{P}} \\ \sqrt{n} b_{ij,n}^{\mathbb{Q}} \end{pmatrix} + o_{\mathbb{P}}(1).$$

The assumed CLT and bias limits give the bivariate limit. The premium estimator is the linear map $(x_P, x_Q) \mapsto x_Q - x_P$, so the second claim follows by the continuous mapping theorem. $\square$

Proposition 9.8 (A simple sufficient rate contract for negligible calibration error). *Suppose the harmonization objects are deterministic or estimated on an auxiliary sample, and there exists a deterministic sequence $a_n \downarrow 0$ such that*

$$\max_{1 \leq d \leq n} \|\hat{\mathbf{C}}_{i,d} - \mathbf{C}_{i,d}^{\mathbb{Q}}\|_{\mathcal{X}} = O_{\mathbb{P}}(a_n), \quad \sqrt{n} a_n \rightarrow 0.$$

Then the calibration contribution to $\hat{\gamma}_{ij,n}^{\mathbb{Q}}$ is $o_{\mathbb{P}}(n^{-1/2})$. In this undersmoothed regime the cleaned surface can be treated as providing the oracle coefficient observations $g_{ij,d}^{\mathbb{Q}}$ at the first-order premium scale.

Proof. By boundedness of $\mathcal{A}_{ij}$, P_{ij} , and the coefficient functional, there is a constant C_{ij} such that

$$|\hat{g}_{ij,d}^{\mathbb{Q}} - g_{ij,d}^{\mathbb{Q}}| \leq C_{ij} \|\hat{\mathbf{C}}_{i,d} - \mathbf{C}_{i,d}^{\mathbb{Q}}\|_{\mathcal{X}}$$

up to the separately controlled harmonization error. Hence

$$\sqrt{n} \left| \frac{1}{n} \sum_{d=1}^n (\hat{g}_{ij,d}^{\mathbb{Q}} - g_{ij,d}^{\mathbb{Q}}) \right| \leq C_{ij} \sqrt{n} \max_d \|\hat{\mathbf{C}}_{i,d} - \mathbf{C}_{i,d}^{\mathbb{Q}}\|_{\mathcal{X}} = o_{\mathbb{P}}(1).$$

$\square$

Remark 9.9 (Operational calibration path for premium coefficients). The premium coefficients are no longer abstract observations. The market data path is quote cloud $\rightarrow$ no-arbitrage cleaned surface $\hat{\mathbf{C}}_{i,d} \rightarrow$ local derivative object $\mathcal{A}_{ij}(\hat{\mathbf{C}}_{i,d}) \rightarrow$ source-screened projection and normalization $\rightarrow$ coefficient observation $\hat{g}_{ij,d}^{\mathbb{Q}} \rightarrow$ matched premium estimator $\hat{\Pi}_{ij,n}$. The transfer theorem identifies where calibration noise enters the asymptotic variance and when undersmoothing or bias correction removes it at first order.

Proposition 9.10 (Matched pairwise premium estimator from a cleaned $\mathbb{P}/\mathbb{Q}$ surface). *Under Assumption 9.3, define*

$$\hat{\Pi}_{ij,n} := \hat{\gamma}_{ij,n}^{\mathbb{Q}} - \hat{\gamma}_{ij,n}^{\mathbb{P}}.$$

Then $\hat{\Pi}_{ij,n} \rightarrow \Pi_{ij}$ in probability and

$$\sqrt{n} (\hat{\Pi}_{ij,n} - \Pi_{ij}) \Rightarrow \mathcal{N} \left(0, \Sigma_{ij}^{\mathbb{Q}} + \Sigma_{ij}^{\mathbb{P}} - 2 \Sigma_{ij}^{\mathbb{P}\mathbb{Q}} \right).$$

In particular, the paired covariance term is the default asymptotic correction when the realized and option-implied inputs are observed on common dates after option-surface cleaning.

Proof. Consistency follows because $\widehat{\Pi}_{ij,n} - \Pi_{ij}$ is the difference of the two consistent coefficient estimators. For the asymptotic law, write

$$\sqrt{n}(\widehat{\Pi}_{ij,n} - \Pi_{ij}) = \begin{pmatrix} -1 & 1 \end{pmatrix} \sqrt{n} \begin{pmatrix} \widehat{\gamma}_{ij,n}^{\mathbb{P}} - \gamma_{ij}^{\mathbb{P}} \\ \widehat{\gamma}_{ij,n}^{\mathbb{Q}} - \gamma_{ij}^{\mathbb{Q}} \end{pmatrix}.$$

The variance is the quadratic form of the joint covariance matrix under the contrast $(-1, 1)$:

$$(-1, 1) \begin{pmatrix} \Sigma_{ij}^{\mathbb{P}} & \Sigma_{ij}^{\mathbb{P}\mathbb{Q}} \\ \Sigma_{ij}^{\mathbb{P}\mathbb{Q}} & \Sigma_{ij}^{\mathbb{Q}} \end{pmatrix} \begin{pmatrix} -1 \\ 1 \end{pmatrix} = \Sigma_{ij}^{\mathbb{Q}} + \Sigma_{ij}^{\mathbb{P}} - 2\Sigma_{ij}^{\mathbb{P}\mathbb{Q}}.$$

This is the multivariate delta method for the linear map $(x, y) \mapsto y - x$. $\square$

Remark 9.11 (Asymptotic independence as a special case). If the $\mathbb{P}$ - and $\mathbb{Q}$ -side inputs are asymptotically independent, or if the paired design is replaced by asymptotically disjoint samples, then $\Sigma_{ij}^{\mathbb{P}\mathbb{Q}} = 0$ and Proposition 9.10 simplifies to the familiar variance sum

$$\Sigma_{ij}^{\mathbb{Q}} + \Sigma_{ij}^{\mathbb{P}}.$$

The paired covariance term is retained because matched dates are the natural design for comparing realized and option-implied structural loading on the same market states.

Corollary 9.12 (Pairwise $\mathbb{P}/\mathbb{Q}$ classification on a matched cleaned basis). *Fix tolerance levels $\eta_P, \eta_Q, \eta_{\Pi} > 0$. For each ordered pair (i, j) , the common projected basis φ_{ij} and the cleaned matched $\mathbb{Q}$ -surface yield the following population classification:*

- (i) dynamically observable: $|\gamma_{ij}^{\mathbb{P}}| \geq \eta_P$ and $|\gamma_{ij}^{\mathbb{Q}}| \geq \eta_Q$;
- (ii) pricing-visible only: $|\gamma_{ij}^{\mathbb{P}}| < \eta_P$ and $|\gamma_{ij}^{\mathbb{Q}}| \geq \eta_Q$;
- (iii) path-detectable only: $|\gamma_{ij}^{\mathbb{P}}| \geq \eta_P$ and $|\gamma_{ij}^{\mathbb{Q}}| < \eta_Q$;
- (iv) neutral/low-premium: $|\gamma_{ij}^{\mathbb{P}}| < \eta_P$, $|\gamma_{ij}^{\mathbb{Q}}| < \eta_Q$, and $|\Pi_{ij}| < \eta_{\Pi}$;
- (v) premium-dominant: $|\Pi_{ij}| \geq \eta_{\Pi}$ on the common basis, even after harmonized source-screening.

Confidence-region analogues follow from Proposition 9.10.

Definition 9.13 (Pairwise premium support). For $\eta_{\Pi} \geq 0$, define the pairwise premium support at level η_{Π} by

$$\mathcal{E}_{\Pi}(\eta_{\Pi}) := \{(i, j) : |\Pi_{ij}| \geq \eta_{\Pi}\}.$$

When $\eta_{\Pi} = 0$, write

$$\mathcal{E}_{\Pi} := \{(i, j) : \Pi_{ij} \neq 0\}.$$

Proposition 9.14 (Separated $\mathbb{P}/\mathbb{Q}$ observability forces premium support). *Fix $\eta_P, \eta_Q > 0$ and an ordered pair (i, j) .*

- (i) *If $\eta_P > \eta_Q$ and*

$$|\gamma_{ij}^{\mathbb{P}}| \geq \eta_P, \quad |\gamma_{ij}^{\mathbb{Q}}| < \eta_Q,$$

then

$$|\Pi_{ij}| \geq \eta_P - \eta_Q.$$

Consequently,

$$(i, j) \in \mathcal{E}_{\Pi}(\eta_P - \eta_Q) \subseteq \mathcal{E}_{\Pi}.$$

- (ii) *If $\eta_Q > \eta_P$ and*

$$|\gamma_{ij}^{\mathbb{Q}}| \geq \eta_Q, \quad |\gamma_{ij}^{\mathbb{P}}| < \eta_P,$$

then

$$|\Pi_{ij}| \geq \eta_Q - \eta_P.$$

Consequently,

$$(i, j) \in \mathcal{E}_{\Pi}(\eta_Q - \eta_P) \subseteq \mathcal{E}_{\Pi}.$$

In particular, every pair whose harmonized projected coefficients are separated across measures by a strict visibility gap on the common source-screened basis belongs to the nontrivial premium support.

Proof. Since

$$\Pi_{ij} = \gamma_{ij}^{\mathbb{Q}} - \gamma_{ij}^{\mathbb{P}},$$

the reverse triangle inequality gives

$$|\Pi_{ij}| \geq ||\gamma_{ij}^{\mathbb{P}}| - |\gamma_{ij}^{\mathbb{Q}}||.$$

Under the hypotheses of part (i),

$$|\Pi_{ij}| > \eta_P - \eta_Q.$$

This proves part (i). Part (ii) is identical after exchanging $\mathbb{P}$ and $\mathbb{Q}$. $\square$

Corollary 9.15 (Structural support inclusion for the pairwise premium). *Fix an ordered pair (i, j) and suppose the harmonized projected basis φ_{ij} realizes the same source-screened structural direction under $\mathbb{P}$ and $\mathbb{Q}$.*

(i) *If there exist $0 < \eta_Q < \eta_P$ such that the $\mathbb{P}$ -side pair is separated from zero on the common basis,*

$$|\gamma_{ij}^{\mathbb{P}}| \geq \eta_P,$$

while the matched cleaned $\mathbb{Q}$ -side surface places the same pair below the pricing-visibility threshold on that basis,

$$|\gamma_{ij}^{\mathbb{Q}}| < \eta_Q,$$

then

$$(i, j) \in \mathcal{E}_{\Pi}(\eta_P - \eta_Q) \subseteq \mathcal{E}_{\Pi}.$$

(ii) *If there exist $0 < \eta_P < \eta_Q$ such that the matched cleaned $\mathbb{Q}$ -side surface is separated from zero on the common basis,*

$$|\gamma_{ij}^{\mathbb{Q}}| \geq \eta_Q,$$

while the $\mathbb{P}$ -side projected-score geometry places the same pair below the physical source-screened threshold on that basis,

$$|\gamma_{ij}^{\mathbb{P}}| < \eta_P,$$

then

$$(i, j) \in \mathcal{E}_{\Pi}(\eta_Q - \eta_P) \subseteq \mathcal{E}_{\Pi}.$$

Equivalently, every harmonized pair with path-detectability alone or pricing-visibility alone must belong to the premium support.

Proof. For part (i), Proposition 9.14(i) says that a pair with $|\gamma_{ij}^{\mathbb{P}}| \geq \eta_P$ and $|\gamma_{ij}^{\mathbb{Q}}| < \eta_Q$ has $|\gamma_{ij}^{\mathbb{Q}} - \gamma_{ij}^{\mathbb{P}}| \geq \eta_P - \eta_Q$. Part (ii) is the same inequality with $\mathbb{P}$ and $\mathbb{Q}$ interchanged, using Proposition 9.14(ii). Hence any support disagreement above the stated thresholds produces a nonzero signed premium coefficient. $\square$

Remark 9.16 (How visibility inputs feed the support inclusion). The structural content of Corollary 9.15 is that premium support is now tied to $\mathbb{P}/\mathbb{Q}$ visibility separation rather than left entirely at the level of a fixed pairwise estimator. On the $\mathbb{P}$ -side, [4, 1] identify the source-screened, physically nondegenerate pairwise channels. On the $\mathbb{Q}$ -side, [2] identifies when a pairwise contagion channel is pricing-visible or screened at the tested maturity scale. Whenever the matched cleaned $\mathbb{Q}$ -surface and the projected-score $\mathbb{P}$ -geometry are harmonized on the same source-screened basis φ_{ij} , those structural inputs feed directly into the support inclusion theorem above.

Proposition 9.17 (One-dimensional harmonized pairwise nondegeneracy is coefficient nonvanishing). *Fix an ordered pair (i, j) , let $\mu \in \{\mathbb{P}, \mathbb{Q}\}$, and suppose the harmonized reduced experiment for that pair is one-dimensional on the common source-screened basis φ_{ij} , so that*

$$D_{ij}^{\mu} = \gamma_{ij}^{\mu} \varphi_{ij}.$$

Then

$$D_{ij}^{\mu} \neq 0 \quad \Longleftrightarrow \quad \gamma_{ij}^{\mu} \neq 0.$$

Proof. Since φ_{ij} has unit norm, the displayed representation gives

$$\|D_{ij}^\mu\|_{\mu,ij} = |\gamma_{ij}^\mu|.$$

Thus $D_{ij}^\mu = 0$ if and only if its norm is zero, and this is equivalent to $\gamma_{ij}^\mu = 0$. $\square$

Corollary 9.18 (Exact premium support from harmonized path-only or pricing-only pairs). *Fix an ordered pair (i, j) and suppose the harmonized reduced experiments under $\mathbb{P}$ and $\mathbb{Q}$ are both one-dimensional on the common source-screened basis φ_{ij} .*

(i) *If the $\mathbb{P}$ -side pair is physically nondegenerate on the harmonized basis,*

$$D_{ij}^\mathbb{P} \neq 0,$$

while the matched cleaned $\mathbb{Q}$ -side pair is screened on that same basis at the tested scale,

$$D_{ij}^\mathbb{Q} = 0,$$

then

$$\Pi_{ij} \neq 0, \quad (i, j) \in \mathcal{E}_\Pi.$$

(ii) *If the matched cleaned $\mathbb{Q}$ -side pair is pricing-nondegenerate on the harmonized basis,*

$$D_{ij}^\mathbb{Q} \neq 0,$$

while the $\mathbb{P}$ -side source-screened reduced loading vanishes on that same basis,

$$D_{ij}^\mathbb{P} = 0,$$

then

$$\Pi_{ij} \neq 0, \quad (i, j) \in \mathcal{E}_\Pi.$$

In particular, exact harmonized path-only pairs and exact harmonized pricing-only pairs must belong to the pairwise premium support.

Proof. Under the one-dimensional reduction, Proposition 9.17 gives

$$D_{ij}^\mathbb{P} \neq 0 \iff \gamma_{ij}^\mathbb{P} \neq 0, \quad D_{ij}^\mathbb{Q} \neq 0 \iff \gamma_{ij}^\mathbb{Q} \neq 0.$$

If $D_{ij}^\mathbb{P} \neq 0$ and $D_{ij}^\mathbb{Q} = 0$, then

$$\Pi_{ij} = \gamma_{ij}^\mathbb{Q} - \gamma_{ij}^\mathbb{P} = -\gamma_{ij}^\mathbb{P} \neq 0.$$

This proves part (i). Part (ii) is identical after exchanging $\mathbb{P}$ and $\mathbb{Q}$. $\square$

Remark 9.19 (Exact versus tolerance-separated premium support). Corollary 9.18 is the exact-support counterpart to Proposition 9.14. The former is the sharp statement when the harmonized pairwise reduction is genuinely one-dimensional and one side vanishes exactly on the common source-screened basis. The latter is the threshold-separated asymptotic statement when only a visibility gap separated by thresholds η_P, η_Q is available.

The one-dimensional projected-information estimator, covariance-normalizer construction, finite-dimensional if-and-only-if support law, and unsigned-information counterexamples are implementation infrastructure. They are collected in Appendix B. The body keeps the premium claim at coefficient level. Signed matched coefficients identify the priced $\mathbb{P}/\mathbb{Q}$ gap on a common source-screened basis; unsigned information is only a support diagnostic unless signs and normalizers are tracked.

10. VARIANCE CLAIMS AND RISK-NEUTRAL COMPRESSION

Index-option summaries are only one public class. Nonlinear claims expose the same compression operator through convexity. Corollary 4.10 makes the public variance-pricing object a compression of the claim's priced structural loading. The section first records the abstract compression theorem for square-integrable claims and then specializes it to the market-wide object

$$Q_M := \int_0^T M(t)^2 dt.$$

The market-variance motivation is adjacent to the log-contract and volatility-derivatives literature [39, 40]. The result is a compression statement, not a variance-swap replication formula.

Assumption 10.1 (Raw visible-plus-synchronized aggregate reduction under $\mathbb{Q}$). There exist a visible process $M^{\text{vis},\mathbb{Q}}$, the structural synchronized amplitude process ϖ , and a raw residual process $\varepsilon^{\mathbb{Q},\text{raw}}$ such that

$$M(t) = M^{\text{vis},\mathbb{Q}}(t) + (w^\top v)\varpi(t) + \varepsilon^{\mathbb{Q},\text{raw}}(t), \quad 0 \leq t \leq T,$$

where $M^{\text{vis},\mathbb{Q}}(t)$ is $\mathcal{G}_t$ -measurable for every t . Assume further that

$$\int_0^T (M^{\text{vis},\mathbb{Q}}(t))^2 dt, \quad \int_0^T \varpi(t)^2 dt, \quad \int_0^T (\varepsilon^{\mathbb{Q},\text{raw}}(t))^2 dt$$

all belong to $L^2(\mathbb{Q}, \mathcal{F}_T)$, and that the conditional Fubini operations below are valid.

Definition 10.2 (Risk-neutral aggregate-centering defect). Under Assumption 10.1, set

$$\mathcal{H}_\varpi := \mathcal{G}_T \vee \mathcal{L}_T, \quad \zeta(t) := (w^\top v)\varpi(t), \quad A^\mathbb{Q}(t) := M^{\text{vis},\mathbb{Q}}(t) + \zeta(t).$$

Define the $\mathbb{Q}$ -aggregate-centering defect by

$$b_\varpi^\mathbb{Q}(t) := \mathbb{E}_\mathbb{Q}[\varepsilon^{\mathbb{Q},\text{raw}}(t) \mid \mathcal{H}_\varpi], \quad 0 \leq t \leq T,$$

and the recentered residual by

$$\eta_\varpi^{\mathbb{Q},\perp}(t) := \varepsilon^{\mathbb{Q},\text{raw}}(t) - b_\varpi^\mathbb{Q}(t).$$

Thus $b_\varpi^\mathbb{Q}$ measures the failure of the aggregate residual to remain centered after the state-price-density reweighting.

Proposition 10.3 (Risk-neutral recentering of the aggregate reduction). *Under Assumption 10.1 and Definition 10.2,*

$$M(t) = A^\mathbb{Q}(t) + b_\varpi^\mathbb{Q}(t) + \eta_\varpi^{\mathbb{Q},\perp}(t), \quad 0 \leq t \leq T,$$

and

$$\mathbb{E}_\mathbb{Q}[\eta_\varpi^{\mathbb{Q},\perp}(t) \mid \mathcal{H}_\varpi] = 0 \quad \text{for almost every } t \in [0, T].$$

Moreover $b_\varpi^\mathbb{Q}$ and $\eta_\varpi^{\mathbb{Q},\perp}$ inherit the required square integrability, and the priced structural loading of integrated market variance is

$$L_{Q_M}^\mathbb{Q} = \int_0^T (A^\mathbb{Q}(t) + b_\varpi^\mathbb{Q}(t))^2 dt + \int_0^T \mathbb{E}_\mathbb{Q}[(\eta_\varpi^{\mathbb{Q},\perp}(t))^2 \mid \mathcal{H}_\varpi] dt.$$

Proof. The decomposition is the definition of $b_\varpi^\mathbb{Q}$ and $\eta_\varpi^{\mathbb{Q},\perp}$. Conditional centering follows from

$$\mathbb{E}_\mathbb{Q}[\eta_\varpi^{\mathbb{Q},\perp}(t) \mid \mathcal{H}_\varpi] = \mathbb{E}_\mathbb{Q}[\varepsilon^{\mathbb{Q},\text{raw}}(t) \mid \mathcal{H}_\varpi] - b_\varpi^\mathbb{Q}(t) = 0.$$

Conditional Jensen and conditional Fubini give

$$\int_0^T (b_\varpi^\mathbb{Q}(t))^2 dt \leq \mathbb{E}_\mathbb{Q}\left[\int_0^T (\varepsilon^{\mathbb{Q},\text{raw}}(t))^2 dt \mid \mathcal{H}_\varpi\right],$$

so $b_\varpi^\mathbb{Q}$ is square-integrable at the required integrated scale; the same follows for $\eta_\varpi^{\mathbb{Q},\perp}$ from $(x - y)^2 \leq 2x^2 + 2y^2$. Finally,

$$Q_M = \int_0^T (A^\mathbb{Q}(t) + b_\varpi^\mathbb{Q}(t) + \eta_\varpi^{\mathbb{Q},\perp}(t))^2 dt.$$

Taking $\mathbb{E}_{\mathbb{Q}}[\cdot | \mathcal{H}_{\varpi}]$, using that $A^{\mathbb{Q}} + b_{\varpi}^{\mathbb{Q}}$ is $\mathcal{H}_{\varpi}$ -measurable, and applying the displayed centering removes the cross term and gives the formula for $L_{Q_M}^{\mathbb{Q}} = C_{\mathcal{H}_{\varpi}}^{\mathbb{Q}}(Q_M)$. $\square$

Proposition 10.4 (State-price-density formula for the centering defect). *Assume Assumption 2.2. Under Assumption 10.1, for almost every t ,*

$$b_{\varpi}^{\mathbb{Q}}(t) = \frac{\mathbb{E}_{\mathbb{P}}[Z_T \varepsilon^{\mathbb{Q}, \text{raw}}(t) | \mathcal{H}_{\varpi}]}{\mathbb{E}_{\mathbb{P}}[Z_T | \mathcal{H}_{\varpi}]}.$$

If, in addition, the raw residual is $\mathbb{P}$ -centered relative to the same structural sigma-field,

$$\mathbb{E}_{\mathbb{P}}[\varepsilon^{\mathbb{Q}, \text{raw}}(t) | \mathcal{H}_{\varpi}] = 0,$$

then

$$b_{\varpi}^{\mathbb{Q}}(t) = \frac{\text{Cov}_{\mathbb{P}}(Z_T, \varepsilon^{\mathbb{Q}, \text{raw}}(t) | \mathcal{H}_{\varpi})}{\mathbb{E}_{\mathbb{P}}[Z_T | \mathcal{H}_{\varpi}]}.$$

Consequently the exact zero-defect branch holds whenever the state-price density is conditionally orthogonal to the raw aggregate residual given $\mathcal{H}_{\varpi}$.

Proof. For every bounded $\mathcal{H}_{\varpi}$ -measurable Y ,

$$\mathbb{E}_{\mathbb{Q}}[Y \varepsilon^{\mathbb{Q}, \text{raw}}(t)] = \mathbb{E}_{\mathbb{P}}[Z_T Y \varepsilon^{\mathbb{Q}, \text{raw}}(t)] = \mathbb{E}_{\mathbb{P}}[Y \mathbb{E}_{\mathbb{P}}[Z_T \varepsilon^{\mathbb{Q}, \text{raw}}(t) | \mathcal{H}_{\varpi}]].$$

Likewise

$$\mathbb{E}_{\mathbb{Q}}[Y b] = \mathbb{E}_{\mathbb{P}}[Z_T Y b] = \mathbb{E}_{\mathbb{P}}[Y b \mathbb{E}_{\mathbb{P}}[Z_T | \mathcal{H}_{\varpi}]]$$

for every $\mathcal{H}_{\varpi}$ -measurable candidate b . Since $Z_T > 0$ $\mathbb{P}$ -a.s. and $\mathbb{Q} \sim \mathbb{P}$, the denominator $\mathbb{E}_{\mathbb{P}}[Z_T | \mathcal{H}_{\varpi}]$ is strictly positive a.s. The displayed ratio is therefore the $\mathbb{Q}$ -conditional expectation. If the raw residual is $\mathbb{P}$ -centered, subtract $\mathbb{E}_{\mathbb{P}}[Z_T | \mathcal{H}_{\varpi}] \mathbb{E}_{\mathbb{P}}[\varepsilon^{\mathbb{Q}, \text{raw}}(t) | \mathcal{H}_{\varpi}] = 0$ from the numerator to obtain the conditional covariance formula. The final statement is the case in which this conditional covariance vanishes. $\square$

Proposition 10.5 (Small centering defect gives a perturbative variance-loading formula). *Under Assumption 10.1 and Definition 10.2, define the zero-defect benchmark*

$$L_{Q_M}^{\mathbb{Q}, 0} := \int_0^T (A^{\mathbb{Q}}(t))^2 dt + \int_0^T \mathbb{E}_{\mathbb{Q}}[(\varepsilon^{\mathbb{Q}, \text{raw}}(t))^2 | \mathcal{H}_{\varpi}] dt.$$

Then

$$L_{Q_M}^{\mathbb{Q}} - L_{Q_M}^{\mathbb{Q}, 0} = 2 \int_0^T A^{\mathbb{Q}}(t) b_{\varpi}^{\mathbb{Q}}(t) dt.$$

Consequently,

$$\|L_{Q_M}^{\mathbb{Q}} - L_{Q_M}^{\mathbb{Q}, 0}\|_{L^2(\mathbb{Q})} \leq 2 \left\| \int_0^T (A^{\mathbb{Q}}(t))^2 dt \right\|_{L^2(\mathbb{Q})}^{1/2} \left\| \int_0^T (b_{\varpi}^{\mathbb{Q}}(t))^2 dt \right\|_{L^2(\mathbb{Q})}^{1/2}.$$

In particular, the exact centered branch is recovered when $b_{\varpi}^{\mathbb{Q}} = 0$, and a small state-price-density centering defect produces only the displayed first-order loading error.

Proof. Since

$$\varepsilon^{\mathbb{Q}, \text{raw}} = b_{\varpi}^{\mathbb{Q}} + \eta_{\varpi}^{\mathbb{Q}, \perp},$$

conditional orthogonality gives

$$\mathbb{E}_{\mathbb{Q}}[(\varepsilon^{\mathbb{Q}, \text{raw}}(t))^2 | \mathcal{H}_{\varpi}] = (b_{\varpi}^{\mathbb{Q}}(t))^2 + \mathbb{E}_{\mathbb{Q}}[(\eta_{\varpi}^{\mathbb{Q}, \perp}(t))^2 | \mathcal{H}_{\varpi}].$$

Subtracting this identity from the loading formula in Proposition 10.3 leaves the claimed difference identity; the $\int (b_{\varpi}^{\mathbb{Q}})^2$ terms cancel. Pathwise Cauchy–Schwarz gives

$$\left| \int_0^T A^{\mathbb{Q}}(t) b_{\varpi}^{\mathbb{Q}}(t) dt \right|^2 \leq \left(\int_0^T (A^{\mathbb{Q}}(t))^2 dt \right) \left(\int_0^T (b_{\varpi}^{\mathbb{Q}}(t))^2 dt \right).$$

Taking $L^2(\mathbb{Q})$ -norms and applying Cauchy–Schwarz to the product gives the displayed bound. $\square$

Assumption 10.6 (Visible aggregate-volatility reduction under $\mathbb{Q}$). There exist a visible process $M^{\text{vis},\mathbb{Q}}$, the structural synchronized amplitude process ϖ , and a residual process $\eta^{\mathbb{Q}}$ such that

$$M(t) = M^{\text{vis},\mathbb{Q}}(t) + (w^\top v)\varpi(t) + \eta^{\mathbb{Q}}(t), \quad 0 \leq t \leq T,$$

where $M^{\text{vis},\mathbb{Q}}(t)$ is $\mathcal{G}_t$ -measurable for every t and

$$\mathbb{E}_{\mathbb{Q}}[\eta^{\mathbb{Q}}(t) \mid \mathcal{G}_T \vee \mathcal{L}_T] = 0 \quad \text{for almost every } t \in [0, T].$$

Assume further that

$$\int_0^T (M^{\text{vis},\mathbb{Q}}(t))^2 dt, \quad \int_0^T \varpi(t)^2 dt, \quad \int_0^T (\eta^{\mathbb{Q}}(t))^2 dt$$

all belong to $L^2(\mathbb{Q}, \mathcal{F}_T)$.

Remark 10.7 (What is assumed and what is derived in the variance block). Assumption 10.6 is the zero-defect specialization of the raw recentered reduction in Propositions 10.3–10.5. It gives the minimal centered variance formula. When exact $\mathbb{Q}$ -centering fails, the market-variance loading remains explicit with the additional defect $b_{\varpi}^{\mathbb{Q}}$, and the error relative to the zero-defect benchmark is controlled by Proposition 10.5.

Proposition 10.8 (Canonical $\mathbb{Q}$ -decomposition of integrated market variance). *Assume Assumption 10.6 and set*

$$\mathcal{H}_{\varpi} := \mathcal{G}_T \vee \mathcal{L}_T.$$

Then

$$Q_M = L_{Q_M}^{\mathbb{Q}} + U_{Q_M}^{\mathbb{Q}},$$

where

$$L_{Q_M}^{\mathbb{Q}} := \mathbb{E}_{\mathbb{Q}}[Q_M \mid \mathcal{H}_{\varpi}] \in \mathcal{M}_{\mathbb{Q}},$$

and

$$U_{Q_M}^{\mathbb{Q}} := Q_M - \mathbb{E}_{\mathbb{Q}}[Q_M \mid \mathcal{H}_{\varpi}].$$

Moreover,

$$\mathbb{E}_{\mathbb{Q}}[U_{Q_M}^{\mathbb{Q}} \mid \mathcal{H}_{\varpi}] = 0.$$

Proof. The decomposition is the tower-property identity

$$Q_M = \mathbb{E}_{\mathbb{Q}}[Q_M \mid \mathcal{H}_{\varpi}] + (Q_M - \mathbb{E}_{\mathbb{Q}}[Q_M \mid \mathcal{H}_{\varpi}]).$$

The first term belongs to $\mathcal{M}_{\mathbb{Q}}$. The second lies in the residual fiber because

$$\mathbb{E}_{\mathbb{Q}}[Q_M - \mathbb{E}_{\mathbb{Q}}[Q_M \mid \mathcal{H}_{\varpi}] \mid \mathcal{H}_{\varpi}] = 0.$$

The normalization is notation. □

Proposition 10.9 (Explicit form of the market-variance loading). *Writing $\zeta(t) := (w^\top v)\varpi(t)$, one has*

$$L_{Q_M}^{\mathbb{Q}} = \int_0^T (M^{\text{vis},\mathbb{Q}}(t) + \zeta(t))^2 dt + \int_0^T \mathbb{E}_{\mathbb{Q}}[\eta^{\mathbb{Q}}(t)^2 \mid \mathcal{H}_{\varpi}] dt.$$

Proof. Writing $\zeta(t) := (w^\top v)\varpi(t)$, one has

$$M(t) = M^{\text{vis},\mathbb{Q}}(t) + \zeta(t) + \eta^{\mathbb{Q}}(t),$$

so

$$Q_M = \int_0^T (M^{\text{vis},\mathbb{Q}}(t) + \zeta(t))^2 dt + 2 \int_0^T (M^{\text{vis},\mathbb{Q}}(t) + \zeta(t))\eta^{\mathbb{Q}}(t) dt + \int_0^T \eta^{\mathbb{Q}}(t)^2 dt.$$

Taking $\mathbb{E}_{\mathbb{Q}}[\cdot \mid \mathcal{H}_{\varpi}]$ and using

$$\mathbb{E}_{\mathbb{Q}}[\eta^{\mathbb{Q}}(t) \mid \mathcal{H}_{\varpi}] = 0$$

for almost every t , together with conditional Fubini and the square-integrability assumptions, eliminates the cross term and yields the displayed formula. □

Remark 10.10 (Interpretation). Proposition 10.9 identifies the market-variance loading as the sum of the visible-plus-synchronized square and the conditional second moment of the residual channel.

Proposition 10.11 (Canonical extraction of the aggregate factor from integrated market variance). *Assume Assumption 10.6. Let*

$$V_{Q_M}^{\mathbb{Q}} := C_{\mathcal{G}_T}^{\mathbb{Q}}(L_{Q_M}^{\mathbb{Q}}), \quad \tilde{L}_{Q_M}^{\mathbb{Q}} := L_{Q_M}^{\mathbb{Q}} - V_{Q_M}^{\mathbb{Q}}.$$

Then

$$\tilde{L}_{Q_M}^{\mathbb{Q}} = \alpha_{Q_M}^{\mathbb{Q},*} \Lambda_M^*, \quad R_{Q_M}^{\mathbb{Q},*} = 0,$$

where

$$\alpha_{Q_M}^{\mathbb{Q},*} := \|\tilde{L}_{Q_M}^{\mathbb{Q}}\|_{L^2(\mathbb{Q})}$$

whenever $\tilde{L}_{Q_M}^{\mathbb{Q}} \neq 0$, with the convention $\alpha_{Q_M}^{\mathbb{Q},*} := 0$ otherwise. Equivalently,

$$L_{Q_M}^{\mathbb{Q}} = V_{Q_M}^{\mathbb{Q}} + \alpha_{Q_M}^{\mathbb{Q},*} \Lambda_M^*,$$

so integrated market variance generates the canonical aggregate factor benchmark exactly in the sense of Proposition 4.6.

Proof. By Definition 4.5, Λ_M^* is obtained by normalizing the canonical centered structural increment of integrated market variance under the present pricing measure. In the notation of Proposition 10.8, that centered increment is exactly

$$\tilde{L}_{Q_M}^{\mathbb{Q}} = L_{Q_M}^{\mathbb{Q}} - C_{\mathcal{G}_T}^{\mathbb{Q}}(L_{Q_M}^{\mathbb{Q}}).$$

If $\tilde{L}_{Q_M}^{\mathbb{Q}} = 0$, the claim is trivial. Otherwise,

$$\Lambda_M^* = \frac{\tilde{L}_{Q_M}^{\mathbb{Q}}}{\|\tilde{L}_{Q_M}^{\mathbb{Q}}\|_{L^2(\mathbb{Q})}},$$

which yields the stated representation and shows that the orthogonal residual in Proposition 4.6 vanishes for Q_M . $\square$

Proposition 10.12 (Canonical $\mathbb{Q}$ -decomposition of variance-linked claims). *Assume Assumption 10.6, let $X = \Psi(Q_M) \in L^2(\mathbb{Q}, \mathcal{F}_T)$, and set*

$$\mathcal{H}_{\varpi} := \mathcal{G}_T \vee \mathcal{L}_T.$$

Define

$$L_X^{\mathbb{Q}} := \mathbb{E}_{\mathbb{Q}}[X \mid \mathcal{H}_{\varpi}], \\ U_X^{\mathbb{Q}} := X - \mathbb{E}_{\mathbb{Q}}[X \mid \mathcal{H}_{\varpi}].$$

Then

$$X = L_X^{\mathbb{Q}} + U_X^{\mathbb{Q}},$$

where $L_X^{\mathbb{Q}} \in \mathcal{M}_{\mathbb{Q}}$ and

$$\mathbb{E}_{\mathbb{Q}}[U_X^{\mathbb{Q}} \mid \mathcal{H}_{\varpi}] = 0.$$

If $L_X^{\mathbb{Q}} \neq 0$, define

$$\alpha_X^{\mathbb{Q}} := \|L_X^{\mathbb{Q}}\|_{L^2(\mathbb{Q})}, \quad \Lambda_X := \frac{L_X^{\mathbb{Q}}}{\alpha_X^{\mathbb{Q}}};$$

otherwise set $\alpha_X^{\mathbb{Q}} := 0$ and $\Lambda_X := 0$. In all cases,

$$\mathbb{E}_{\mathbb{Q}}[U_X^{\mathbb{Q}} \mid \mathcal{G}_T, \Lambda_X] = 0.$$

Proof. The decomposition is again the tower-property identity

$$X = \mathbb{E}_{\mathbb{Q}}[X \mid \mathcal{H}_{\varpi}] + (X - \mathbb{E}_{\mathbb{Q}}[X \mid \mathcal{H}_{\varpi}]).$$

The first term belongs to $\mathcal{M}_{\mathbb{Q}}$ exactly as in Proposition 10.8. The definition of $U_X^{\mathbb{Q}}$ gives $\mathbb{E}_{\mathbb{Q}}[U_X^{\mathbb{Q}} \mid \mathcal{H}_{\varpi}] = 0$. Since Λ_X is $\mathcal{H}_{\varpi}$ -measurable and $\sigma(\mathcal{G}_T, \Lambda_X) \subseteq \mathcal{H}_{\varpi}$, the tower property gives

$$\mathbb{E}_{\mathbb{Q}}[U_X^{\mathbb{Q}} \mid \mathcal{G}_T, \Lambda_X] = \mathbb{E}_{\mathbb{Q}}[\mathbb{E}_{\mathbb{Q}}[U_X^{\mathbb{Q}} \mid \mathcal{H}_{\varpi}] \mid \mathcal{G}_T, \Lambda_X] = 0.$$

$\square$

Corollary 10.13 (Affine variance-linked claims inherit aggregate-factor alignment). *Assume Assumption 10.6, and let*

$$X = a + bQ_M$$

for constants $a, b \in \mathbb{R}$ such that $X \in L^2(\mathbb{Q}, \mathcal{F}_T)$. Then

$$L_X^{\mathbb{Q}} = C_{\mathcal{G}_T}^{\mathbb{Q}}(L_X^{\mathbb{Q}}) + b\alpha_{Q_M}^{\mathbb{Q},*}\Lambda_M^*,$$

and $R_X^{\mathbb{Q},*} = 0$. Hence every affine variance-linked claim is aggregate-factor aligned.

Proof. Linearity of conditional expectation gives

$$L_X^{\mathbb{Q}} = a + bL_{Q_M}^{\mathbb{Q}}.$$

Apply Proposition 10.11 to $L_{Q_M}^{\mathbb{Q}}$:

$$L_X^{\mathbb{Q}} = a + bV_{Q_M}^{\mathbb{Q}} + b\alpha_{Q_M}^{\mathbb{Q},*}\Lambda_M^*.$$

Since $a + bV_{Q_M}^{\mathbb{Q}} \in L^2(\mathbb{Q}, \mathcal{G}_T)$, this is exactly the decomposition from Proposition 4.6 with zero orthogonal centered remainder. $\square$

Remark 10.14 (Why nonlinear variance-linked claims require a hierarchy). Corollary 10.13 is exact because affine transformations preserve the one-factor structure of the centered market-variance loading. A general nonlinear claim $X = \Psi(Q_M)$ need not do so. After conditioning onto $\mathcal{H}_{\varpi}$, the centered structural loading of $\Psi(Q_M)$ can acquire higher-order centered components beyond the span of Λ_M^* . Thus exact one-factor inheritance is the affine branch, while nonlinear variance-linked claims are handled by the general compression and hedging theorems and, in the one-factor experiment below, by the explicit centered-power hierarchy in Proposition 10.22.

Lemma 10.15 (Nonlinearity plus compression of the priced structural loading). *Let $L \in L^2(\mathbb{Q}, \mathcal{H}_{\varpi})$, let $\mathcal{H} \subseteq \mathcal{G}_T$, and let $\Phi: \mathbb{R} \rightarrow \mathbb{R}$ be convex and twice continuously differentiable with*

$$\inf_x \Phi''(x) \geq \underline{\kappa} > 0.$$

Assume $\Phi(L), \Phi(C_{\mathcal{H}}^{\mathbb{Q}}(L)) \in L^1(\mathbb{Q})$. Then

$$\mathbb{E}_{\mathbb{Q}}[\Phi(L)] - \mathbb{E}_{\mathbb{Q}}[\Phi(C_{\mathcal{H}}^{\mathbb{Q}}(L))] \geq \frac{\underline{\kappa}}{2} \|L - C_{\mathcal{H}}^{\mathbb{Q}}(L)\|_{L^2(\mathbb{Q})}^2.$$

If $L = \alpha\Lambda$ with $\alpha \geq 0$ and $\|\Lambda\|_{L^2(\mathbb{Q})} = 1$, then

$$\mathbb{E}_{\mathbb{Q}}[\Phi(\alpha\Lambda)] - \mathbb{E}_{\mathbb{Q}}[\Phi(\alpha C_{\mathcal{H}}^{\mathbb{Q}}(\Lambda))] \geq \frac{\underline{\kappa}}{2} \|\alpha\Lambda - \alpha C_{\mathcal{H}}^{\mathbb{Q}}(\Lambda)\|_{L^2(\mathbb{Q})}^2.$$

Proof. Since $C_{\mathcal{H}}^{\mathbb{Q}}(L) = \mathbb{E}_{\mathbb{Q}}[L | \mathcal{H}]$, conditional Jensen together with the strong convexity bound gives

$$\mathbb{E}_{\mathbb{Q}}[\Phi(L)] - \mathbb{E}_{\mathbb{Q}}[\Phi(C_{\mathcal{H}}^{\mathbb{Q}}(L))] \geq \frac{\underline{\kappa}}{2} \mathbb{E}_{\mathbb{Q}}[\text{Var}_{\mathbb{Q}}(L | \mathcal{H})].$$

The first claim follows from

$$\mathbb{E}_{\mathbb{Q}}[\text{Var}_{\mathbb{Q}}(L | \mathcal{H})] = \|L - C_{\mathcal{H}}^{\mathbb{Q}}(L)\|_{L^2(\mathbb{Q})}^2.$$

The second is the specialization $L = \alpha\Lambda$. $\square$

The variance-compression theorem identifies the Jensen gap left after public conditioning of the priced structural loading.

Theorem 10.16 (Risk-neutral compression of priced structural claims). *Let $X \in L^2(\mathbb{Q}, \mathcal{F}_T)$, and let*

$$X = L_X^{\mathbb{Q}} + U_X^{\mathbb{Q}}$$

be the decomposition of Proposition 4.2. Let $\Phi: \mathbb{R} \rightarrow \mathbb{R}$ be convex and twice continuously differentiable with

$$\inf_x \Phi''(x) \geq \underline{\kappa} > 0,$$

and assume $\Phi(X), \Phi(C_{\mathcal{H}}^{\mathbb{Q}}(X)) \in L^1(\mathbb{Q})$ for every sub- σ -field $\mathcal{H} \subseteq \mathcal{G}_T$ considered below. Then for every such $\mathcal{H}$,

$$\mathbb{E}_{\mathbb{Q}}[\Phi(X)] - \mathbb{E}_{\mathbb{Q}}[\Phi(C_{\mathcal{H}}^{\mathbb{Q}}(X))] \geq \frac{\kappa}{2} \|X - C_{\mathcal{H}}^{\mathbb{Q}}(X)\|_{L^2(\mathbb{Q})}^2.$$

Equivalently,

$$\mathbb{E}_{\mathbb{Q}}[\Phi(X)] - \mathbb{E}_{\mathbb{Q}}[\Phi(C_{\mathcal{H}}^{\mathbb{Q}}(X))] \geq \frac{\kappa}{2} \left(\|L_X^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}}(L_X^{\mathbb{Q}})\|_{L^2(\mathbb{Q})}^2 + \|U_X^{\mathbb{Q}}\|_{L^2(\mathbb{Q})}^2 \right).$$

If $L_X^{\mathbb{Q}} = \alpha_X^{\mathbb{Q}} \Lambda_X$, the same lower bound is

$$\frac{\kappa}{2} \left((\alpha_X^{\mathbb{Q}})^2 \|\Lambda_X - C_{\mathcal{H}}^{\mathbb{Q}}(\Lambda_X)\|_{L^2(\mathbb{Q})}^2 + \|U_X^{\mathbb{Q}}\|_{L^2(\mathbb{Q})}^2 \right).$$

In particular, compression to any visible pricing summary $\sigma(\mathfrak{S}) \subseteq \mathcal{G}_T$ leaves a strict positive gap whenever either

$$\|L_X^{\mathbb{Q}} - C_{\sigma(\mathfrak{S})}^{\mathbb{Q}}(L_X^{\mathbb{Q}})\|_{L^2(\mathbb{Q})}^2 > 0$$

or $\|U_X^{\mathbb{Q}}\|_{L^2(\mathbb{Q})}^2 > 0$.

Proof. For every $x, y \in \mathbb{R}$, the curvature lower bound gives

$$\Phi(y) \geq \Phi(x) + \Phi'(x)(y - x) + \frac{\kappa}{2}(y - x)^2.$$

Set $x = C_{\mathcal{H}}^{\mathbb{Q}}(X)$ and $y = X$, and condition on $\mathcal{H}$. The affine term vanishes because $\mathbb{E}_{\mathbb{Q}}[X - C_{\mathcal{H}}^{\mathbb{Q}}(X) \mid \mathcal{H}] = 0$. Therefore

$$\mathbb{E}_{\mathbb{Q}}[\Phi(X)] - \mathbb{E}_{\mathbb{Q}}[\Phi(C_{\mathcal{H}}^{\mathbb{Q}}(X))] \geq \frac{\kappa}{2} \mathbb{E}_{\mathbb{Q}}[\text{Var}_{\mathbb{Q}}(X \mid \mathcal{H})] = \frac{\kappa}{2} \|X - C_{\mathcal{H}}^{\mathbb{Q}}(X)\|_{L^2(\mathbb{Q})}^2.$$

Now $L_X^{\mathbb{Q}}$ is $\mathcal{H}_{\varpi}$ -measurable and Proposition 4.2 gives $\mathbb{E}_{\mathbb{Q}}[U_X^{\mathbb{Q}} \mid \mathcal{H}_{\varpi}] = 0$. Since $\mathcal{H} \subseteq \mathcal{G}_T \subseteq \mathcal{H}_{\varpi}$ and $\sigma(\mathcal{H}, L_X^{\mathbb{Q}}) \subseteq \mathcal{H}_{\varpi}$,

$$\mathbb{E}_{\mathbb{Q}}[U_X^{\mathbb{Q}} \mid \mathcal{H}, L_X^{\mathbb{Q}}] = 0.$$

Equivalently, $L_X^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}}(L_X^{\mathbb{Q}}) \in L^2(\mathbb{Q}, \mathcal{H}_{\varpi})$ while $U_X^{\mathbb{Q}} \perp L^2(\mathbb{Q}, \mathcal{H}_{\varpi})$. Since

$$X - C_{\mathcal{H}}^{\mathbb{Q}}(X) = L_X^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}}(L_X^{\mathbb{Q}}) + U_X^{\mathbb{Q}},$$

Pythagoras gives

$$\|X - C_{\mathcal{H}}^{\mathbb{Q}}(X)\|_{L^2(\mathbb{Q})}^2 = \|L_X^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}}(L_X^{\mathbb{Q}})\|_{L^2(\mathbb{Q})}^2 + \|U_X^{\mathbb{Q}}\|_{L^2(\mathbb{Q})}^2.$$

Combining the displays proves the second inequality. The scalar-factor form follows from $L_X^{\mathbb{Q}} = \alpha_X^{\mathbb{Q}} \Lambda_X$. $\square$

Corollary 10.17 (Factor-level decomposition of the compression gap). *Under the assumptions of Proposition 4.6 and 10.16, for every public sub- σ -field $\mathcal{H} \subseteq \mathcal{G}_T$,*

$$\|L_X^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}}(L_X^{\mathbb{Q}})\|_{L^2(\mathbb{Q})}^2 = \|V_X^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}}(V_X^{\mathbb{Q}})\|_{L^2(\mathbb{Q})}^2 + (\alpha_X^{\mathbb{Q},*})^2 + \|R_X^{\mathbb{Q},*}\|_{L^2(\mathbb{Q})}^2.$$

If X is aggregate-factor aligned, so $R_X^{\mathbb{Q},*} = 0$, this reduces to

$$\|L_X^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}}(L_X^{\mathbb{Q}})\|_{L^2(\mathbb{Q})}^2 = \|V_X^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}}(V_X^{\mathbb{Q}})\|_{L^2(\mathbb{Q})}^2 + (\alpha_X^{\mathbb{Q},*})^2.$$

Consequently,

$$\mathbb{E}_{\mathbb{Q}}[\Phi(X)] - \mathbb{E}_{\mathbb{Q}}[\Phi(C_{\mathcal{H}}^{\mathbb{Q}}(X))] \geq \frac{\kappa}{2} \left(\|V_X^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}}(V_X^{\mathbb{Q}})\|_{L^2(\mathbb{Q})}^2 + (\alpha_X^{\mathbb{Q},*})^2 + \|R_X^{\mathbb{Q},*}\|_{L^2(\mathbb{Q})}^2 \right).$$

Proof. By Proposition 4.6,

$$L_X^{\mathbb{Q}} = V_X^{\mathbb{Q}} + \alpha_X^{\mathbb{Q},*} \Lambda_M^* + R_X^{\mathbb{Q},*}.$$

Because $\mathcal{H} \subseteq \mathcal{G}_T$, Corollary 4.7 yields

$$C_{\mathcal{H}}^{\mathbb{Q}}(L_X^{\mathbb{Q}}) = C_{\mathcal{H}}^{\mathbb{Q}}(V_X^{\mathbb{Q}}).$$

Therefore

$$L_X^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}}(L_X^{\mathbb{Q}}) = (V_X^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}}(V_X^{\mathbb{Q}})) + \alpha_X^{\mathbb{Q},*} \Lambda_M^* + R_X^{\mathbb{Q},*}.$$

The first term belongs to $L^2(\mathbb{Q}, \mathcal{G}_T)$, while $\alpha_X^{\mathbb{Q},*} \Lambda_M^* + R_X^{\mathbb{Q},*} \in \mathcal{M}_{\mathbb{Q}}^{\circ}$; hence it is orthogonal to the first term. Proposition 4.6 also gives $R_X^{\mathbb{Q},*} \perp \Lambda_M^*$. The displayed norm identity follows, and the final lower bound is Theorem 10.16 combined with that identity. $\square$

Corollary 10.18 (Market-variance application of the $\mathbb{Q}$ -side compression theorem). *Under Assumption 10.6, Theorem 10.16 applies in particular to Q_M and to variance-linked claims $X = \Psi(Q_M)$. By Proposition 10.11 and Corollary 10.13, the centered market-variance loading itself and every affine variance-linked claim are exactly aggregate-factor aligned. This is the risk-neutral analogue of the compression result in [1]. The option block and the market-variance block therefore use the same projection structure. Public prices are visible $\mathbb{Q}$ -compressions of claims that may still carry priced structural loadings in $\mathcal{M}_{\mathbb{Q}}$.*

Assumption 10.19 (One-factor aggregate variance experiment). Let $\Xi \in L^4(\mathbb{Q})$ be independent of $\mathcal{G}_T$, with

$$\mathbb{E}_{\mathbb{Q}}[\Xi] = 0, \quad \mathbb{E}_{\mathbb{Q}}[\Xi^2] = 1.$$

Let $f \in L^2([0, T])$ be deterministic and nonzero, and let $m(t)$ be a square-integrable $\mathcal{G}_T$ -measurable visible process. Fix $a \in \mathbb{R}$. The aggregate volatility is

$$M(t) = m(t) + af(t)\Xi, \quad 0 \leq t \leq T.$$

Set

$$Q_M := \int_0^T M(t)^2 dt, \quad A_G := \int_0^T m(t)f(t) dt, \quad B_f := \int_0^T f(t)^2 dt,$$

and

$$V_Q := \int_0^T m(t)^2 dt + a^2 B_f.$$

Proposition 10.20 (One-factor market-variance compression gap). *Under Assumption 10.19,*

$$\mathbb{E}_{\mathbb{Q}}[Q_M \mid \mathcal{G}_T] = V_Q$$

and

$$Q_M - V_Q = 2aA_G\Xi + a^2B_f(\Xi^2 - 1).$$

Consequently, for every public $\mathcal{H} \subseteq \mathcal{G}_T$,

$$\|Q_M - \mathbb{E}_{\mathbb{Q}}[Q_M \mid \mathcal{H}]\|_{L^2(\mathbb{Q})}^2 = \|V_Q - \mathbb{E}_{\mathbb{Q}}[V_Q \mid \mathcal{H}]\|_{L^2(\mathbb{Q})}^2 + \Delta_Q^2,$$

where

$$\Delta_Q^2 := 4a^2\mathbb{E}_{\mathbb{Q}}[A_G^2] + 4a^3B_f\mathbb{E}_{\mathbb{Q}}[A_G]\mathbb{E}_{\mathbb{Q}}[\Xi(\Xi^2 - 1)] + a^4B_f^2\text{Var}_{\mathbb{Q}}(\Xi^2).$$

If Ξ is symmetric, this reduces to

$$\Delta_Q^2 = 4a^2\mathbb{E}_{\mathbb{Q}}[A_G^2] + a^4B_f^2\text{Var}_{\mathbb{Q}}(\Xi^2).$$

Let $\Phi : \mathbb{R} \rightarrow \mathbb{R}$ be twice continuously differentiable and strongly convex with $\inf_x \Phi''(x) \geq \underline{\kappa} > 0$, and assume the displayed expectations are finite. Then

$$\mathbb{E}_{\mathbb{Q}}[\Phi(Q_M)] - \mathbb{E}_{\mathbb{Q}}[\Phi(\mathbb{E}_{\mathbb{Q}}[Q_M \mid \mathcal{H}])] \geq \frac{\underline{\kappa}}{2} \left(\|V_Q - \mathbb{E}_{\mathbb{Q}}[V_Q \mid \mathcal{H}]\|_{L^2(\mathbb{Q})}^2 + \Delta_Q^2 \right).$$

In particular, even when $\mathcal{H} = \mathcal{G}_T$ and all visible variance information is used, the strongly convex compression gap is bounded below by

$$\frac{\underline{\kappa}}{2} \Delta_Q^2.$$

Proof. Expanding the square gives

$$Q_M = \int_0^T m(t)^2 dt + 2a\Xi \int_0^T m(t)f(t) dt + a^2\Xi^2 \int_0^T f(t)^2 dt.$$

Since Ξ is independent of $\mathcal{G}_T$, centered, and has unit variance,

$$\mathbb{E}_{\mathbb{Q}}[Q_M \mid \mathcal{G}_T] = \int_0^T m(t)^2 dt + a^2B_f = V_Q,$$

and the centered hidden part is

$$Q_M - V_Q = 2aA_G\Xi + a^2B_f(\Xi^2 - 1).$$

For $\mathcal{H} \subseteq \mathcal{G}_T$, conditional independence gives

$$\mathbb{E}_{\mathbb{Q}}[Q_M \mid \mathcal{H}] = \mathbb{E}_{\mathbb{Q}}[V_Q \mid \mathcal{H}].$$

Thus

$$Q_M - \mathbb{E}_{\mathbb{Q}}[Q_M \mid \mathcal{H}] = (V_Q - \mathbb{E}_{\mathbb{Q}}[V_Q \mid \mathcal{H}]) + 2aA_G\Xi + a^2B_f(\Xi^2 - 1).$$

The first term is $\mathcal{G}_T$ -measurable. It is orthogonal to both hidden terms because $\mathbb{E}_{\mathbb{Q}}[\Xi] = 0$, $\mathbb{E}_{\mathbb{Q}}[\Xi^2 - 1] = 0$, and Ξ is independent of $\mathcal{G}_T$. Taking the L^2 -norm gives the norm identity; the cross term between $2aA_G\Xi$ and $a^2B_f(\Xi^2 - 1)$ is exactly

$$4a^3B_f\mathbb{E}_{\mathbb{Q}}[A_G]\mathbb{E}_{\mathbb{Q}}[\Xi(\Xi^2 - 1)].$$

The strong-convexity bound follows from conditional Jensen with the variance correction:

$$\mathbb{E}_{\mathbb{Q}}[\Phi(Q_M)] - \mathbb{E}_{\mathbb{Q}}[\Phi(\mathbb{E}_{\mathbb{Q}}[Q_M \mid \mathcal{H}])] \geq \frac{\kappa}{2}\mathbb{E}_{\mathbb{Q}}[\text{Var}_{\mathbb{Q}}(Q_M \mid \mathcal{H})] = \frac{\kappa}{2}\|Q_M - \mathbb{E}_{\mathbb{Q}}[Q_M \mid \mathcal{H}]\|_{L^2(\mathbb{Q})}^2.$$

□

Corollary 10.21 (Pure hidden-factor variance floor). *If, in Proposition 10.20, $m \equiv 0$, then*

$$Q_M = a^2B_f\Xi^2, \quad \mathbb{E}_{\mathbb{Q}}[Q_M \mid \mathcal{G}_T] = a^2B_f,$$

and for every $\mathcal{H} \subseteq \mathcal{G}_T$,

$$\|Q_M - \mathbb{E}_{\mathbb{Q}}[Q_M \mid \mathcal{H}]\|_{L^2(\mathbb{Q})}^2 = a^4B_f^2 \text{Var}_{\mathbb{Q}}(\Xi^2).$$

If $\Xi \sim \mathcal{N}(0, 1)$, this is

$$\|Q_M - \mathbb{E}_{\mathbb{Q}}[Q_M \mid \mathcal{H}]\|_{L^2(\mathbb{Q})}^2 = 2a^4B_f^2.$$

Proposition 10.22 (One-factor nonlinear variance-linked compression hierarchy). *Assume Assumption 10.19, and set*

$$H_Q := 2aA_G\Xi + a^2B_f(\Xi^2 - 1).$$

Then

$$Q_M = V_Q + H_Q, \quad \mathbb{E}_{\mathbb{Q}}[H_Q \mid \mathcal{G}_T] = 0.$$

Let $\Psi : \mathbb{R} \rightarrow \mathbb{R}$ be Borel with $\Psi(Q_M) \in L^2(\mathbb{Q})$, and define

$$V_{\Psi} := \mathbb{E}_{\mathbb{Q}}[\Psi(V_Q + H_Q) \mid \mathcal{G}_T], \quad R_{\Psi} := \Psi(V_Q + H_Q) - V_{\Psi}.$$

Then

$$\mathbb{E}_{\mathbb{Q}}[\Psi(Q_M) \mid \mathcal{G}_T] = V_{\Psi}, \quad \mathbb{E}_{\mathbb{Q}}[R_{\Psi} \mid \mathcal{G}_T] = 0.$$

Consequently, for every public σ -field $\mathcal{H} \subseteq \mathcal{G}_T$,

$$\mathbb{E}_{\mathbb{Q}}[\Psi(Q_M) \mid \mathcal{H}] = \mathbb{E}_{\mathbb{Q}}[V_{\Psi} \mid \mathcal{H}],$$

and for every closed visible hedge class $\mathcal{S} \subseteq L^2(\mathbb{Q}, \mathcal{G}_T)$,

$$\inf_{Y \in \mathcal{S}} \mathbb{E}_{\mathbb{Q}}[(\Psi(Q_M) - Y)^2] = \text{dist}_{L^2(\mathbb{Q})}(V_{\Psi}, \mathcal{S})^2 + \|R_{\Psi}\|_{L^2(\mathbb{Q})}^2.$$

If, in addition,

$$\Psi(x) = \sum_{r=0}^p c_r x^r$$

and the displayed polynomial terms are square-integrable, then

$$V_{\Psi} = \sum_{r=0}^p c_r \sum_{k=0}^r \binom{r}{k} V_Q^{r-k} \mathbb{E}_{\mathbb{Q}}[H_Q^k \mid \mathcal{G}_T],$$

while the full-visible hidden residual is

$$R_{\Psi} = \sum_{r=0}^p c_r \sum_{k=1}^r \binom{r}{k} V_Q^{r-k} (H_Q^k - \mathbb{E}_{\mathbb{Q}}[H_Q^k \mid \mathcal{G}_T]).$$

Proof. The identity $Q_M = V_Q + H_Q$ and the centering $\mathbb{E}_\mathbb{Q}[H_Q \mid \mathcal{G}_T] = 0$ are exactly Proposition 10.20. The definitions of V_Ψ and R_Ψ give

$$\Psi(Q_M) = V_\Psi + R_\Psi, \quad \mathbb{E}_\mathbb{Q}[R_\Psi \mid \mathcal{G}_T] = 0.$$

The tower property gives, for $\mathcal{H} \subseteq \mathcal{G}_T$,

$$\mathbb{E}_\mathbb{Q}[\Psi(Q_M) \mid \mathcal{H}] = \mathbb{E}_\mathbb{Q}[\mathbb{E}_\mathbb{Q}[\Psi(Q_M) \mid \mathcal{G}_T] \mid \mathcal{H}] = \mathbb{E}_\mathbb{Q}[V_\Psi \mid \mathcal{H}].$$

If $Y \in \mathcal{S} \subseteq L^2(\mathbb{Q}, \mathcal{G}_T)$, then $V_\Psi - Y$ is $\mathcal{G}_T$ -measurable and therefore orthogonal to R_Ψ . Hence

$$\mathbb{E}_\mathbb{Q}[(\Psi(Q_M) - Y)^2] = \|V_\Psi - Y\|_{L^2(\mathbb{Q})}^2 + \|R_\Psi\|_{L^2(\mathbb{Q})}^2.$$

Taking the infimum over $Y \in \mathcal{S}$ gives the hedging identity. The polynomial formulas follow by expanding $(V_Q + H_Q)^r$, using that V_Q is $\mathcal{G}_T$ -measurable, and subtracting the conditional expectation from the full polynomial. $\square$

Corollary 10.23 (Quadratic variance-linked claims expose the next hidden component). *Assume Assumption 10.19, and assume $Q_M^2 \in L^2(\mathbb{Q})$. For*

$$\Psi(x) = \alpha + \beta x + \gamma x^2$$

one has

$$V_\Psi = \alpha + \beta V_Q + \gamma(V_Q^2 + \mathbb{E}_\mathbb{Q}[H_Q^2 \mid \mathcal{G}_T]),$$

and

$$R_\Psi = (\beta + 2\gamma V_Q)H_Q + \gamma(H_Q^2 - \mathbb{E}_\mathbb{Q}[H_Q^2 \mid \mathcal{G}_T]).$$

Moreover,

$$\mathbb{E}_\mathbb{Q}[H_Q^2 \mid \mathcal{G}_T] = 4a^2 A_G^2 + 4a^3 A_G B_f \mathbb{E}_\mathbb{Q}[\Xi(\Xi^2 - 1)] + a^4 B_f^2 \text{Var}_\mathbb{Q}(\Xi^2).$$

If Ξ is symmetric, this conditional second moment reduces to

$$4a^2 A_G^2 + a^4 B_f^2 \text{Var}_\mathbb{Q}(\Xi^2).$$

Thus the affine case $\gamma = 0$ is the single hidden-factor branch, while a genuinely quadratic claim exposes the additional centered-square component

$$H_Q^2 - \mathbb{E}_\mathbb{Q}[H_Q^2 \mid \mathcal{G}_T]$$

with nonzero contribution whenever $\gamma \neq 0$ and this centered square is nonzero in $L^2(\mathbb{Q})$.

Proof. Apply Proposition 10.22 with $p = 2$. Since $\mathbb{E}_\mathbb{Q}[H_Q \mid \mathcal{G}_T] = 0$, the $k = 1$ term contributes only to the centered residual. Finally, write

$$H_Q = 2aA_G\Xi + a^2B_f(\Xi^2 - 1)$$

and use independence of Ξ from $\mathcal{G}_T$, together with $\mathbb{E}_\mathbb{Q}[\Xi^2] = 1$, to compute $\mathbb{E}_\mathbb{Q}[H_Q^2 \mid \mathcal{G}_T]$. Symmetry of Ξ gives $\mathbb{E}_\mathbb{Q}[\Xi(\Xi^2 - 1)] = 0$. $\square$

Assumption 10.24 (Conditionally Gaussian aggregate-variance reduction). Let $\mathcal{G}$ be the visible market sigma-field. Assume the integrated market-variance loading has the one-factor form

$$Q_M = V + \alpha Z + S\xi,$$

where V and $S \geq 0$ are $\mathcal{G}$ -measurable, $\alpha \geq 0$, and (Z, ξ) are independent standard normal variables independent of $\mathcal{G}$. The canonical hidden aggregate factor is Z . For a Borel payoff Ψ , define $X = \Psi(Q_M)$, assume $X \in L^2(\mathbb{Q})$, and write

$$P_s \Psi(x) := \mathbb{E}[\Psi(x + s\xi)]$$

whenever the expectation is finite.

Proposition 10.25 (Exact one-factor Hermite hierarchy for nonlinear variance-linked claims). *Under Assumption 10.24, the priced structural loading of $X = \Psi(Q_M)$ relative to $\mathcal{G} \vee \sigma(Z)$ is*

$$L_X^{\mathbb{Q}} := \mathbb{E}_{\mathbb{Q}}[X \mid \mathcal{G}, Z] = P_S \Psi(V + \alpha Z).$$

For each fixed (V, S) , it has the Hermite expansion

$$L_X^{\mathbb{Q}} = c_0(V, S) + \sum_{m \geq 1} c_m(V, S) \text{He}_m(Z)$$

in L^2 , where

$$c_m(V, S) = \frac{1}{m!} \mathbb{E}[P_S \Psi(V + \alpha \zeta) \text{He}_m(\zeta) \mid V, S], \quad \zeta \sim N(0, 1).$$

If $P_S \Psi$ has m weak derivatives with the required integrability, then

$$c_m(V, S) = \frac{\alpha^m}{m!} \mathbb{E}[(P_S \Psi)^{(m)}(V + \alpha \zeta) \mid V, S].$$

Thus arbitrary nonlinear Ψ contributes through the explicit Hermite coefficient sequence $(c_m)_{m \geq 0}$, as well as through the implicit conditional expectation.

Proof. Conditioning on $\mathcal{G} \vee \sigma(Z)$ freezes V, S, Z and integrates only the independent normal residual ξ , giving $L_X^{\mathbb{Q}} = P_S \Psi(V + \alpha Z)$. For fixed (V, S) , the function $z \mapsto P_S \Psi(V + \alpha z)$ belongs to L^2 under the standard normal law by $X \in L^2$ and Jensen's inequality. Expanding it in the Hermite basis gives the coefficient formula. The derivative representation follows from Gaussian integration by parts:

$$\mathbb{E}[f(\zeta) \text{He}_m(\zeta)] = \mathbb{E}[f^{(m)}(\zeta)]$$

with $f(z) = P_S \Psi(V + \alpha z)$. □

Corollary 10.26 (Visible-beta first-factor law). *Assume the hypotheses of Proposition 10.25. If $P_S \Psi$ has two derivatives and the corresponding second-derivative envelope is square-integrable for small α , then*

$$L_X^{\mathbb{Q}} = C_{\mathcal{G}}^{\mathbb{Q}}(L_X^{\mathbb{Q}}) + \alpha(P_S \Psi)'(V)Z + R_{\Psi, \alpha},$$

with

$$\|R_{\Psi, \alpha}\|_{L^2(\mathbb{Q})} = O(\alpha^2).$$

Consequently the first hidden variance factor is the same aggregate factor Z , with visible beta $(P_S \Psi)'(V)$. The nonlinear residual beyond that first factor is curvature-order.

Proof. The Hermite hierarchy gives

$$L_X^{\mathbb{Q}} - C_{\mathcal{G}}^{\mathbb{Q}}(L_X^{\mathbb{Q}}) = \sum_{m \geq 1} c_m(V, S) \text{He}_m(Z).$$

The first coefficient is

$$c_1(V, S) = \alpha \mathbb{E}[(P_S \Psi)'(V + \alpha \zeta) \mid V, S] = \alpha(P_S \Psi)'(V) + O_{L^2}(\alpha^2),$$

where the last step is Taylor's formula under the assumed second-derivative envelope. For $m \geq 2$, the Hermite coefficient formula and the same envelope give the required $O_{L^2}(\alpha^m)$ bounds at this order. Orthogonality of the Hermite polynomials then makes the sum of all terms beyond the displayed first factor an L^2 -residual of order α^2 . □

Corollary 10.27 (Simple nonlinear compression gap). *Under the assumptions of Corollary 10.26, if the visible hedge or price summary is contained in $L^2(\mathbb{Q}, \mathcal{G})$, then*

$$\|L_X^{\mathbb{Q}} - C_{\mathcal{G}}^{\mathbb{Q}}(L_X^{\mathbb{Q}})\|_{L^2(\mathbb{Q})}^2 = \alpha^2 \mathbb{E}_{\mathbb{Q}}[((P_S \Psi)'(V))^2] + O(\alpha^3).$$

If $(P_S \Psi)'(V)$ is nonzero with positive probability, the nonlinear variance-linked claim has a positive first-order hidden aggregate-factor gap.

Proof. The factor Z is independent of $\mathcal{G}$ and centered, so

$$C_{\mathcal{G}}^{\mathbb{Q}}((P_S\Psi)'(V)Z) = 0.$$

The visible-beta law gives

$$L_X^{\mathbb{Q}} - C_{\mathcal{G}}^{\mathbb{Q}}(L_X^{\mathbb{Q}}) = \alpha(P_S\Psi)'(V)Z + O_{L^2}(\alpha^2).$$

Taking squared norms and using $\mathbb{E}[Z^2] = 1$ gives the expansion. $\square$

Corollary 10.28 (Affine claims are the only exact deterministic-beta scalar one-factor family). *In the model of Assumption 10.24 with $S = 0$ and nondegenerate V , suppose*

$$\Psi(V + \alpha Z) = A(V) + bZ$$

for some $\mathcal{G}$ -measurable $A(V)$ and deterministic scalar b , for all sufficiently small $\alpha > 0$. Then Ψ is affine on the support of $V + \alpha Z$. Hence nonlinear claims cannot generally have an exact deterministic-beta scalar one-factor law; the exact law is the Hermite hierarchy, and the first-order law has visible beta.

Proof. For fixed $V = v$, the map $z \mapsto \Psi(v + \alpha z)$ is affine in z on the support of the standard normal law. Therefore $x \mapsto \Psi(x)$ is affine on the translated line $x = v + \alpha z$. Since the support is an interval and V is nondegenerate, these intervals cover the support of $V + \alpha Z$ under the stated nondegeneracy. Thus a version of Ψ is affine on that support, up to null sets. Non-affine Ψ produces higher Hermite coefficients in Proposition 10.25. $\square$

Remark 10.29 (Nonlinear variance-linked law). For nonlinear variance-linked claims, the correct statement is not that every Ψ is scalar one-factor. That statement is false except in affine cases. The correct formulation is: under a conditionally Gaussian aggregate-variance reduction, every square-integrable nonlinear variance-linked claim has an exact one-factor Hermite hierarchy, and its leading hidden component is the same aggregate factor with beta $(P_S\Psi)'(V)$, up to curvature-order error.

11. VISIBLE HEDGING AND RISK-NEUTRAL INCOMPLETENESS

Visible hedging changes the space from public information to tradeable span after the variance compression of Section 10. By Corollary 4.10, a visible hedge acts on the same priced structural loading $L_X^{\mathbb{Q}}$, but only through its projection onto a closed public hedge class. Incompleteness is therefore a projection residual. A visible hedge class cannot eliminate structural loading outside its span. This is the Hilbert-space compression form of incomplete-market mean-square hedging [10, 11, 12, 13, 9].

The hedging theorem identifies that residual exactly.

Theorem 11.1 (Visible hedging as projection failure of priced structural loading). *Let $X \in L^2(\mathbb{Q}, \mathcal{F}_T)$, and let*

$$X = L_X^{\mathbb{Q}} + U_X^{\mathbb{Q}}$$

be the decomposition of Proposition 4.2. If $\mathcal{S} \subseteq L^2(\mathbb{Q}, \mathcal{G}_T)$ is a visible hedge class, then

$$\inf_{H \in \mathcal{S}} \mathbb{E}_{\mathbb{Q}}[(X - H)^2] = \text{dist}_{L^2(\mathbb{Q})}(L_X^{\mathbb{Q}}, \mathcal{S})^2 + \|U_X^{\mathbb{Q}}\|_{L^2(\mathbb{Q})}^2.$$

Equivalently,

$$\inf_{H \in \mathcal{S}} \mathbb{E}_{\mathbb{Q}}[(X - H)^2] = \|L_X^{\mathbb{Q}} - P_{\mathcal{S}}^{\mathbb{Q}} L_X^{\mathbb{Q}}\|_{L^2(\mathbb{Q})}^2 + \|U_X^{\mathbb{Q}}\|_{L^2(\mathbb{Q})}^2.$$

If $L_X^{\mathbb{Q}} = \alpha_X^{\mathbb{Q}} \Lambda_X$, this becomes

$$\inf_{H \in \mathcal{S}} \mathbb{E}_{\mathbb{Q}}[(X - H)^2] = (\alpha_X^{\mathbb{Q}})^2 \|\Lambda_X - P_{\mathcal{S}}^{\mathbb{Q}}(\Lambda_X)\|_{L^2(\mathbb{Q})}^2 + \|U_X^{\mathbb{Q}}\|_{L^2(\mathbb{Q})}^2.$$

Hence incompleteness is exactly the projection failure of the priced structural loading.

Proof. Fix $H \in \mathcal{S}$. Since $H \in \mathcal{S} \subseteq L^2(\mathbb{Q}, \mathcal{G}_T) \subseteq \mathcal{M}_{\mathbb{Q}}$ and $L_X^{\mathbb{Q}} \in \mathcal{M}_{\mathbb{Q}}$, one has $L_X^{\mathbb{Q}} - H \in \mathcal{M}_{\mathbb{Q}}$. Because $U_X^{\mathbb{Q}} \perp \mathcal{M}_{\mathbb{Q}}$, one has

$$\mathbb{E}_{\mathbb{Q}}[(L_X^{\mathbb{Q}} - H)U_X^{\mathbb{Q}}] = 0.$$

Therefore

$$\mathbb{E}_{\mathbb{Q}}[(X - H)^2] = \|L_X^{\mathbb{Q}} - H\|_{L^2(\mathbb{Q})}^2 + \|U_X^{\mathbb{Q}}\|_{L^2(\mathbb{Q})}^2.$$

The second term is independent of H . Taking the infimum over $\mathcal{S}$ gives the distance identity, and closedness of $\mathcal{S}$ identifies the minimizer as $P_S^{\mathbb{Q}}L_X^{\mathbb{Q}}$. If $L_X^{\mathbb{Q}} = \alpha_X^{\mathbb{Q}}\Lambda_X$, linearity of orthogonal projection gives $P_S^{\mathbb{Q}}L_X^{\mathbb{Q}} = \alpha_X^{\mathbb{Q}}P_S^{\mathbb{Q}}\Lambda_X$, which yields the scalar-factor formula. $\square$

Proposition 11.2 (One-factor visible hedging residual for market variance). *Assume Assumption 10.19. Let $\mathcal{S} \subseteq L^2(\mathbb{Q}, \mathcal{G}_T)$ be any closed visible hedge class. Then*

$$\inf_{H \in \mathcal{S}} \mathbb{E}_{\mathbb{Q}}[(Q_M - H)^2] = \text{dist}_{L^2(\mathbb{Q})}(V_Q, \mathcal{S})^2 + \Delta_Q^2,$$

where Δ_Q^2 is the hidden variance-residual constant from Proposition 10.20. In particular, if the visible hedge class is the entire visible space $\mathcal{S} = L^2(\mathbb{Q}, \mathcal{G}_T)$, then

$$\inf_{H \in L^2(\mathbb{Q}, \mathcal{G}_T)} \mathbb{E}_{\mathbb{Q}}[(Q_M - H)^2] = \Delta_Q^2.$$

If Ξ is symmetric, the full-visible hedge residual is

$$4a^2\mathbb{E}_{\mathbb{Q}}[A_G^2] + a^4B_f^2\text{Var}_{\mathbb{Q}}(\Xi^2).$$

If moreover $m \equiv 0$ and $\Xi \sim \mathcal{N}(0, 1)$, the residual is the explicit floor

$$2a^4B_f^2.$$

Proof. For any $H \in \mathcal{S} \subseteq L^2(\mathbb{Q}, \mathcal{G}_T)$, write

$$Q_M - H = (V_Q - H) + (2aA_G\Xi + a^2B_f(\Xi^2 - 1)).$$

The visible term $V_Q - H$ is orthogonal to the hidden centered term by the same independence and centering argument used in Proposition 10.20. Hence

$$\mathbb{E}_{\mathbb{Q}}[(Q_M - H)^2] = \|V_Q - H\|_{L^2(\mathbb{Q})}^2 + \Delta_Q^2.$$

Taking the infimum over $H \in \mathcal{S}$ gives the projection identity. The special cases follow from Proposition 10.20 and Corollary 10.21. $\square$

Remark 11.3 (Reading of the visible hedging residual). The abstract hedging theorem characterizes visible hedging failure as projection failure of the priced structural loading. Proposition 11.2 identifies that failure in a one-factor aggregate-variance model. The irreducible hidden piece is

$$2aA_G\Xi + a^2B_f(\Xi^2 - 1),$$

which no $\mathcal{G}_T$ -measurable hedge can span. The visible hedge class can improve only the visible component V_Q ; it cannot reduce Δ_Q^2 .

Corollary 11.4 (Factor-level decomposition of the visible-hedging residual). *Under the assumptions of Proposition 4.6 and 11.1, if $\mathcal{S} \subseteq L^2(\mathbb{Q}, \mathcal{G}_T)$ is a visible hedge class, then*

$$\inf_{H \in \mathcal{S}} \mathbb{E}_{\mathbb{Q}}[(X - H)^2] = \|V_X^{\mathbb{Q}} - P_S^{\mathbb{Q}}(V_X^{\mathbb{Q}})\|_{L^2(\mathbb{Q})}^2 + (\alpha_X^{\mathbb{Q},*})^2 + \|R_X^{\mathbb{Q},*}\|_{L^2(\mathbb{Q})}^2 + \|U_X^{\mathbb{Q}}\|_{L^2(\mathbb{Q})}^2.$$

If X is aggregate-factor aligned, so $R_X^{\mathbb{Q},*} = 0$, this becomes

$$\inf_{H \in \mathcal{S}} \mathbb{E}_{\mathbb{Q}}[(X - H)^2] = \|V_X^{\mathbb{Q}} - P_S^{\mathbb{Q}}(V_X^{\mathbb{Q}})\|_{L^2(\mathbb{Q})}^2 + (\alpha_X^{\mathbb{Q},*})^2 + \|U_X^{\mathbb{Q}}\|_{L^2(\mathbb{Q})}^2.$$

Thus the canonical centered aggregate factor contributes an exact irreducible term to the visible hedging error whenever $\alpha_X^{\mathbb{Q},*} \neq 0$.

Proof. By Proposition 4.6,

$$L_X^{\mathbb{Q}} = V_X^{\mathbb{Q}} + \alpha_X^{\mathbb{Q},*} \Lambda_M^* + R_X^{\mathbb{Q},*}.$$

Since $\Lambda_M^* \in \mathcal{M}_{\mathbb{Q}}^{\circ} = \mathcal{M}_{\mathbb{Q}} \ominus L^2(\mathbb{Q}, \mathcal{G}_T)$, the factor direction is orthogonal to $L^2(\mathbb{Q}, \mathcal{G}_T)$, hence to $\mathcal{S} \subseteq L^2(\mathbb{Q}, \mathcal{G}_T)$. Therefore

$$P_{\mathcal{S}}^{\mathbb{Q}}(\Lambda_M^*) = 0.$$

Also $R_X^{\mathbb{Q},*} \in \mathcal{M}_{\mathbb{Q}}^{\circ}$, so $R_X^{\mathbb{Q},*} \perp \mathcal{S}$ and therefore

$$P_{\mathcal{S}}^{\mathbb{Q}}(R_X^{\mathbb{Q},*}) = 0.$$

Hence

$$L_X^{\mathbb{Q}} - P_{\mathcal{S}}^{\mathbb{Q}}(L_X^{\mathbb{Q}}) = (V_X^{\mathbb{Q}} - P_{\mathcal{S}}^{\mathbb{Q}}(V_X^{\mathbb{Q}})) + \alpha_X^{\mathbb{Q},*} \Lambda_M^* + R_X^{\mathbb{Q},*}.$$

These three terms are mutually orthogonal for the same reason as in Corollary 10.17. Combining the resulting norm identity with Theorem 11.1 gives the claim. $\square$

Corollary 11.5 (Option first-variation claims inherit factor control of compression and hedging). *Assume the hypotheses of Theorem D.14. Fix a public sub- σ -field $\mathcal{H} \subseteq \mathcal{G}_T$, let $\mathcal{S} \subseteq L^2(\mathbb{Q}, \mathcal{G}_T)$ be a visible hedge class, and let*

$$V_{M,\tau}^{\mathbb{Q}} := C_{\mathcal{G}_T}^{\mathbb{Q}}(L_{M,\tau}^{\mathbb{Q}}).$$

Then

$$\|L_{M,\tau}^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}}(L_{M,\tau}^{\mathbb{Q}})\|_{L^2(\mathbb{Q})}^2 \geq \|V_{M,\tau}^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}}(V_{M,\tau}^{\mathbb{Q}})\|_{L^2(\mathbb{Q})}^2 + (|a_{M,\tau}^*| - \delta_{M,\tau}^*)_+^2,$$

where $(x)_+ := \max\{x, 0\}$. Consequently, if $\Phi: \mathbb{R} \rightarrow \mathbb{R}$ is convex and twice continuously differentiable with $\inf_x \Phi''(x) \geq \kappa > 0$, and if $\Phi(\Xi_{M,\tau}), \Phi(C_{\mathcal{H}}^{\mathbb{Q}}(\Xi_{M,\tau})) \in L^1(\mathbb{Q})$, then

$$\mathbb{E}_{\mathbb{Q}}[\Phi(\Xi_{M,\tau})] - \mathbb{E}_{\mathbb{Q}}[\Phi(C_{\mathcal{H}}^{\mathbb{Q}}(\Xi_{M,\tau}))] \geq \frac{\kappa}{2} \left(\|V_{M,\tau}^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}}(V_{M,\tau}^{\mathbb{Q}})\|_{L^2(\mathbb{Q})}^2 + (|a_{M,\tau}^*| - \delta_{M,\tau}^*)_+^2 \right).$$

Moreover,

$$\inf_{H \in \mathcal{S}} \mathbb{E}_{\mathbb{Q}}[(\Xi_{M,\tau} - H)^2] \geq \|V_{M,\tau}^{\mathbb{Q}} - P_{\mathcal{S}}^{\mathbb{Q}}(V_{M,\tau}^{\mathbb{Q}})\|_{L^2(\mathbb{Q})}^2 + (|a_{M,\tau}^*| - \delta_{M,\tau}^*)_+^2 + \|U_{M,\tau}^{\mathbb{Q}}\|_{L^2(\mathbb{Q})}^2.$$

If $\delta_{M,\tau}^* = 0$, both bounds become exact identities with $(|a_{M,\tau}^*| - \delta_{M,\tau}^*)_+^2 = (a_{M,\tau}^*)^2$.

Proof. Under Theorem D.14,

$$\tilde{L}_{M,\tau}^{\mathbb{Q}} = \alpha_{M,\tau}^{\mathbb{Q},*} \Lambda_M^* + R_{M,\tau}^{\mathbb{Q},*},$$

with

$$|\alpha_{M,\tau}^{\mathbb{Q},*} - a_{M,\tau}^*| \leq \delta_{M,\tau}^*, \quad \|R_{M,\tau}^{\mathbb{Q},*}\|_{L^2(\mathbb{Q})} \leq \delta_{M,\tau}^*.$$

Hence

$$|\alpha_{M,\tau}^{\mathbb{Q},*}| \geq (|a_{M,\tau}^*| - \delta_{M,\tau}^*)_+.$$

Corollary 10.17 gives

$$\|L_{M,\tau}^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}}(L_{M,\tau}^{\mathbb{Q}})\|_{L^2(\mathbb{Q})}^2 = \|V_{M,\tau}^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}}(V_{M,\tau}^{\mathbb{Q}})\|_{L^2(\mathbb{Q})}^2 + (\alpha_{M,\tau}^{\mathbb{Q},*})^2 + \|R_{M,\tau}^{\mathbb{Q},*}\|_{L^2(\mathbb{Q})}^2,$$

which implies the first lower bound. The convex-compression estimate is then Theorem 10.16 combined with that lower bound.

Likewise Corollary 11.4 yields

$$\inf_{H \in \mathcal{S}} \mathbb{E}_{\mathbb{Q}}[(\Xi_{M,\tau} - H)^2] = \|V_{M,\tau}^{\mathbb{Q}} - P_{\mathcal{S}}^{\mathbb{Q}}(V_{M,\tau}^{\mathbb{Q}})\|_{L^2(\mathbb{Q})}^2 + (\alpha_{M,\tau}^{\mathbb{Q},*})^2 + \|R_{M,\tau}^{\mathbb{Q},*}\|_{L^2(\mathbb{Q})}^2 + \|U_{M,\tau}^{\mathbb{Q}}\|_{L^2(\mathbb{Q})}^2,$$

which implies the visible-hedging lower bound. If $\delta_{M,\tau}^* = 0$, then Corollary D.15 gives

$$\alpha_{M,\tau}^{\mathbb{Q},*} = a_{M,\tau}^*, \quad R_{M,\tau}^{\mathbb{Q},*} = 0,$$

and the stated exact identities follow. $\square$

Corollary 11.6 (Visible-hedge-class residual). *Under the assumptions of Theorem 11.1,*

$$\inf_{H \in \mathcal{S}} \mathbb{E}_{\mathbb{Q}}[(X - H)^2] \geq \text{dist}_{L^2(\mathbb{Q})}(L_X^{\mathbb{Q}}, \mathcal{S})^2.$$

Visible markets therefore do not complete structural claim content whenever the priced structural loading does not belong to the visible hedge class.

Proof. Theorem 11.1 gives the displayed lower bound after dropping the nonnegative term $\|U_X^{\mathbb{Q}}\|_{L^2(\mathbb{Q})}^2$. If $L_X^{\mathbb{Q}} \notin \mathcal{S}$, the distance term is strictly positive by closedness of $\mathcal{S}$; the visible hedge class leaves a residual fiber in the priced loading. $\square$

Corollary 11.7 (Structural compression across input geometries). *Suppose the synchronized-perturbation contribution is screened in a visible summary family of the type treated in Theorem D.62, but $L_X^{\mathbb{Q}} \neq 0$ for a variance-linked claim X . Then:*

- (i) *the short-maturity pricing asymmetry of [2] persists at market scale;*
- (ii) *that short-maturity import is made either under a direct $\mathbb{Q}$ -Volterra specification with the stated roughness exponents, or under an equivalent change of measure whose induced drift correction is lower order at the short-maturity scale used in the skew expansion;*
- (iii) *public screening under $\mathbb{Q}$ remains consistent with the observable-screening logic of [3];*
- (iv) *the harmonized pairwise premium language of Section 9 states that the surviving $\mathbb{P}$ -side screened object of [4] need not be recoverable from the chosen visible $\mathbb{Q}$ -summary family, and that one-dimensional harmonized projected information already determines premium support exactly once the pairwise normalizers and signs are tracked;*
- (v) *visible prices and visible hedges remain compressions of structural claim content in the sense of Corollary 4.10 and Theorems 10.16 and 11.1; if the claim is aggregate-factor aligned in the sense of Proposition 4.6, then Corollaries 10.17 and 11.4 identify the canonical centered aggregate factor as an exact additive source of the compression and hedging residuals.*

Proof. Part (i) is the combination of Theorem D.62 with the threshold logic of [2]. Part (ii) is the short-maturity compatibility restriction needed when the threshold is read under $\mathbb{Q}$. Part (iii) follows from the separation between structural geometry and public visibility established in Sections 3 and 4. Part (iv) combines the existence of a nontrivial $\mathbb{P}$ -side screened object with the possibility of risk-neutral screening at the public summary level and then expresses the resulting gap on the harmonized projected basis of Section 9. Part (v) is the common compression and projection mechanism spelled out in Corollary 4.10, Theorem 10.16, and Theorem 11.1. $\square$

Corollary 11.8 (One structural loading, three public consequences). *Fix a claim $X \in L^2(\mathbb{Q}, \mathcal{F}_T)$ with priced structural loading $L_X^{\mathbb{Q}} = \alpha_X^{\mathbb{Q}} \Lambda_X$, fix a public sub- σ -field $\mathcal{H} \subseteq \mathcal{G}_T$, and let $\mathcal{S} \subseteq L^2(\mathbb{Q}, \mathcal{G}_T)$ be a visible hedge class. Let $\Phi: \mathbb{R} \rightarrow \mathbb{R}$ be convex and twice continuously differentiable with*

$$\inf_x \Phi''(x) \geq \underline{\kappa} > 0,$$

and assume $\Phi(X), \Phi(C_{\mathcal{H}}^{\mathbb{Q}}(X)) \in L^1(\mathbb{Q})$. Then the three public mechanisms treated here are generated by the same priced structural loading:

- (i) *every visible summary functional $\Psi_{\mathcal{H}}: L^2(\mathbb{Q}, \mathcal{H}) \rightarrow \mathbb{R}^m$ acts on*

$$\Psi_{\mathcal{H}}(C_{\mathcal{H}}^{\mathbb{Q}}(X)) = \Psi_{\mathcal{H}}(\alpha_X^{\mathbb{Q}} C_{\mathcal{H}}^{\mathbb{Q}}(\Lambda_X));$$

- (ii) *every strongly convex pricing functional leaves a compression gap bounded below by*

$$\frac{\underline{\kappa}}{2} (\alpha_X^{\mathbb{Q}})^2 \|\Lambda_X - C_{\mathcal{H}}^{\mathbb{Q}}(\Lambda_X)\|_{L^2(\mathbb{Q})}^2;$$

- (iii) *every visible hedge class leaves the exact residual*

$$\inf_{H \in \mathcal{S}} \mathbb{E}_{\mathbb{Q}}[(X - H)^2] = (\alpha_X^{\mathbb{Q}})^2 \|\Lambda_X - P_{\mathcal{S}}^{\mathbb{Q}}(\Lambda_X)\|_{L^2(\mathbb{Q})}^2 + \|U_X^{\mathbb{Q}}\|_{L^2(\mathbb{Q})}^2.$$

Thus options, variance compression, and visible hedging are linear, nonlinear, and orthogonal faces of the same public treatment of $L_X^{\mathbb{Q}} = \alpha_X^{\mathbb{Q}} \Lambda_X$. If one further resolves $L_X^{\mathbb{Q}}$ through Proposition 4.6, then Corollaries 10.17 and 11.4 show that the canonical centered aggregate factor contributes an exact additive term $(\alpha_X^{\mathbb{Q},})^2$ to the nonlinear compression gap and to the visible-hedging residual whenever X is aggregate-factor aligned.*

Proof. Part (i) is Corollary 4.10(iv). Part (ii) is the scalar-factor form of Theorem 10.16. Part (iii) is the scalar-factor form of Theorem 11.1. $\square$

12. SCOPE AND INTERPRETATION OF THE MAIN RESULTS

The body uses one Hilbert-space structure. A priced claim is first reduced to its structural loading $L_X^{\mathbb{Q}}$. Public option summaries, premium comparisons, variance claims, and visible hedges contain only the corresponding $\mathbb{Q}$ -compression or projection of that loading. The result is narrower than a complete pricing model, a calibration theorem, or a market-completion statement.

Each endpoint has its own transfer condition. The option branch requires a cumulant-class, Hermite, jump, Gaussian/Bachelier, or local Black-window hypothesis connecting the quoted coefficient to the compressed loading. The premium branch requires signed coefficient vectors on a common harmonized basis; unsigned projected information alone cannot give an unrestricted support law. The variance branch is exact for affine or explicitly reduced one-factor claims, while nonlinear variance-linked claims expose a Hermite hierarchy rather than a universal scalar beta. Black-implied-volatility transfer is local in maturity, moneyness, and vega.

Appendix C records the endpoint obstructions and exact scope conditions. The structural relation remains fixed: option prices, premium comparisons, variance-linked claims, and visible hedges are different public images of the same priced structural loading.

13. CONCLUSION

The $\mathbb{P}/\mathbb{Q}$ compression link is generated by the Perron state. Directed Volterra kernels, roughness exponents, signed edge support, and the normalized Perron mode are primitive geometric objects. Equivalent pricing changes weighted projections of claims, not that geometry. Finite-sample visibility weights and screened edge selections remain observational projections unless they are imposed structurally. The inherited object from [1] is the Perron-generated sigma-field $\mathcal{H}_{\varpi}$; for market variance it also supplies the centered increment

$$\Delta_{Q_M} = C_{\mathcal{H}_{\varpi}}(Q_M) - C_{\mathcal{G}_T}(Q_M).$$

The market-scale construction of [1] derives the synchronized aggregate state. This paper prices its public images under $\mathbb{Q}$.

Theorem 4.12 is the organizing result. Every treated public pricing or hedging object is generated from the $\mathbb{Q}$ -side priced structural loading $L_X^{\mathbb{Q}}$ by compression, projection, or measurable transformation. Option summaries encode conditional quote coefficients; premium results compare matched signed $\mathbb{P}/\mathbb{Q}$ projected loadings; variance claims expose nonlinear compression gaps; visible hedges leave projection residuals. Public prices and visible hedges do not recover the full latent state. They select the component that survives the specified risk-neutral compression.

TECHNICAL APPENDICES

The appendices collect the option-transfer calculations and Gaussian benchmark reductions used above. They keep the option-surface branch checkable and leave the body focused on the compression theorem, endpoint scope, and empirical diagnostic.

APPENDIX A. SUPPORTING NON-GAUSSIAN AND JUMP OPTION-TRANSFER VARIANTS

The option-transfer branch is stated through the compact cumulant-class theorem. This appendix records the Hermite, primitive non-Gaussian Volterra, and jump calculations that delimit the theorem's scope.

Assumption A.1 (Hermite-admissible local non-Gaussian perturbation). For every sufficiently small maturity τ , let

$$R_{\tau}^{\varepsilon} = X_{\tau} + \varepsilon Z_{\tau} + o_{L^2}(\varepsilon)$$

under $\mathbb{Q}$, where

$$X_{\tau} \sim N(0, v_{\tau}), \quad v_{\tau} = \vartheta \tau + o(\tau), \quad \vartheta > 0.$$

Set $Y_\tau := X_\tau / \sqrt{v_\tau}$. Assume $Z_\tau \in L^2(\mathbb{Q})$, $\mathbb{E}_\mathbb{Q}[Z_\tau] = 0$, and that its conditional projection onto $\sigma(X_\tau)$ admits the finite Hermite expansion

$$\mathbb{E}_\mathbb{Q}[Z_\tau | X_\tau] = \sum_{m=1}^{M_\tau} a_{m,\tau} \text{He}_m(Y_\tau) + r_\tau(Y_\tau),$$

where $M_\tau \geq 2$ and $r_\tau \perp \text{He}_m(Y_\tau)$ for $0 \leq m \leq M_\tau$. For the option-price statements below assume either $r_\tau = 0$, or that the displayed expansion holds in L^2 and the residual term is lower order on the moneyness window under consideration.

Lemma A.2 (Option-price derivative generated by the conditional projection). *Under Assumption A.1, for every strike coordinate k at which the baseline distribution is continuous,*

$$\partial_\varepsilon \mathbb{E}_\mathbb{Q}[(R_\tau^\varepsilon - k)_+] \big|_{\varepsilon=0} = \mathbb{E}_\mathbb{Q}[Z_\tau \mathbf{1}_{\{X_\tau > k\}}].$$

Writing $z = k / \sqrt{v_\tau}$, this equals

$$\phi(z) \sum_{m=1}^{M_\tau} a_{m,\tau} \text{He}_{m-1}(z) + \mathbb{E}_\mathbb{Q}[r_\tau(Y_\tau) \mathbf{1}_{\{Y_\tau > z\}}].$$

Proof. The directional derivative of $x \mapsto (x - k)_+$ is $\mathbf{1}_{\{x > k\}}$ at continuity points of the baseline law. The assumed L^2 -differentiability gives the first identity by dominated localization. Taking conditional expectation onto $\sigma(X_\tau)$, substituting the Hermite expansion, and using

$$\mathbb{E}[\text{He}_m(Y) \mathbf{1}_{\{Y > z\}}] = \phi(z) \text{He}_{m-1}(z), \quad Y \sim N(0, 1),$$

gives the second identity. $\square$

Theorem A.3 (Hermite option-price transfer and the cumulant obstruction). *Assume Assumption A.1 with $r_\tau = 0$. Let $\sigma_N(\tau, k; \varepsilon)$ be the Bachelier implied normal volatility, and let $\sigma_{N,0}(\tau) := \sqrt{v_\tau / \tau}$. Then, for fixed τ ,*

$$\partial_\varepsilon \sigma_N(\tau, k; \varepsilon) \big|_{\varepsilon=0} = \frac{1}{\sqrt{\tau}} \sum_{m=1}^{M_\tau} a_{m,\tau} \text{He}_{m-1}\left(\frac{k}{\sqrt{v_\tau}}\right).$$

Consequently the ATM normal-skew derivative is

$$\partial_k \partial_\varepsilon \sigma_N(\tau, k; \varepsilon) \big|_{k=0, \varepsilon=0} = \frac{1}{\sqrt{\tau} \sqrt{v_\tau}} \sum_{m=2}^{M_\tau} (m-1) a_{m,\tau} \text{He}_{m-2}(0).$$

The third-cumulant derivative is

$$\partial_\varepsilon \kappa_3^\mathbb{Q}(R_\tau^\varepsilon) \big|_{\varepsilon=0} = 6v_\tau a_{2,\tau}.$$

Hence

$$\begin{aligned} \partial_k \partial_\varepsilon \sigma_N(\tau, k; \varepsilon) \big|_{k=0, \varepsilon=0} &= \frac{\partial_\varepsilon \kappa_3^\mathbb{Q}(R_\tau^\varepsilon) \big|_{\varepsilon=0}}{6v_\tau^{3/2} \sqrt{\tau}} \\ &\quad + \frac{1}{\sqrt{\tau} \sqrt{v_\tau}} \sum_{\substack{m \geq 4 \\ m \leq M_\tau}} (m-1) a_{m,\tau} \text{He}_{m-2}(0), \end{aligned}$$

because the odd m -terms with $m \geq 3$ vanish at the money. Therefore the cumulant-only Bachelier transfer is exact if the conditional projection has no even Hermite components beyond He_2 , and it is asymptotically exact at scale b_τ if

$$\frac{1}{\sqrt{\tau} \sqrt{v_\tau}} \sum_{\substack{m \geq 4 \\ m \leq M_\tau}} (m-1) a_{m,\tau} \text{He}_{m-2}(0) = o(b_\tau).$$

Proof. At the baseline normal price, Bachelier vega is

$$\partial_\sigma C_N(\tau, k; \sigma_{N,0}) = \sqrt{\tau} \phi(k / \sqrt{v_\tau}).$$

Dividing the price derivative in Lemma A.2 by this vega gives the normal-volatility derivative. Differentiating in k and using $\frac{d}{dz} \text{He}_{m-1}(z) = (m-1) \text{He}_{m-2}(z)$ gives the ATM formula. Finally,

$$\partial_\varepsilon \kappa_3^\mathbb{Q}(X_\tau + \varepsilon Z_\tau)|_{\varepsilon=0} = 3\mathbb{E}_\mathbb{Q}[X_\tau^2 Z_\tau] = 3v_\tau \mathbb{E}_\mathbb{Q}[Y_\tau^2 \mathbb{E}_\mathbb{Q}[Z_\tau | Y_\tau]].$$

Since $Y^2 = \text{He}_2(Y) + 1$ and $\mathbb{E}_\mathbb{Q}[Z_\tau] = 0$, only the He_2 -coefficient survives, and $\mathbb{E}[\text{He}_2(Y)^2] = 2$. Thus the cumulant derivative is $6v_\tau a_{2,\tau}$. Substitution proves the decomposition, and the parity statement follows from $\text{He}_n(0) = 0$ for odd n . $\square$

Corollary A.4 (Second-chaos closure of the cumulant transfer). *If*

$$\mathbb{E}_\mathbb{Q}[Z_\tau | X_\tau] = a_{2,\tau} \text{He}_2(X_\tau / \sqrt{v_\tau})$$

for all sufficiently small τ , then the Bachelier ATM skew-cumulant transfer is exact:

$$\partial_k \partial_\varepsilon \sigma_N(\tau, k; \varepsilon)|_{k=0, \varepsilon=0} = \frac{\partial_\varepsilon \kappa_3^\mathbb{Q}(R_\tau^\varepsilon)|_{\varepsilon=0}}{6v_\tau^{3/2} \sqrt{\tau}}.$$

This is the continuous-return second-chaos branch generated by the Brownian-correlated Volterra leverage terms.

Proof. Under the displayed assumption all Hermite coefficients in Theorem A.3 vanish except the second one. The skew derivative in that theorem therefore reduces to the $m = 2$ term, while the cumulant identity in the same proof gives $\partial_\varepsilon \kappa_3^\mathbb{Q}(R_\tau^\varepsilon)|_{\varepsilon=0} = 6v_\tau a_{2,\tau}$. Substitution into the ATM Bachelier skew formula gives the displayed exact transfer. $\square$

Remark A.5 (Sharp meaning of the non-Gaussian limitation). A non-Gaussian option-price theorem cannot be universally reduced to the third cumulant: higher even Hermite projections of the return perturbation also affect ATM skew. The cumulant-class theorem is therefore the correct theorem for the branch where the second Hermite projection dominates, not a universal law for every non-Gaussian return family.

Assumption A.6 (Non-Gaussian Volterra leverage primitives). Fix a finite active channel set $\mathcal{E}^{\text{ng}}$. On $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{Q})$, let B be the option-return Brownian motion. For each source index j , let $B^{\perp,j}$ be a Brownian motion independent of B , and let $\tilde{J}^j$ be a compensated compound-Poisson martingale,

$$\tilde{J}_t^j = \sum_{m=1}^{N_t^j} \xi_m^j - \nu_j t \mathbb{E}_\mathbb{Q}[\xi_1^j],$$

where N^j is a Poisson process of intensity ν_j , the jumps ξ_m^j have finite fourth moment, and the families are mutually independent except through the common Brownian motion stated below. Let

$$L_t^j := \rho_j B_t + \bar{\rho}_j B_t^{\perp,j} + \chi_j \tilde{J}_t^j, \quad \rho_j \in [-1, 1], \quad \bar{\rho}_j \geq 0, \quad \rho_j^2 + \bar{\rho}_j^2 \leq 1, \quad \chi_j \geq 0.$$

Whenever $\chi_j > 0$ and the jump law is nondegenerate, L^j is non-Gaussian.

For each active channel $e \in \mathcal{E}^{\text{ng}}$, write $j(e)$ for its source index and assume

$$G_e(t, s) = d_e a_e(t, s)(t - s)^{H_e - 1/2}, \quad 0 \leq s < t \leq T,$$

where $d_e \in \mathbb{R}$, $H_e \in (0, 1)$, and a_e is continuous with $a_e(0, 0) > 0$ when $d_e \neq 0$. The baseline normal volatility σ_0 is deterministic, strictly positive near zero, and continuous. There are constants $C, t_0 > 0$ and $\eta \in (0, 1]$ such that, for $0 \leq r \leq t \leq t_0$,

$$|\sigma_0(t) - \sigma_0(0)| \leq Ct^\eta, \quad |\sigma_0(r) a_e(t, r) - \sigma_0(0) a_e(0, 0)| \leq Ct^\eta \quad (e \in \mathcal{E}^{\text{ng}}).$$

Define the predictable Volterra leverage field by

$$Y_t := \sum_{e \in \mathcal{E}^{\text{ng}}} \int_0^{t-} G_e(t, s) dL_s^{j(e)},$$

and, for a local synchronized perturbation parameter ε , define

$$R_\tau^\varepsilon := \int_0^\tau (\sigma_0(t) + \varepsilon Y_t) dB_t, \quad X_\tau := R_\tau^0, \quad Z_\tau := \int_0^\tau Y_t dB_t.$$

Assume $X_\tau^2 Z_\tau \in L^1(\mathbb{Q})$ for sufficiently small τ .

Definition A.7 (Primitive non-Gaussian Volterra coefficients). Under Assumption A.6, set

$$\vartheta := \sigma_0(0)^2, \quad \mathcal{E}_\rho^{\text{ng}} := \{e \in \mathcal{E}^{\text{ng}} : d_e \rho_{j(e)} \neq 0\}.$$

For $e \in \mathcal{E}_\rho^{\text{ng}}$, define

$$K_e := \frac{6\sigma_0(0)^2 d_e \rho_{j(e)} a_e(0,0)}{(H_e + 1/2)(H_e + 3/2)}.$$

Let

$$\mathcal{B}_J := \{H_e : e \in \mathcal{E}_\rho^{\text{ng}}\}.$$

For $\beta \in \mathcal{B}_J$, set

$$A_J(\beta) := \sum_{e \in \mathcal{E}_\rho^{\text{ng}} : H_e = \beta} \frac{d_e \rho_{j(e)} a_e(0,0)}{\sigma_0(0)(H_e + 1/2)(H_e + 3/2)} = \frac{1}{6\vartheta^{3/2}} \sum_{e \in \mathcal{E}_\rho^{\text{ng}} : H_e = \beta} K_e.$$

The primitive compressed support selected by the cumulant condition is

$$\text{supp}_J := \{\beta \in \mathcal{B}_J : A_J(\beta) \neq 0\},$$

with selected exponent

$$\beta_J^{\text{comp}} := \inf \text{supp}_J$$

when the support is nonempty.

Theorem A.8 (Primitive non-Gaussian Volterra cumulant expansion). *Assume Assumption A.6. Then*

$$v(\tau) := \text{Var}_{\mathbb{Q}}(X_\tau) = \vartheta\tau + O(\tau^{1+\eta}).$$

If $\mathcal{E}_\rho^{\text{ng}} = \emptyset$, then

$$\partial_\varepsilon \kappa_3^{\mathbb{Q}}(R_\tau^\varepsilon)|_{\varepsilon=0} = 0 \quad \text{for all sufficiently small } \tau.$$

If $\mathcal{E}_\rho^{\text{ng}} \neq \emptyset$, then, as $\tau \downarrow 0$,

$$\partial_\varepsilon \kappa_3^{\mathbb{Q}}(R_\tau^\varepsilon)|_{\varepsilon=0} = \sum_{e \in \mathcal{E}_\rho^{\text{ng}}} K_e \tau^{H_e+3/2} + O\left(\sum_{e \in \mathcal{E}_\rho^{\text{ng}}} |d_e \rho_{j(e)}| \tau^{H_e+3/2+\eta}\right).$$

Equivalently,

$$\partial_\varepsilon \kappa_3^{\mathbb{Q}}(R_\tau^\varepsilon)|_{\varepsilon=0} = 6\vartheta^{3/2} \sum_{\beta \in \mathcal{B}_J} A_J(\beta) \tau^{\beta+3/2} + O(\tau^{\beta_0+3/2+\eta}),$$

where $\beta_0 := \min \mathcal{B}_J$. In particular, if

$$\text{supp}_J \neq \emptyset \quad \text{and} \quad \beta_0 + \eta > \beta_J^{\text{comp}},$$

then the cumulant part of Assumption 6.5 holds with support supp_J , selected exponent β_J^{comp} , and remainder

$$o(\tau^{\beta_J^{\text{comp}}+3/2}).$$

A sufficient simpler condition is $A_J(\beta_0) \neq 0$; in that case $\beta_J^{\text{comp}} = \beta_0$, and the displayed remainder is automatically lower order.

Proof. The variance identity follows from Itô isometry:

$$v(\tau) = \int_0^\tau \sigma_0(t)^2 dt = \sigma_0(0)^2 \tau + O(\tau^{1+\eta}).$$

Since $R_\tau^\varepsilon = X_\tau + \varepsilon Z_\tau$, X_τ is centered Gaussian, and $\mathbb{E}_{\mathbb{Q}}[Z_\tau] = 0$,

$$\partial_\varepsilon \kappa_3^{\mathbb{Q}}(R_\tau^\varepsilon)|_{\varepsilon=0} = 3\mathbb{E}_{\mathbb{Q}}[X_\tau^2 Z_\tau].$$

Itô's product formula gives

$$\mathbb{E}_{\mathbb{Q}}[X_\tau^2 Z_\tau] = 2 \int_0^\tau \sigma_0(t) \mathbb{E}_{\mathbb{Q}}[X_t Y_t] dt.$$

Only the $\rho_j B$ -component of L^j has nonzero covariance with the return Brownian motion. The orthogonal Brownian components and the compensated jump martingales have zero contribution to $\mathbb{E}_{\mathbb{Q}}[X_t Y_t]$. Hence

$$\mathbb{E}_{\mathbb{Q}}[X_t Y_t] = \sum_{e \in \mathcal{E}_\rho^{\text{ng}}} d_e \rho_{j(e)} \int_0^t \sigma_0(r) a_e(t, r) (t - r)^{H_e - 1/2} dr.$$

The Hölder condition implies, uniformly over the finite active set,

$$\int_0^t \sigma_0(r) a_e(t, r) (t - r)^{H_e - 1/2} dr = \frac{\sigma_0(0) a_e(0, 0)}{H_e + 1/2} t^{H_e + 1/2} + O(t^{H_e + 1/2 + \eta}).$$

Multiplying by the outer factor $2\sigma_0(t)$, integrating in t , and multiplying by the third-cumulant factor 3 yields

$$3\mathbb{E}_{\mathbb{Q}}[X_\tau^2 Z_\tau] = \sum_{e \in \mathcal{E}_\rho^{\text{ng}}} \frac{6\sigma_0(0)^2 d_e \rho_{j(e)} a_e(0, 0)}{(H_e + 1/2)(H_e + 3/2)} \tau^{H_e + 3/2} + O\left(\sum_{e \in \mathcal{E}_\rho^{\text{ng}}} |d_e \rho_{j(e)}| \tau^{H_e + 3/2 + \eta}\right),$$

which is the stated expansion. Grouping equal powers gives the coefficient form. If the grouped support is nonempty and $\beta_0 + \eta > \beta_J^{\text{comp}}$, then every lower-order remainder, including remainders attached to powers whose leading coefficients cancel, is $o(\tau^{\beta_J^{\text{comp}} + 3/2})$. The final statement follows. $\square$

Remark A.9 (Role of the non-Gaussian jump component). The jump sources make the latent Volterra driver non-Gaussian. At first order in the synchronized volatility perturbation, the third cumulant is selected by the predictable covariance between the latent source and the return Brownian motion. Thus the compound-Poisson component changes the law and higher moments of the leverage field, while the leading first-order skew coefficient is carried by the Brownian-correlated channel. The pure rare-jump mechanism below gives a complementary class in which the jump component itself carries the leading third cumulant.

Assumption A.10 (Self-similar rare-jump option-facing return). Let $X_\tau = \int_0^\tau \sigma_0(t) dB_t$ be the same centered Gaussian baseline as above, with $\sigma_0(0) > 0$. For each channel $e \in \mathcal{E}^J$, fix $\beta_e > 0$, an intensity coefficient $q_e \geq 0$, and a jump-size law F_e on $\mathbb{R}$ with finite fourth moment and third moment

$$m_{3,e} := \int x^3 F_e(dx).$$

For a local intensity perturbation $\varepsilon \geq 0$, let N_e^ε be a Poisson random measure on $(0, T] \times \mathbb{R}$, independent of B , with compensator

$$\varepsilon q_e s^{\beta_e - 1} ds F_e(dx).$$

The rare-jump perturbation return is

$$J_\tau^\varepsilon := \sum_{e \in \mathcal{E}^J} \int_{(0, \tau] \times \mathbb{R}} s^{1/2} x (N_e^\varepsilon(ds, dx) - \varepsilon q_e s^{\beta_e - 1} ds F_e(dx)),$$

and

$$R_\tau^\varepsilon := X_\tau + J_\tau^\varepsilon.$$

Theorem A.11 (Exact rare-jump cumulant expansion). *Under Assumption A.10,*

$$\text{Var}_{\mathbb{Q}}(X_\tau) = \sigma_0(0)^2 \tau + o(\tau).$$

Moreover, for the right derivative at $\varepsilon = 0$ of the third cumulant of R_τ^ε ,

$$\partial_\varepsilon^+ \kappa_3^{\mathbb{Q}}(R_\tau^\varepsilon)|_{\varepsilon=0} = \sum_{e \in \mathcal{E}^J} \frac{q_e m_{3,e}}{\beta_e + 3/2} \tau^{\beta_e + 3/2}.$$

Thus the cumulant support is

$$\text{supp}_{J, \text{rare}} = \left\{ \beta : \sum_{e: \beta_e = \beta} \frac{q_e m_{3,e}}{\beta + 3/2} \neq 0 \right\},$$

and, with $\vartheta = \sigma_0(0)^2$, the corresponding coefficients are

$$A_{J,\text{rare}}(\beta) = \frac{1}{6\vartheta^{3/2}} \sum_{e:\beta_e=\beta} \frac{q_e m_{3,e}}{\beta + 3/2}.$$

There is no asymptotic cumulant remainder beyond grouping equal powers.

Proof. The baseline variance follows from Itô's isometry and continuity of σ_0 at zero. Cumulants of independent random variables add. The centered Gaussian variable X_τ has zero third cumulant, and the compensated Poisson integral has third cumulant equal to the integral of the third power of its jump amplitude against its compensator, the standard compound-Poisson cumulant identity [41, 42]. Therefore

$$\kappa_3^{\mathbb{Q}}(J_\tau^\varepsilon) = \varepsilon \sum_{e \in \mathcal{E}^J} \int_0^\tau \int_{\mathbb{R}} (s^{1/2}x)^3 q_e s^{\beta_e-1} ds F_e(dx),$$

which is exactly

$$\varepsilon \sum_{e \in \mathcal{E}^J} \frac{q_e m_{3,e}}{\beta_e + 3/2} \tau^{\beta_e+3/2}.$$

Taking the right derivative at $\varepsilon = 0$ gives the formula. $\square$

Remark A.12 (Use of the two primitive non-Gaussian mechanisms). The Volterra leverage class is the natural non-Gaussian extension of the option-facing dynamics in the Gaussian block. It supplies the cumulant power law from explicit primitive sources while preserving the continuous-return second-chaos structure used in the Bachelier transfer below. The rare-jump class is a separate primitive mechanism for the cumulant condition. It shows that the same power structure can also arise from genuinely jump-carried third cumulants, with exact remainder control; this is the local use of standard jump-process cumulant calculus [41, 42]. For large rare jumps, option-price first variations need not be determined only by the third cumulant, so the self-contained Bachelier transfer is stated for the continuous-return second-chaos leverage class.

Remark A.13 (Scope of the Poisson mechanisms). The Poisson terms above are local primitive mechanisms, not a change in the baseline Gaussian or risk-neutral compression model. Setting the jump amplitudes to zero recovers the Brownian Volterra leverage class. The compensated compound-Poisson term in Assumption A.6 is used only to show that genuinely non-Gaussian latent Volterra drivers can coexist with the same public Brownian leverage channel, while Assumption A.10 is used only to exhibit an exact rare-jump cumulant power law. Neither Poisson mechanism is invoked in the scalar Gaussian seam, the local maturity-anchor closure, the premium transfer, or the variance and hedging compression theorems unless it is explicitly referenced.

Assumption A.14 (Compensated rare-jump perturbation of a normal baseline). Fix a maturity τ . Let $X_\tau \sim N(0, v_\tau)$, $v_\tau > 0$. For a local intensity parameter $\varepsilon \geq 0$, let the perturbation be a compensated compound-Poisson jump measure on $\mathbb{R}$ with finite first absolute moment and finite third absolute moment. Its intensity measure is $\varepsilon \mu_\tau(dy)$, and

$$J_\tau^\varepsilon = \sum_{n=1}^{N_\tau^\varepsilon} Y_n - \varepsilon \int_{\mathbb{R}} y \mu_\tau(dy),$$

where N_τ^ε has intensity $\varepsilon \mu_\tau(\mathbb{R})$, the marks have law $\mu_\tau/\mu_\tau(\mathbb{R})$ when $\mu_\tau(\mathbb{R}) > 0$, and J_τ^ε is independent of X_τ . Set $R_\tau^\varepsilon := X_\tau + J_\tau^\varepsilon$.

Theorem A.15 (Exact rare-jump option-price first variation). *Under Assumption A.14, let*

$$C_0(k) := \mathbb{E}_{\mathbb{Q}}[(X_\tau - k)_+].$$

Then

$$\partial_\varepsilon \mathbb{E}_{\mathbb{Q}}[(R_\tau^\varepsilon - k)_+] \big|_{\varepsilon=0} = \int_{\mathbb{R}} [C_0(k - y) - C_0(k) + y C_0'(k)] \mu_\tau(dy).$$

Consequently the Bachelier implied-normal-volatility derivative is

$$\partial_\varepsilon \sigma_N(\tau, k; \varepsilon)|_{\varepsilon=0} = \frac{1}{\sqrt{\tau} \phi(k/\sqrt{v_\tau})} \int_{\mathbb{R}} [C_0(k-y) - C_0(k) + yC'_0(k)] \mu_\tau(dy).$$

The ATM normal-skew derivative is the exact tail-kernel expression

$$\partial_k \partial_\varepsilon \sigma_N(\tau, k; \varepsilon)|_{k=0, \varepsilon=0} = \frac{1}{\sqrt{\tau} \phi(0)} \int_{\mathbb{R}} \left[-\Phi\left(\frac{y}{\sqrt{v_\tau}}\right) + \frac{1}{2} + \frac{y}{\sqrt{v_\tau}} \phi(0) \right] \mu_\tau(dy).$$

Proof. To first order in intensity, either no jump occurs or one jump occurs. The compensation contributes the deterministic shift $-\varepsilon \int y \mu_\tau(dy)$, whose first-order price effect is $-\int y \mu_\tau(dy) \mathbb{Q}(X_\tau > k)$. Since $C'_0(k) = -\mathbb{Q}(X_\tau > k)$, the price derivative is

$$\int_{\mathbb{R}} \{ \mathbb{E}_{\mathbb{Q}}[(X_\tau + y - k)_+] - \mathbb{E}_{\mathbb{Q}}[(X_\tau - k)_+] + yC'_0(k) \} \mu_\tau(dy).$$

The identity $\mathbb{E}_{\mathbb{Q}}[(X_\tau + y - k)_+] = C_0(k - y)$ gives the first formula. Dividing by Bachelier vega $\sqrt{\tau} \phi(k/\sqrt{v_\tau})$ gives the implied-volatility derivative. Differentiating the price kernel in k and evaluating at zero gives

$$C'_0(-y) - C'_0(0) + yC''_0(0) = -\Phi(y/\sqrt{v_\tau}) + \frac{1}{2} + \frac{y}{\sqrt{v_\tau}} \phi(0),$$

which proves the ATM formula. $\square$

Corollary A.16 (Small-amplitude jumps recover the cumulant transfer). *Assume the jump measure is concentrated on $|y| \leq \delta_\tau \sqrt{v_\tau}$, where $\delta_\tau \downarrow 0$, and*

$$\int |y|^5 \mu_\tau(dy) < \infty.$$

Then

$$\begin{aligned} \partial_k \partial_\varepsilon \sigma_N(\tau, k; \varepsilon)|_{k=0, \varepsilon=0} &= \frac{1}{6v_\tau^{3/2} \sqrt{\tau}} \int y^3 \mu_\tau(dy) \\ &\quad + O\left(\frac{1}{v_\tau^{5/2} \sqrt{\tau}} \int |y|^5 \mu_\tau(dy)\right). \end{aligned}$$

Since

$$\partial_\varepsilon \kappa_3^{\mathbb{Q}}(R_\tau^\varepsilon)|_{\varepsilon=0} = \int y^3 \mu_\tau(dy),$$

the leading term is exactly the cumulant transfer.

Proof. For $u = y/\sqrt{v_\tau}$, Taylor expansion gives

$$-\Phi(u) + \frac{1}{2} + u\phi(0) = \frac{\phi(0)}{6} u^3 + O(|u|^5)$$

uniformly for $|u| \leq \delta_\tau$. Substitute this into the kernel of Theorem A.15 and divide by $\sqrt{\tau} \phi(0)$. The cumulant identity is the standard compound-Poisson cumulant formula for a compensated jump integral [41, 42]. $\square$

Proposition A.17 (Large jumps are not determined by the third cumulant). *Fix $v_\tau > 0$. There exist two finite jump intensity measures $\mu_\tau^{(1)}$ and $\mu_\tau^{(2)}$ with*

$$\int y^3 \mu_\tau^{(1)}(dy) = \int y^3 \mu_\tau^{(2)}(dy)$$

but with different ATM jump-skew derivatives in Theorem A.15. Hence no large-jump option-price transfer can depend only on the third cumulant.

Proof. The ATM jump-skew derivative is the integral of

$$K_v(y) := -\Phi(y/\sqrt{v}) + \frac{1}{2} + \frac{y}{\sqrt{v}} \phi(0)$$

against the jump measure, up to the fixed factor $1/(\sqrt{\tau} \phi(0))$. The function K_v is not a constant multiple of y^3 on $\mathbb{R}$. Hence there is a finite signed measure ν with $\int y^3 \nu(dy) = 0$

but $\int K_v(y)\nu(dy) \neq 0$. Adding a sufficiently small multiple of ν to any strictly positive finite measure with finite third moment gives two positive finite measures with equal third moment and different kernel integrals. $\square$

Remark A.18 (Jump-transfer scope). Rare or small jumps are primitive mechanisms for cumulant powers and, under the small-amplitude condition of Corollary A.16, also fall into the cumulant option-transfer branch. Large jumps have the exact transfer in Theorem A.15; the governing coefficient is the tail-kernel integral, not the third cumulant.

APPENDIX B. PREMIUM IMPLEMENTATION AND SUPPORT CLASSIFICATION DETAILS

This appendix collects the estimator, covariance-normalizer, finite-dimensional support, and unsigned-information results that support the premium section without serving as main-text theorem endpoints.

Definition B.1 (One-dimensional harmonized projected information). Fix an ordered pair (i, j) and suppose the harmonized reduced experiment under $\mu \in \{\mathbb{P}, \mathbb{Q}\}$ is one-dimensional on the common source-screened basis φ_{ij} . Let $\omega_{ij}^{\mu, \text{harm}} > 0$ denote the corresponding one-dimensional covariance normalizer on that basis. Define the harmonized one-dimensional projected information by

$$\mathcal{I}_{ij}^{\mu, \text{harm}} := \frac{(\gamma_{ij}^\mu)^2}{\omega_{ij}^{\mu, \text{harm}}}.$$

This is the one-dimensional harmonized counterpart of the projected information quantities used in the operational $\mathbb{P}/\mathbb{Q}$ classification template of [4].

Assumption B.2 (Estimated harmonized covariance normalizers). In the setup of Assumption 9.3, suppose moreover that for each ordered pair (i, j) in the genuinely one-dimensional harmonized regime there exist positive estimators

$$\hat{\omega}_{ij,n}^{\mathbb{P}, \text{harm}}, \quad \hat{\omega}_{ij,n}^{\mathbb{Q}, \text{harm}},$$

such that

$$\sqrt{n} \begin{pmatrix} \hat{\gamma}_{ij,n}^{\mathbb{P}} - \gamma_{ij}^{\mathbb{P}} \\ \hat{\gamma}_{ij,n}^{\mathbb{Q}} - \gamma_{ij}^{\mathbb{Q}} \\ \hat{\omega}_{ij,n}^{\mathbb{P}, \text{harm}} - \omega_{ij}^{\mathbb{P}, \text{harm}} \\ \hat{\omega}_{ij,n}^{\mathbb{Q}, \text{harm}} - \omega_{ij}^{\mathbb{Q}, \text{harm}} \end{pmatrix} \Rightarrow \mathcal{N}(0, \Omega_{ij}^{\text{info}})$$

for some positive semidefinite 4×4 matrix $\Omega_{ij}^{\text{info}}$.

Proposition B.3 (Empirical harmonized covariance normalizers from matched coefficient observations). Fix an ordered pair (i, j) in the genuinely one-dimensional harmonized regime and suppose the matched coefficient-observation vectors

$$(\hat{g}_{ij,d}^{\mathbb{P}}, \hat{g}_{ij,d}^{\mathbb{Q}}), \quad d = 1, \dots, n,$$

are i.i.d. across matched dates with finite fourth moments, with

$$\mathbb{E}[\hat{g}_{ij,d}^\mu] = \gamma_{ij}^\mu, \quad \text{Var}(\hat{g}_{ij,d}^\mu) = \omega_{ij}^{\mu, \text{harm}}, \quad \mu \in \{\mathbb{P}, \mathbb{Q}\}.$$

Define the empirical one-dimensional covariance normalizers by

$$\tilde{\omega}_{ij,n}^{\mu, \text{harm}} := \frac{1}{n} \sum_{d=1}^n (\hat{g}_{ij,d}^\mu - \hat{\gamma}_{ij,n}^\mu)^2, \quad \mu \in \{\mathbb{P}, \mathbb{Q}\}.$$

Then

$$\tilde{\omega}_{ij,n}^{\mu, \text{harm}} \rightarrow \omega_{ij}^{\mu, \text{harm}} \quad \text{in probability,} \quad \mu \in \{\mathbb{P}, \mathbb{Q}\},$$

and

$$\sqrt{n} \begin{pmatrix} \hat{\gamma}_{ij,n}^{\mathbb{P}} - \gamma_{ij}^{\mathbb{P}} \\ \hat{\gamma}_{ij,n}^{\mathbb{Q}} - \gamma_{ij}^{\mathbb{Q}} \\ \tilde{\omega}_{ij,n}^{\mathbb{P},\text{harm}} - \omega_{ij}^{\mathbb{P},\text{harm}} \\ \tilde{\omega}_{ij,n}^{\mathbb{Q},\text{harm}} - \omega_{ij}^{\mathbb{Q},\text{harm}} \end{pmatrix} \Rightarrow \mathcal{N}(0, \tilde{\Omega}_{ij}^{\text{info}}),$$

where

$$\tilde{\Omega}_{ij}^{\text{info}} := \text{Cov} \left(\begin{pmatrix} \hat{g}_{ij,1}^{\mathbb{P}} - \gamma_{ij}^{\mathbb{P}} \\ \hat{g}_{ij,1}^{\mathbb{Q}} - \gamma_{ij}^{\mathbb{Q}} \\ (\hat{g}_{ij,1}^{\mathbb{P}} - \gamma_{ij}^{\mathbb{P}})^2 - \omega_{ij}^{\mathbb{P},\text{harm}} \\ (\hat{g}_{ij,1}^{\mathbb{Q}} - \gamma_{ij}^{\mathbb{Q}})^2 - \omega_{ij}^{\mathbb{Q},\text{harm}} \end{pmatrix} \right).$$

Consequently Assumption B.2 holds with

$$\tilde{\omega}_{ij,n}^{\mu,\text{harm}} = \tilde{\omega}_{ij,n}^{\mu,\text{harm}}, \quad \Omega_{ij}^{\text{info}} = \tilde{\Omega}_{ij}^{\text{info}}.$$

Proof. Write

$$\tilde{\omega}_{ij,n}^{\mu,\text{harm}} := \frac{1}{n} \sum_{d=1}^n (\hat{g}_{ij,d}^{\mu} - \gamma_{ij}^{\mu})^2.$$

Then the sample-variance identity gives

$$\tilde{\omega}_{ij,n}^{\mu,\text{harm}} = \tilde{\omega}_{ij,n}^{\mu,\text{harm}} - (\hat{\gamma}_{ij,n}^{\mu} - \gamma_{ij}^{\mu})^2.$$

Since $\hat{\gamma}_{ij,n}^{\mu} - \gamma_{ij}^{\mu} = O_{\mathbb{P}}(n^{-1/2})$, we have

$$\sqrt{n} (\tilde{\omega}_{ij,n}^{\mu,\text{harm}} - \tilde{\omega}_{ij,n}^{\mu,\text{harm}}) = -\sqrt{n} (\hat{\gamma}_{ij,n}^{\mu} - \gamma_{ij}^{\mu})^2 = o_{\mathbb{P}}(1).$$

By the multivariate CLT applied to the i.i.d. vector

$$\begin{pmatrix} \hat{g}_{ij,d}^{\mathbb{P}} - \gamma_{ij}^{\mathbb{P}} \\ \hat{g}_{ij,d}^{\mathbb{Q}} - \gamma_{ij}^{\mathbb{Q}} \\ (\hat{g}_{ij,d}^{\mathbb{P}} - \gamma_{ij}^{\mathbb{P}})^2 - \omega_{ij}^{\mathbb{P},\text{harm}} \\ (\hat{g}_{ij,d}^{\mathbb{Q}} - \gamma_{ij}^{\mathbb{Q}})^2 - \omega_{ij}^{\mathbb{Q},\text{harm}} \end{pmatrix},$$

the sample averages built from $(\hat{\gamma}_{ij,n}^{\mathbb{P}}, \hat{\gamma}_{ij,n}^{\mathbb{Q}}, \tilde{\omega}_{ij,n}^{\mathbb{P},\text{harm}}, \tilde{\omega}_{ij,n}^{\mathbb{Q},\text{harm}})$ satisfy the stated joint CLT. Replacing $\tilde{\omega}$ by $\tilde{\omega}$ does not change the limit by the displayed $o_{\mathbb{P}}(1)$ estimate. $\square$

Theorem B.4 (Operational harmonized projected-information estimator and support classifier). *Under Assumptions 9.3 and B.2, define*

$$\hat{\mathcal{I}}_{ij,n}^{\mu,\text{harm}} := \frac{(\hat{\gamma}_{ij,n}^{\mu})^2}{\tilde{\omega}_{ij,n}^{\mu,\text{harm}}}, \quad \hat{\sigma}_{ij,n}^{\mu} := \text{sgn}(\hat{\gamma}_{ij,n}^{\mu}), \quad \mu \in \{\mathbb{P}, \mathbb{Q}\}.$$

We use the convention $\text{sgn}(0) = 0$. Then:

(i)

$$\hat{\mathcal{I}}_{ij,n}^{\mu,\text{harm}} \rightarrow \mathcal{I}_{ij}^{\mu,\text{harm}} \quad \text{in probability,} \quad \mu \in \{\mathbb{P}, \mathbb{Q}\}.$$

(ii) With

$$J_{ij}^{\text{info}} := \begin{pmatrix} \frac{2\gamma_{ij}^{\mathbb{P}}}{\omega_{ij}^{\mathbb{P},\text{harm}}} & 0 & -\frac{(\gamma_{ij}^{\mathbb{P}})^2}{(\omega_{ij}^{\mathbb{P},\text{harm}})^2} & 0 \\ 0 & \frac{2\gamma_{ij}^{\mathbb{Q}}}{\omega_{ij}^{\mathbb{Q},\text{harm}}} & 0 & -\frac{(\gamma_{ij}^{\mathbb{Q}})^2}{(\omega_{ij}^{\mathbb{Q},\text{harm}})^2} \end{pmatrix},$$

we have

$$\sqrt{n} \begin{pmatrix} \hat{\mathcal{I}}_{ij,n}^{\mathbb{P},\text{harm}} - \mathcal{I}_{ij}^{\mathbb{P},\text{harm}} \\ \hat{\mathcal{I}}_{ij,n}^{\mathbb{Q},\text{harm}} - \mathcal{I}_{ij}^{\mathbb{Q},\text{harm}} \end{pmatrix} \Rightarrow \mathcal{N}(0, J_{ij}^{\text{info}} \Omega_{ij}^{\text{info}} (J_{ij}^{\text{info}})^{\top}).$$

(iii) If

$$\hat{\Pi}_{ij,n}^{\text{info}} := \hat{\sigma}_{ij,n}^{\mathbb{Q}} \sqrt{\hat{\omega}_{ij,n}^{\mathbb{Q},\text{harm}} \hat{\mathcal{I}}_{ij,n}^{\mathbb{Q},\text{harm}}} - \hat{\sigma}_{ij,n}^{\mathbb{P}} \sqrt{\hat{\omega}_{ij,n}^{\mathbb{P},\text{harm}} \hat{\mathcal{I}}_{ij,n}^{\mathbb{P},\text{harm}}},$$

then

$$\hat{\Pi}_{ij,n}^{\text{info}} = \hat{\Pi}_{ij,n} \quad \text{identically.}$$

(iv) If $\gamma_{ij}^{\mu} \neq 0$, then

$$\hat{\sigma}_{ij,n}^{\mu} \rightarrow \sigma_{ij}^{\mu} \quad \text{in probability.}$$

(v) For every deterministic sequence $\eta_{\Pi,n} \downarrow 0$ satisfying

$$\sqrt{n} \eta_{\Pi,n} \rightarrow \infty,$$

we have

$$\Pi_{ij} = 0 \implies \Pr\left(|\hat{\Pi}_{ij,n}^{\text{info}}| < \eta_{\Pi,n}\right) \rightarrow 1,$$

and

$$\Pi_{ij} \neq 0 \implies \Pr\left(|\hat{\Pi}_{ij,n}^{\text{info}}| \geq \eta_{\Pi,n}\right) \rightarrow 1.$$

Thus the exact information-side premium-support laws admit consistent plug-in empirical versions on the matched cleaned $\mathbb{P}/\mathbb{Q}$ surface.

Proof. Part (i) follows from the continuous mapping theorem applied to the quotient map $(x, u) \mapsto x^2/u$, with $\hat{\omega}_{ij,n}^{\mu,\text{harm}} > 0$ by assumption. Part (ii) is the multivariate delta method applied to the map

$$(x_P, x_Q, u_P, u_Q) \mapsto \left(\frac{x_P^2}{u_P}, \frac{x_Q^2}{u_Q} \right),$$

whose Jacobian at the population point is J_{ij}^{info} .

For part (iii), use the definition of $\hat{\mathcal{I}}_{ij,n}^{\mu,\text{harm}}$ to write

$$\hat{\sigma}_{ij,n}^{\mu} \sqrt{\hat{\omega}_{ij,n}^{\mu,\text{harm}} \hat{\mathcal{I}}_{ij,n}^{\mu,\text{harm}}} = \hat{\sigma}_{ij,n}^{\mu} \sqrt{(\hat{\gamma}_{ij,n}^{\mu})^2} = \hat{\gamma}_{ij,n}^{\mu}, \quad \mu \in \{\mathbb{P}, \mathbb{Q}\},$$

by the sign convention. Subtracting the $\mathbb{P}$ -term from the $\mathbb{Q}$ -term yields $\hat{\Pi}_{ij,n}^{\text{info}} = \hat{\Pi}_{ij,n}$.

For part (iv), if $\gamma_{ij}^{\mu} \neq 0$, consistency of $\hat{\gamma}_{ij,n}^{\mu}$ implies sign stability, so $\hat{\sigma}_{ij,n}^{\mu} \rightarrow \sigma_{ij}^{\mu}$ in probability. For part (v), if $\Pi_{ij} = 0$, Proposition 9.10 and part (iii) give

$$\sqrt{n} \hat{\Pi}_{ij,n}^{\text{info}} = \sqrt{n} \hat{\Pi}_{ij,n} = O_{\mathbb{P}}(1),$$

so $|\hat{\Pi}_{ij,n}^{\text{info}}| < \eta_{\Pi,n}$ with probability tending to one because $\sqrt{n} \eta_{\Pi,n} \rightarrow \infty$. If $\Pi_{ij} \neq 0$, part (iii) and consistency of $\hat{\Pi}_{ij,n}$ yield

$$\hat{\Pi}_{ij,n}^{\text{info}} \rightarrow \Pi_{ij},$$

while $\eta_{\Pi,n} \rightarrow 0$, so eventually $|\hat{\Pi}_{ij,n}^{\text{info}}| \geq \eta_{\Pi,n}$ with probability tending to one. $\square$

Remark B.5 (Operational exact information-side support classification). Theorem B.4 is the estimator-side counterpart to Corollaries B.13 and B.14, while Corollary B.6 is the weak-dependence operational extension. The population exact-support laws carry more than algebraic content. Once the one-dimensional covariance normalizers are estimated on the matched cleaned surface, the same support conclusions can be implemented empirically either through the coefficient premium $\hat{\Pi}_{ij,n}$ or through the equivalent projected-information representation $\hat{\Pi}_{ij,n}^{\text{info}}$, and they remain valid under the matched weak-dependence extension once the long-run covariance is controlled.

Corollary B.6 (Weakly dependent matched-date extension of the operational information-side premium theorem). *Fix an ordered pair (i, j) in the genuinely one-dimensional harmonized regime and define*

$$Y_{i,j,d} := \begin{pmatrix} \hat{g}_{ij,d}^{\mathbb{P}} - \gamma_{ij}^{\mathbb{P}} \\ \hat{g}_{ij,d}^{\mathbb{Q}} - \gamma_{ij}^{\mathbb{Q}} \\ (\hat{g}_{ij,d}^{\mathbb{P}} - \gamma_{ij}^{\mathbb{P}})^2 - \omega_{ij}^{\mathbb{P},\text{harm}} \\ (\hat{g}_{ij,d}^{\mathbb{Q}} - \gamma_{ij}^{\mathbb{Q}})^2 - \omega_{ij}^{\mathbb{Q},\text{harm}} \end{pmatrix}.$$

Assume that $(Y_{ij,d})_{d \geq 1}$ is strictly stationary, strongly mixing with summable mixing coefficients, and has finite $(4 + \delta)$ -moments for some $\delta > 0$. Suppose moreover that the long-run covariance matrix

$$\Omega_{ij}^{\text{lr,info}} := \sum_{h \in \mathbb{Z}} \text{Cov}(Y_{ij,0}, Y_{ij,h})$$

is well defined and that the corresponding multivariate central limit theorem holds:

$$\sqrt{n} \begin{pmatrix} \hat{\gamma}_{ij,n}^{\mathbb{P}} - \gamma_{ij}^{\mathbb{P}} \\ \hat{\gamma}_{ij,n}^{\mathbb{Q}} - \gamma_{ij}^{\mathbb{Q}} \\ \tilde{\omega}_{ij,n}^{\mathbb{P},\text{harm}} - \omega_{ij}^{\mathbb{P},\text{harm}} \\ \tilde{\omega}_{ij,n}^{\mathbb{Q},\text{harm}} - \omega_{ij}^{\mathbb{Q},\text{harm}} \end{pmatrix} \Rightarrow \mathcal{N}(0, \Omega_{ij}^{\text{lr,info}}).$$

Then the conclusions of Proposition B.3 and Theorem B.4 continue to hold with $\tilde{\Omega}_{ij}^{\text{info}}$ and $\Omega_{ij}^{\text{info}}$ replaced by $\Omega_{ij}^{\text{lr,info}}$. In particular,

$$\sqrt{n} \begin{pmatrix} \hat{\mathcal{I}}_{ij,n}^{\mathbb{P},\text{harm}} - \mathcal{I}_{ij}^{\mathbb{P},\text{harm}} \\ \hat{\mathcal{I}}_{ij,n}^{\mathbb{Q},\text{harm}} - \mathcal{I}_{ij}^{\mathbb{Q},\text{harm}} \end{pmatrix} \Rightarrow \mathcal{N}\left(0, J_{ij}^{\text{info}} \Omega_{ij}^{\text{lr,info}} (J_{ij}^{\text{info}})^{\top}\right),$$

and the information-side premium-support classifier based on $\hat{\Pi}_{ij,n}^{\text{info}} = \hat{\Pi}_{ij,n}$ remains consistent for any deterministic threshold sequence $\eta_{\Pi,n} \downarrow 0$ with $\sqrt{n}\eta_{\Pi,n} \rightarrow \infty$.

Proof. The proof is the same as for Proposition B.3 and Theorem B.4, except that the i.i.d. multivariate CLT is replaced by the assumed long-run covariance CLT for the strictly stationary mixing array. The sample-variance identity

$$\tilde{\omega}_{ij,n}^{\mu,\text{harm}} = \bar{\omega}_{ij,n}^{\mu,\text{harm}} - (\hat{\gamma}_{ij,n}^{\mu} - \gamma_{ij}^{\mu})^2$$

still gives the $o_{\mathbb{P}}(1)$ replacement of $\bar{\omega}$ by $\tilde{\omega}$ after multiplication by $\sqrt{n}$, and the multivariate delta method applies unchanged. The identity $\hat{\Pi}_{ij,n}^{\text{info}} = \hat{\Pi}_{ij,n}$ is algebraic, so the support classification conclusion again follows from the asymptotic law of $\hat{\Pi}_{ij,n}$ and the threshold condition $\sqrt{n}\eta_{\Pi,n} \rightarrow \infty$. $\square$

Remark B.7 (Long-run covariance estimation for operational premium support). Corollary B.6 shows that the operational information-side premium-support theorem is not tied to an i.i.d. matched-date design. Under weak dependence, the same plug-in classifier is valid once the long-run covariance matrix is estimated consistently, for instance by a standard HAC procedure applied to the matched vector $Y_{ij,d}$ [43, 44].

Remark B.8 (Preferred operational reading and premium extension boundary). For implementation and empirical interpretation, Corollary B.6 is the preferred operational theorem. It contains Proposition B.3 and Theorem B.4 as the special case in which the long-run covariance reduces to the contemporaneous covariance, while keeping the premium section inside the genuinely one-dimensional harmonized regime solved here. Extending the exact information-side support law beyond that regime, for instance to low-rank or unrestricted matched local geometry under $\mathbb{Q}$, would require new normalizer and sign tracking and is outside the present theorem.

Proposition B.9 (One-dimensional harmonized projected information is coefficient-equivalent). *Fix an ordered pair (i, j) and suppose the harmonized reduced experiment under $\mu \in \{\mathbb{P}, \mathbb{Q}\}$ is one-dimensional on the common source-screened basis φ_{ij} . Then, with $\mathcal{I}_{ij}^{\mu,\text{harm}}$ defined by Definition B.1,*

$$\mathcal{I}_{ij}^{\mu,\text{harm}} = 0 \quad \Longleftrightarrow \quad \gamma_{ij}^{\mu} = 0 \quad \Longleftrightarrow \quad D_{ij}^{\mu} = 0.$$

If, moreover,

$$0 < \omega_{ij}^{\mu} \leq \omega_{ij}^{\mu,\text{harm}} \leq \bar{\omega}_{ij}^{\mu} < \infty,$$

then for every $\kappa_{\mu} > 0$,

$$\mathcal{I}_{ij}^{\mu,\text{harm}} \geq \kappa_{\mu} \quad \Longrightarrow \quad |\gamma_{ij}^{\mu}| \geq \sqrt{\kappa_{\mu} \omega_{ij}^{\mu}},$$

and

$$\mathcal{I}_{ij}^{\mu, \text{harm}} < \kappa_\mu \implies |\gamma_{ij}^\mu| < \sqrt{\kappa_\mu \bar{\omega}_{ij}^\mu}.$$

Proof. Since $\omega_{ij}^{\mu, \text{harm}} > 0$,

$$\mathcal{I}_{ij}^{\mu, \text{harm}} = 0 \iff \gamma_{ij}^\mu = 0.$$

The equivalence with $D_{ij}^\mu = 0$ is Proposition 9.17. The lower and upper bounds follow from

$$(\gamma_{ij}^\mu)^2 = \mathcal{I}_{ij}^{\mu, \text{harm}} \omega_{ij}^{\mu, \text{harm}}$$

and the assumed bounds on $\omega_{ij}^{\mu, \text{harm}}$. $\square$

Corollary B.10 (Separated harmonized projected information forces premium support). *Fix an ordered pair (i, j) and suppose the harmonized reduced experiments under $\mathbb{P}$ and $\mathbb{Q}$ are both one-dimensional on the common source-screened basis φ_{ij} , with positive covariance normalizers satisfying*

$$0 < \underline{\omega}_{ij}^\mathbb{P} \leq \omega_{ij}^{\mathbb{P}, \text{harm}} \leq \bar{\omega}_{ij}^\mathbb{P}, \quad 0 < \underline{\omega}_{ij}^\mathbb{Q} \leq \omega_{ij}^{\mathbb{Q}, \text{harm}} \leq \bar{\omega}_{ij}^\mathbb{Q}.$$

(i) If

$$\mathcal{I}_{ij}^{\mathbb{P}, \text{harm}} \geq \kappa_P, \quad \mathcal{I}_{ij}^{\mathbb{Q}, \text{harm}} < \kappa_Q,$$

and

$$\sqrt{\kappa_P \underline{\omega}_{ij}^\mathbb{P}} > \sqrt{\kappa_Q \bar{\omega}_{ij}^\mathbb{Q}},$$

then

$$|\Pi_{ij}| \geq \sqrt{\kappa_P \underline{\omega}_{ij}^\mathbb{P}} - \sqrt{\kappa_Q \bar{\omega}_{ij}^\mathbb{Q}} > 0,$$

so $(i, j) \in \mathcal{E}_\Pi$.

(ii) If

$$\mathcal{I}_{ij}^{\mathbb{Q}, \text{harm}} \geq \kappa_Q, \quad \mathcal{I}_{ij}^{\mathbb{P}, \text{harm}} < \kappa_P,$$

and

$$\sqrt{\kappa_Q \underline{\omega}_{ij}^\mathbb{Q}} > \sqrt{\kappa_P \bar{\omega}_{ij}^\mathbb{P}},$$

then

$$|\Pi_{ij}| \geq \sqrt{\kappa_Q \underline{\omega}_{ij}^\mathbb{Q}} - \sqrt{\kappa_P \bar{\omega}_{ij}^\mathbb{P}} > 0,$$

so $(i, j) \in \mathcal{E}_\Pi$.

In particular, separated projected-information regimes on the harmonized source-screened basis force nontrivial premium support.

Proof. Proposition B.9 converts the projected-information bounds into the coefficient bounds

$$|\gamma_{ij}^\mathbb{P}| \geq \sqrt{\kappa_P \underline{\omega}_{ij}^\mathbb{P}}, \quad |\gamma_{ij}^\mathbb{Q}| < \sqrt{\kappa_Q \bar{\omega}_{ij}^\mathbb{Q}}$$

in part (i), and the analogous bounds with $\mathbb{P}$ and $\mathbb{Q}$ exchanged in part (ii). Apply the reverse triangle inequality

$$|\Pi_{ij}| \geq ||\gamma_{ij}^\mathbb{Q}| - |\gamma_{ij}^\mathbb{P}||$$

to conclude. $\square$

Remark B.11 (Relation to the operational $\mathbb{P}/\mathbb{Q}$ classification template). Corollary B.10 is the premium-support analogue of the operational $\mathbb{P}/\mathbb{Q}$ classification template in [4]. In the genuinely one-dimensional harmonized regime, a path-detectable-only or pricing-visible-only information pattern on the common source-screened basis is no longer only a qualitative label. It forces membership in the premium support once the one-dimensional covariance normalizers are controlled.

Proposition B.12 (Exact premium reconstruction from harmonized projected information). *Fix an ordered pair (i, j) and suppose the harmonized reduced experiments under $\mathbb{P}$ and $\mathbb{Q}$ are both one-dimensional on the common source-screened basis φ_{ij} . Define the sign indicators*

$$\sigma_{ij}^\mu := \begin{cases} 1, & \gamma_{ij}^\mu > 0, \\ 0, & \gamma_{ij}^\mu = 0, \\ -1, & \gamma_{ij}^\mu < 0, \end{cases} \quad \mu \in \{\mathbb{P}, \mathbb{Q}\}.$$

Then

$$\gamma_{ij}^\mu = \sigma_{ij}^\mu \sqrt{\omega_{ij}^{\mu, \text{harm}} \mathcal{I}_{ij}^{\mu, \text{harm}}}, \quad \mu \in \{\mathbb{P}, \mathbb{Q}\},$$

and therefore

$$\Pi_{ij} = \sigma_{ij}^{\mathbb{Q}} \sqrt{\omega_{ij}^{\mathbb{Q}, \text{harm}} \mathcal{I}_{ij}^{\mathbb{Q}, \text{harm}}} - \sigma_{ij}^{\mathbb{P}} \sqrt{\omega_{ij}^{\mathbb{P}, \text{harm}} \mathcal{I}_{ij}^{\mathbb{P}, \text{harm}}}.$$

Proof. By Definition B.1,

$$(\gamma_{ij}^\mu)^2 = \omega_{ij}^{\mu, \text{harm}} \mathcal{I}_{ij}^{\mu, \text{harm}}.$$

Multiplying the positive square root by the sign indicator σ_{ij}^μ recovers γ_{ij}^μ . Substituting that representation into $\Pi_{ij} = \gamma_{ij}^{\mathbb{Q}} - \gamma_{ij}^{\mathbb{P}}$ gives the displayed formula. $\square$

Corollary B.13 (Exact premium-support law from harmonized projected information and normalizers). *Under the hypotheses of Proposition B.12,*

$$\Pi_{ij} = 0 \quad \Longleftrightarrow \quad \sigma_{ij}^{\mathbb{Q}} \sqrt{\omega_{ij}^{\mathbb{Q}, \text{harm}} \mathcal{I}_{ij}^{\mathbb{Q}, \text{harm}}} = \sigma_{ij}^{\mathbb{P}} \sqrt{\omega_{ij}^{\mathbb{P}, \text{harm}} \mathcal{I}_{ij}^{\mathbb{P}, \text{harm}}}.$$

Equivalently,

$$(i, j) \in \mathcal{E}_\Pi \quad \Longleftrightarrow \quad \sigma_{ij}^{\mathbb{Q}} \sqrt{\omega_{ij}^{\mathbb{Q}, \text{harm}} \mathcal{I}_{ij}^{\mathbb{Q}, \text{harm}}} \neq \sigma_{ij}^{\mathbb{P}} \sqrt{\omega_{ij}^{\mathbb{P}, \text{harm}} \mathcal{I}_{ij}^{\mathbb{P}, \text{harm}}}.$$

In particular, if $\sigma_{ij}^{\mathbb{Q}} = -\sigma_{ij}^{\mathbb{P}}$ and

$$\mathcal{I}_{ij}^{\mathbb{Q}, \text{harm}} + \mathcal{I}_{ij}^{\mathbb{P}, \text{harm}} > 0,$$

then $(i, j) \in \mathcal{E}_\Pi$.

Proof. The equivalence is exactly the identity in Proposition B.12. If $\sigma_{ij}^{\mathbb{Q}} = -\sigma_{ij}^{\mathbb{P}}$ and at least one information term is positive, the two signed square-root terms cannot cancel. $\square$

Corollary B.14 (Common-normalizer exact premium-support law from harmonized projected information). *Under the hypotheses of Proposition B.12, assume in addition that*

$$\omega_{ij}^{\mathbb{P}, \text{harm}} = \omega_{ij}^{\mathbb{Q}, \text{harm}} =: \omega_{ij}^{\text{harm}} > 0.$$

Then:

(i) if $\sigma_{ij}^{\mathbb{Q}} = \sigma_{ij}^{\mathbb{P}}$, then

$$|\Pi_{ij}| = \sqrt{\omega_{ij}^{\text{harm}}} \left| \sqrt{\mathcal{I}_{ij}^{\mathbb{Q}, \text{harm}}} - \sqrt{\mathcal{I}_{ij}^{\mathbb{P}, \text{harm}}} \right|,$$

so

$$\Pi_{ij} = 0 \quad \Longleftrightarrow \quad \mathcal{I}_{ij}^{\mathbb{Q}, \text{harm}} = \mathcal{I}_{ij}^{\mathbb{P}, \text{harm}}.$$

Consequently,

$$(i, j) \in \mathcal{E}_\Pi \quad \Longleftrightarrow \quad \mathcal{I}_{ij}^{\mathbb{Q}, \text{harm}} \neq \mathcal{I}_{ij}^{\mathbb{P}, \text{harm}}.$$

(ii) if $\sigma_{ij}^{\mathbb{Q}} = -\sigma_{ij}^{\mathbb{P}}$ and

$$\mathcal{I}_{ij}^{\mathbb{Q}, \text{harm}} + \mathcal{I}_{ij}^{\mathbb{P}, \text{harm}} > 0,$$

then

$$|\Pi_{ij}| = \sqrt{\omega_{ij}^{\text{harm}}} \left(\sqrt{\mathcal{I}_{ij}^{\mathbb{Q}, \text{harm}}} + \sqrt{\mathcal{I}_{ij}^{\mathbb{P}, \text{harm}}} \right) > 0,$$

so $(i, j) \in \mathcal{E}_\Pi$.

Proof. Under the common-normalizer hypothesis, Corollary B.13 simplifies to the displayed formulas. If $\sigma_{ij}^{\mathbb{Q}} = \sigma_{ij}^{\mathbb{P}}$, factor out the common sign and take absolute values. If $\sigma_{ij}^{\mathbb{Q}} = -\sigma_{ij}^{\mathbb{P}}$, the two signed square-root terms add in absolute value. $\square$

Remark B.15 (From support inclusion to an exact information-side support law). Corollary B.10 is the threshold separation theorem: a strict projected-information gap forces premium support without requiring exact coefficient matching. Corollary B.13 is the exact law in the genuinely one-dimensional harmonized regime once the covariance normalizers and coefficient signs are tracked explicitly, and Corollary B.14 is the common-normalizer reduction: there, premium support is characterized exactly by projected-information mismatch on the common basis when the two projected coefficients have the same sign, and sign opposition forces support whenever one side is nondegenerate.

Remark B.16 (What the premium support theorem does and does not identify). Section 9 is therefore no longer only a fixed-pair estimator theorem. Definition 9.13, Proposition 9.14, and Corollaries 9.15, 9.18, and B.10 identify structural subsets of the pairwise premium support directly from the $\mathbb{P}/\mathbb{Q}$ visibility separation on the harmonized source-screened basis, while Corollaries B.13 and B.14 give genuine if-and-only-if support laws in the one-dimensional harmonized regime once the covariance normalizers and coefficient signs are tracked, with a reduced expression in the common-normalizer case. What is *not* claimed here is a full unrestricted if-and-only-if description of $\mathcal{E}_{\Pi}$ without those extra one-dimensional normalization and sign primitives. Such a characterization would require a complete matched $\mathbb{Q}$ -side pairwise local geometry on the same harmonized basis, beyond the support inclusion supplied by separated pricing/path visibility or the exact information-side law available in the normalizer-controlled one-dimensional case.

Assumption B.17 (Finite-dimensional matched harmonized premium experiment). For a fixed ordered pair (i, j) , let

$$\Phi_{ij} = (\varphi_{ij,1}, \dots, \varphi_{ij,r_{ij}})$$

be a common source-screened basis used under both $\mathbb{P}$ and $\mathbb{Q}$. The reduced projected structural-loading representatives have coefficient vectors

$$\gamma_{ij}^{\mathbb{P}} = (\gamma_{ij,1}^{\mathbb{P}}, \dots, \gamma_{ij,r_{ij}}^{\mathbb{P}})^{\top}, \quad \gamma_{ij}^{\mathbb{Q}} = (\gamma_{ij,1}^{\mathbb{Q}}, \dots, \gamma_{ij,r_{ij}}^{\mathbb{Q}})^{\top}.$$

The finite-dimensional pairwise premium vector is

$$\Pi_{ij}^{\text{vec}} := \gamma_{ij}^{\mathbb{Q}} - \gamma_{ij}^{\mathbb{P}}.$$

Let W_{ij}^{μ} be the positive-definite covariance normalizer on the same basis under $\mu \in \{\mathbb{P}, \mathbb{Q}\}$. The signed normalized coefficient vector is

$$\theta_{ij}^{\mu} := (W_{ij}^{\mu})^{-1/2} \gamma_{ij}^{\mu}.$$

Definition B.18 (Finite-dimensional premium support). Under Assumption B.17, define

$$(i, j) \in \mathcal{E}_{\Pi}^{\text{vec}} \iff \Pi_{ij}^{\text{vec}} \neq 0.$$

At level $\eta \geq 0$, define

$$(i, j) \in \mathcal{E}_{\Pi}^{\text{vec}}(\eta) \iff \|\Pi_{ij}^{\text{vec}}\|_2 \geq \eta.$$

Theorem B.19 (Finite-dimensional if-and-only-if premium support law). *Under Assumption B.17,*

$$(i, j) \in \mathcal{E}_{\Pi}^{\text{vec}} \iff \gamma_{ij}^{\mathbb{Q}} \neq \gamma_{ij}^{\mathbb{P}}.$$

Equivalently, in signed normalized information coordinates,

$$(i, j) \in \mathcal{E}_{\Pi}^{\text{vec}} \iff (W_{ij}^{\mathbb{Q}})^{1/2} \theta_{ij}^{\mathbb{Q}} \neq (W_{ij}^{\mathbb{P}})^{1/2} \theta_{ij}^{\mathbb{P}}.$$

Moreover,

$$\|\Pi_{ij}^{\text{vec}}\|_2^2 = \|\gamma_{ij}^{\mathbb{Q}}\|_2^2 + \|\gamma_{ij}^{\mathbb{P}}\|_2^2 - 2\langle \gamma_{ij}^{\mathbb{Q}}, \gamma_{ij}^{\mathbb{P}} \rangle.$$

Thus a complete finite-dimensional support decision requires the relative orientation of the two coefficient vectors as well as their separate squared lengths.

Proof. The first equivalence is the definition of Π_{ij}^{vec} . The normalized-coordinate form follows from $\gamma_{ij}^\mu = (W_{ij}^\mu)^{1/2} \theta_{ij}^\mu$. The norm identity is the polarization identity in $\mathbb{R}^{r_{ij}}$. $\square$

Corollary B.20 (Operational finite-dimensional premium estimator). *Suppose estimators satisfy the joint central limit theorem*

$$\sqrt{n} \begin{pmatrix} \hat{\gamma}_{ij,n}^{\mathbb{P}} - \gamma_{ij}^{\mathbb{P}} \\ \hat{\gamma}_{ij,n}^{\mathbb{Q}} - \gamma_{ij}^{\mathbb{Q}} \end{pmatrix} \Rightarrow N(0, \Sigma_{ij}),$$

where Σ_{ij} is a $2r_{ij} \times 2r_{ij}$ covariance matrix. Then

$$\hat{\Pi}_{ij,n}^{\text{vec}} := \hat{\gamma}_{ij,n}^{\mathbb{Q}} - \hat{\gamma}_{ij,n}^{\mathbb{P}}$$

is consistent and

$$\sqrt{n}(\hat{\Pi}_{ij,n}^{\text{vec}} - \Pi_{ij}^{\text{vec}}) \Rightarrow N(0, D\Sigma_{ij}D^\top), \quad D = (-I_{r_{ij}}, I_{r_{ij}}).$$

If $\Pi_{ij}^{\text{vec}} \neq 0$, then

$$\sqrt{n}(\|\hat{\Pi}_{ij,n}^{\text{vec}}\|_2 - \|\Pi_{ij}^{\text{vec}}\|_2) \Rightarrow N\left(0, \frac{(\Pi_{ij}^{\text{vec}})^\top D\Sigma_{ij}D^\top \Pi_{ij}^{\text{vec}}}{\|\Pi_{ij}^{\text{vec}}\|_2^2}\right).$$

Proof. The vector premium estimator is a linear map of the joint estimator, so the first limit follows from the multivariate delta method with derivative D . Away from zero, the Euclidean norm is differentiable with gradient $x/\|x\|_2$, giving the second limit. $\square$

Proposition B.21 (Unsigned information cannot give unrestricted support). *No unrestricted if-and-only-if premium-support law can be written using only the two unsigned projected information levels*

$$\|\theta_{ij}^{\mathbb{P}}\|_2^2, \quad \|\theta_{ij}^{\mathbb{Q}}\|_2^2,$$

even when the normalizers are known.

Proof. It is enough to take $r_{ij} = 1$ and common normalizer $W = 1$. The two coefficient pairs

$$(\gamma^{\mathbb{P}}, \gamma^{\mathbb{Q}}) = (1, 1), \quad (\gamma^{\mathbb{P}}, \gamma^{\mathbb{Q}}) = (1, -1)$$

have the same unsigned information levels under $\mathbb{P}$ and $\mathbb{Q}$, namely 1 and 1. In the first case $\Pi = 0$; in the second case $\Pi = -2$. Hence unsigned information alone cannot distinguish no support from nontrivial support. In higher dimension the same obstruction appears as an unknown relative angle between $\gamma^{\mathbb{P}}$ and $\gamma^{\mathbb{Q}}$. $\square$

Remark B.22 (Finite-dimensional premium scope). The one-dimensional exact law above is the first identifiable case using scalar projected information plus normalizers and signs. Theorem B.19 is the natural unrestricted statement once the full signed matched coefficient vector is part of the data contract. Proposition B.21 explains why a purely information-magnitude theorem cannot be unrestricted.

Remark B.23 (Implementation data contract). The theorem above is a compression theorem with an explicit data contract. To instantiate it one needs common observation dates, a maturity grid, a strike or moneyness grid, an arbitrage-regularized option surface on each date, a local derivative extraction map on that cleaned surface, and the same harmonized source-screened projection basis as on the realized-volatility side. Market microstructure is outside the theorem. The matched $\mathbb{P}/\mathbb{Q}$ comparison is therefore specified as a data contract rather than only as an abstract possibility.

APPENDIX C. DETAILED SCOPE CONDITIONS AND ENDPOINT OBSTRUCTIONS

The endpoint theorems are conditional. Each endpoint requires its own regularity, model-class, or data-alignment assumptions. The material below records how those assumptions are used. It adds no primitive model. Its role is to state when risk-neutral compression can be converted into a third-cumulant skew law, a Hermite price law, a jump-kernel price law, a local Black surface law, a signed premium comparison, or a nonlinear variance-linked Hermite hierarchy.

The assumptions have four roles. The Perron reduction assumptions import the synchronized state and the Volterra market geometry; in the empirical section that reduction is rebuilt from raw Oxford–Man data. The risk-neutral regularity assumptions make state-price-density changes, $L^2(\mathbb{Q})$ projections, and option first variations well-defined. The option-transfer assumptions specify which surface coordinate is read from the compressed loading: a cumulant-class skew, a Hermite price coefficient, an exact jump-kernel coefficient, or a local Black-implied-volatility coefficient. The premium, variance, and hedging assumptions state how the same loading is compared on a signed $\mathbb{P}/\mathbb{Q}$ basis or passed through nonlinear and projection-residual functionals.

The Gaussian Volterra branch has a separate role. Assumptions D.52 and D.55 are not empirical calibration knobs; they connect the body to Companion A. Companion A proves the raw Volterra estimates and source-family verification results. The article uses only the scalar seam, the common-channel reduction, and the Gaussian option-factor consequence. This separation keeps the body focused on risk-neutral compression and the option-surface diagnostic.

C.1. Non-Gaussian option pricing: cumulant scope and obstructions. The non-Gaussian option branch has two levels. The abstract compression theorem is model-agnostic. The scalar third-cumulant formula is not. The latter is the correct endpoint only when the Hermite and jump conditions leave no leading term outside the third-cumulant channel. This keeps the skew language close to the option skew-law literature [38] and makes explicit the extra compression and quote-transfer conditions needed here.

Definition C.1 (Hermite ATM-skew defect). Under Assumption A.1 with $r_\tau = 0$, define

$$\mathfrak{h}_\tau := \frac{1}{\sqrt{\tau}\sqrt{v_\tau}} \sum_{\substack{m \geq 4 \\ m \leq M_\tau}} (m-1)a_{m,\tau} \text{He}_{m-2}(0).$$

Then Theorem A.3 says that

$$\partial_k \partial_\varepsilon \sigma_N(\tau, k; \varepsilon) \big|_{k=0, \varepsilon=0} = \frac{\partial_\varepsilon \kappa_3^\mathbb{Q}(R_\tau^\varepsilon) \big|_{\varepsilon=0}}{6v_\tau^{3/2}\sqrt{\tau}} + \mathfrak{h}_\tau.$$

The cumulant transfer is exact when $\mathfrak{h}_\tau = 0$, and it is asymptotically exact at scale b_τ when $\mathfrak{h}_\tau = o(b_\tau)$.

Proposition C.2 (Same third cumulant, different ATM skew). Fix $v_\tau > 0$ and let $Y_\tau = X_\tau/\sqrt{v_\tau} \sim N(0, 1)$. For any $a_2 \in \mathbb{R}$ and any nonzero $a_4 \in \mathbb{R}$, the two conditional projections

$$\mathbb{E}_\mathbb{Q}[Z_\tau^{(1)} \mid X_\tau] = a_2 \text{He}_2(Y_\tau), \quad \mathbb{E}_\mathbb{Q}[Z_\tau^{(2)} \mid X_\tau] = a_2 \text{He}_2(Y_\tau) + a_4 \text{He}_4(Y_\tau)$$

have the same third-cumulant derivative but different Bachelier ATM-skew derivatives. Consequently the third cumulant alone cannot determine non-Gaussian ATM skew even in the finite Hermite-admissible class.

Proof. By Theorem A.3, the third-cumulant derivative is $6v_\tau a_2$ for both perturbations, since it depends only on the He_2 -coefficient. The ATM-skew derivative differs by

$$\frac{3a_4 \text{He}_2(0)}{\sqrt{\tau}\sqrt{v_\tau}} = -\frac{3a_4}{\sqrt{\tau}\sqrt{v_\tau}},$$

because $\text{He}_2(0) = -1$. This is nonzero when $a_4 \neq 0$. $\square$

Definition C.3 (Cumulant-valid continuous perturbations). Suppose the cumulant term in Definition C.1 has selected scale b_τ . The continuous non-Gaussian perturbation is cumulant-valid at that scale if

$$\mathfrak{h}_\tau = o(b_\tau).$$

If $\mathfrak{h}_\tau$ is of the same or larger order, the cumulant reduction is not valid at the selected scale. In that case the option-price endpoint is still explicit, but the selected coefficient is the Hermite sum, not the third cumulant alone.

Definition C.4 (Jump tail-kernel defect). For a jump intensity measure μ_τ in Theorem A.15, set

$$K_{v_\tau}(y) := -\Phi(y/\sqrt{v_\tau}) + \frac{1}{2} + \frac{y}{\sqrt{v_\tau}}\phi(0),$$

and define the exact jump ATM-skew contribution, up to the common factor $(\sqrt{\tau}\phi(0))^{-1}$, by

$$\mathfrak{j}_\tau := \int K_{v_\tau}(y) \mu_\tau(dy).$$

The small-jump cumulant approximation replaces $K_{v_\tau}(y)$ by $\phi(0)y^3/(6v_\tau^{3/2})$. Its residual is

$$\mathfrak{r}_\tau^J := \int \left(K_{v_\tau}(y) - \frac{\phi(0)}{6v_\tau^{3/2}}y^3 \right) \mu_\tau(dy).$$

The jump branch is cumulant-valid when $\mathfrak{r}_\tau^J$ is lower order at the selected scale. Otherwise the exact coefficient is $\mathfrak{j}_\tau$, not the third cumulant.

Remark C.5 (Correct non-Gaussian formulation). The formulation is therefore not “non-Gaussian skew is always third-cumulant skew.” Rather, the third cumulant is the option-price coefficient on the cumulant-class branch, i.e. when higher even Hermite projections and large-jump tail-kernel effects are absent or lower order. Outside that branch there is still an option pricing statement, but the coefficient is the visible Hermite or jump-kernel coefficient.

C.2. Premium recovery: what signed vectors identify and unsigned scalars do not.

The finite-dimensional premium theorem is a genuine strengthening because it moves from a scalar support statement to a signed matched vector statement. The price of that strengthening is a data contract: without the signed relative orientation of the $\mathbb{P}$ - and $\mathbb{Q}$ -coefficient vectors, there is no unrestricted support law.

Proposition C.6 (Sharp unsigned premium interval). *Let $u, v \in \mathbb{R}^r$ be two possible signed coefficient vectors and suppose that only their lengths $p := \|u\|_2$ and $q := \|v\|_2$ are known. Then every possible premium magnitude $\|v - u\|_2$ lies in*

$$[|p - q|, p + q].$$

If $r \geq 2$, every value in this interval is attained by a suitable relative angle between u and v . If $r = 1$, the only possible values are $|p - q|$ and $p + q$, corresponding to the two relative signs.

Proof. The bounds are the reverse triangle inequality and the triangle inequality. For $r \geq 2$, choose unit vectors e_1, e_2 and write

$$u = pe_1, \quad v = q(\cos \theta e_1 + \sin \theta e_2).$$

Then

$$\|v - u\|_2^2 = p^2 + q^2 - 2pq \cos \theta,$$

which ranges continuously from $(p - q)^2$ to $(p + q)^2$. In dimension one only $\cos \theta \in \{-1, 1\}$ is available. $\square$

Corollary C.7 (Unsigned information is one-sided). *In the finite-dimensional premium experiment, unequal known signed-vector lengths can force premium support, because $\|\gamma_{ij}^{\mathbb{Q}} - \gamma_{ij}^{\mathbb{P}}\|_2 \geq \left| \|\gamma_{ij}^{\mathbb{Q}}\|_2 - \|\gamma_{ij}^{\mathbb{P}}\|_2 \right|$. Equal lengths do not rule out premium support. Thus unsigned projected information can give sufficient separation tests, but not an unrestricted if-and-only-if support law.*

Proof. The lower bound is Proposition C.6. When the two lengths are equal and positive, the same proposition permits both zero premium, obtained by equal signed vectors, and nonzero premium, obtained by rotating or changing sign. This is exactly the nonidentification mechanism of Proposition B.21. $\square$

Remark C.8 (Premium data contract). The unrestricted finite-dimensional support theorem therefore requires the signed matched vector $\Pi_{ij}^{\text{vec}} = \gamma_{ij}^{\mathbb{Q}} - \gamma_{ij}^{\mathbb{P}}$. Scalar or unsigned information may be used for conservative screening, but the full support decision requires the relative orientation term

$$\langle \gamma_{ij}^{\mathbb{Q}}, \gamma_{ij}^{\mathbb{P}} \rangle$$

in the polarization identity of Theorem B.19. This orientation term is the exact obstruction; removability would require additional structure.

C.3. Nonlinear variance branch and scalar-factor scope. The nonlinear variance-linked branch is a nonlinear theorem with model-class scope. The exact object under Assumption 10.24 is the Hermite coefficient sequence of Proposition 10.25. A scalar one-factor representation is a special collapse of that sequence rather than the generic law.

Definition C.9 (Exact visible-beta one-factor collapse). Under Assumption 10.24, the nonlinear claim $X = \Psi(Q_M)$ has an exact visible-beta one-factor collapse if there exists a $\mathcal{G}$ -measurable random variable $b(V, S) \in L^2(\mathbb{Q})$ such that

$$L_X^{\mathbb{Q}} - C_{\mathcal{G}}^{\mathbb{Q}}(L_X^{\mathbb{Q}}) = b(V, S)Z.$$

This is weaker than a deterministic-beta scalar law, because the coefficient may depend on visible state variables. It is still stronger than the general Hermite hierarchy.

Proposition C.10 (Hermite criterion for exact one-factor collapse). *Under Assumption 10.24, the exact visible-beta one-factor collapse in Definition C.9 holds if and only if*

$$c_m(V, S) = 0 \quad \text{for every } m \geq 2$$

for the Hermite coefficients of Proposition 10.25. Equivalently, for almost every realized visible state (V, S) , the function

$$z \mapsto P_S \Psi(V + \alpha z)$$

is affine as an element of L^2 under the standard normal law. In that case $b(V, S) = c_1(V, S)$.

Proof. Proposition 10.25 gives

$$L_X^{\mathbb{Q}} - C_{\mathcal{G}}^{\mathbb{Q}}(L_X^{\mathbb{Q}}) = \sum_{m \geq 1} c_m(V, S) \text{He}_m(Z).$$

The Hermite polynomials are an orthogonal basis of L^2 under the standard normal law. Hence the right-hand side belongs to the span of $\text{He}_1(Z) = Z$ if and only if all coefficients of He_m , $m \geq 2$, vanish. Completeness of the Hermite basis gives the equivalent affine statement for the conditional function of z . $\square$

Corollary C.11 (Why the affine branch is the exact scalar branch). *If $S = 0$, $\alpha > 0$, and Ψ is non-affine on the support reached by $V + \alpha Z$, then the exact visible-beta one-factor collapse fails. The leading first-factor approximation of Corollary 10.26 may still hold for small α , but the exact law is the Hermite hierarchy.*

Proof. When $S = 0$, the conditional function is $z \mapsto \Psi(V + \alpha z)$. If it were affine in z for almost every visible state, Ψ would be affine on the corresponding translated support intervals. This is the obstruction isolated in Corollary 10.28. If Ψ is not affine there, at least one higher Hermite coefficient is nonzero by Proposition C.10. $\square$

Remark C.12 (Variance-branch scope). The conditionally Gaussian theorem is therefore a hierarchy theorem with a first-factor asymptotic, not a universal exact one-factor scalar theorem. Outside Assumption 10.24, the valid statement is the general compression and hedging residual statement. Inside that assumption, the exact nonlinear object is the full sequence $(c_m(V, S))_{m \geq 0}$.

C.4. Local Black surface transfer. The Black-implied-volatility conversion is local. It preserves visibility on a regular window where the price-to-volatility inverse is smooth and uniformly controlled. It does not claim a global implied-volatility surface theorem. This is the conservative local analogue of the Black and local-surface conventions [33, 35].

Definition C.13 (Black-window regularity condition). A Black-window regularity condition at scale a_τ consists of

$$(\mathcal{D}, m, a_\tau, \underline{v}_B, K_{m+1}),$$

where $\mathcal{D} \subset (0, T] \times \mathbb{R}$ is compact, $m \geq 0$,

$$\inf_{(\tau, k) \in \mathcal{D}} V_B(\tau, k) \geq \underline{v}_B > 0,$$

the inverse Black map has C^{m+1} -norm bounded by K_{m+1} on a neighborhood of $C^0(\mathcal{D})$, and the price expansion satisfies

$$C^\varepsilon = C^0 + \varepsilon \dot{C} + \rho^\varepsilon, \quad \|\rho^\varepsilon\|_{C^m(\mathcal{D})} = o(\varepsilon a_\tau).$$

This condition is exactly the operational content of Assumption E.4.

Proposition C.14 (Why the Black theorem is not global). *A scale-uniform price remainder does not imply a scale-uniform implied-volatility remainder without a vega lower bound. In particular, if along a sequence of points the Black vega satisfies $V_{B,n} \downarrow 0$, then there are price remainders $r_n = o(a_n)$ for which*

$$\frac{r_n}{V_{B,n}}$$

is not $o(a_n)$. Thus the local vega condition in Assumption E.4 is the condition that makes the Taylor transfer scale-uniform, rather than a cosmetic regularity assumption.

Proof. Take $V_{B,n} = 1/n$ and $r_n = a_n/\sqrt{n}$. Then $r_n = o(a_n)$, but

$$\frac{r_n}{V_{B,n}} = \sqrt{n} a_n,$$

which is not $o(a_n)$. Since first-order implied-volatility perturbations divide price perturbations by vega, a lower bound on vega is necessary for a uniform price remainder to remain lower order after inversion. $\square$

Remark C.15 (Black-conversion scope). The correct Black statement is local: on a regular window satisfying Assumption E.4, Black and normal visibility agree for the surface jets whose linear conversion coefficient does not vanish. The theorem does not extend across maturity-strike regions where vega degenerates, the inverse map loses uniform smoothness, or the Taylor remainder is not controlled at the selected scale.

C.5. Summary of scope.

Remark C.16 (Net effect of the scope analysis). The core structure remains risk-neutral compression of priced structural loading. The surrounding assumptions state which additional regularity, basis, tail, or model-class conditions are needed to express that structure as a cumulant skew law, a local Black-implied-volatility law, a finite-dimensional premium support law, or a nonlinear variance-linked Hermite law. When one of those conditions is unavailable, the compression theorem itself remains in force, but the corresponding endpoint conversion should be replaced by the more general projection or price-kernel statement proved earlier.

APPENDIX D. MARKET-SCALE SHORT-MATURITY SKEW ASYMPTOTICS

The structural option statement is Theorem 6.3; the explicit public skew law is Theorem 6.7. This section verifies the non-Gaussian cumulant-class hypotheses in the Gaussian-Volterra/Bachelier market-index model from Section 6 and computes the aggregated coefficients in closed form. At factor level, the Gaussian block reduces option alignment to the centered structural loading of the perturbation return Z_τ and the centered quadratic-fluctuation remainder generated by $X_\tau^2 - v_\tau$.

The internal ladder is proof infrastructure. The centered option-facing loading is first reduced to a canonical Perron-generated proxy claim, then to a benchmark-profile claim built from the leading Perron kernels, then to the benchmark-generated centered ϖ -profile generator, and only at the last step to the aggregate factor benchmark Λ_M^* . The consequence for this article is the Companion A source-family verification. Its role is to prove the scalar seam, the common-channel reduction, and the option-factor theorem, not to introduce additional endpoints. The centered mixed visible/ ϖ remainder and the visible-screening defect are fixed structural obstructions. The off-Perron and quadratic terms remain the maturity-scale defects.

Definition D.1 (Gaussian specialization of the option first-variation pricing object). Under Assumption 6.8, define

$$\Xi_{M,\tau} := \frac{1}{2(\sigma_0^M(0))^3\tau^2} X_\tau^2 Z_\tau.$$

This is the Gaussian-Volterra specialization of the abstract option first-variation pricing object from Assumption 6.1; its priced structural loading is the corresponding $L_{M,\tau}^\mathbb{Q} = C_{\mathcal{H}_\varpi}^\mathbb{Q}(\Xi_{M,\tau})$ from Definition 6.2.

Lemma D.2 (Market leverage covariance asymptotics). *Under Assumption 6.8, if $\mathcal{E}_M^{\text{ret}} = \emptyset$, then*

$$\text{Cov}_\mathbb{Q}(Y_M(t), B_t^\mathbb{Q}) = 0 \quad \text{for every } t \in [0, T].$$

If $\mathcal{E}_M^{\text{ret}} \neq \emptyset$, then, as $t \downarrow 0$,

$$\text{Cov}_\mathbb{Q}(Y_M(t), B_t^\mathbb{Q}) = \sum_{(i,j) \in \mathcal{E}_M^{\text{ret}}} \ell_{ij}^M t^{H_{ij}+1/2} + o(t^{\beta_M+1/2}).$$

Proof. By the covariance formula for correlated Brownian integrals,

$$\text{Cov}_\mathbb{Q}(Y_M(t), B_t^\mathbb{Q}) = \sum_{j=1}^N \rho_j \int_0^t G_j(t, s) ds.$$

If $\mathcal{E}_M^{\text{ret}} = \emptyset$, every coefficient $b_i \dot{\beta}_{ij} \rho_j \mathbf{1}_{\{(i,j) \in \mathcal{E}_M\}}$ vanishes, so this sum is identically zero. Assume now that $\mathcal{E}_M^{\text{ret}} \neq \emptyset$. Substituting the definition of G_j and using continuity of a_{ij} along the diagonal gives

$$\int_0^t a_{ij}(t, s) (t-s)^{H_{ij}-1/2} ds = \frac{a_{ij}(0, 0)}{H_{ij} + 1/2} t^{H_{ij}+1/2} + o(t^{H_{ij}+1/2}),$$

uniformly over the finitely many active channels. Summing over $(i, j) \in \mathcal{E}_M$ yields the stated expansion, and channels with $b_i \dot{\beta}_{ij} \rho_j = 0$ drop out by definition of $\mathcal{E}_M^{\text{ret}}$. $\square$

Proposition D.3 (First variation of the third cumulant of the perturbed market return). *Under Assumption 6.8, if $\mathcal{E}_M^{\text{ret}} = \emptyset$, then*

$$\partial_\varepsilon \kappa_3^\mathbb{Q}(R_\tau^\varepsilon)|_{\varepsilon=0} = 0 \quad \text{for every } \tau \in [0, T].$$

If $\mathcal{E}_M^{\text{ret}} \neq \emptyset$, then, as $\tau \downarrow 0$,

$$\partial_\varepsilon \kappa_3^\mathbb{Q}(R_\tau^\varepsilon)|_{\varepsilon=0} = \sum_{(i,j) \in \mathcal{E}_M^{\text{ret}}} \kappa_{ij}^M \tau^{H_{ij}+3/2} + o(\tau^{\beta_M+3/2}).$$

Proof. Because $R_\tau^\varepsilon = X_\tau + \varepsilon Z_\tau$ and X_τ is centered Gaussian,

$$\partial_\varepsilon \kappa_3^\mathbb{Q}(R_\tau^\varepsilon)|_{\varepsilon=0} = 3 \mathbb{E}_\mathbb{Q}[X_\tau^2 Z_\tau].$$

Apply Itô's product formula to $X_t^2 Z_t$. Since

$$dX_t = \sigma_0^M(t) dB_t^\mathbb{Q}, \quad dZ_t = Y_M(t) dB_t^\mathbb{Q},$$

taking expectations yields

$$\mathbb{E}_\mathbb{Q}[X_\tau^2 Z_\tau] = 2 \int_0^\tau \sigma_0^M(u) \mathbb{E}_\mathbb{Q}[X_u Y_M(u)] du.$$

By correlated Itô isometry,

$$\mathbb{E}_{\mathbb{Q}}[X_u Y_M(u)] = \sum_{j=1}^N \rho_j \int_0^u \sigma_0^M(r) G_j(u, r) dr.$$

If $\mathcal{E}_M^{\text{ret}} = \emptyset$, the last sum is identically zero, so the first claim follows. Assume now that $\mathcal{E}_M^{\text{ret}} \neq \emptyset$. Using the continuity of σ_0^M at 0, the continuity of a_{ij} along the diagonal, and the definition of G_j , one obtains

$$\mathbb{E}_{\mathbb{Q}}[X_u Y_M(u)] = \sigma_0^M(0) \sum_{(i,j) \in \mathcal{E}_M^{\text{ret}}} \frac{b_i \dot{\beta}_{ij} a_{ij}(0,0) \rho_j}{H_{ij} + 1/2} u^{H_{ij}+1/2} + o(u^{\beta_M+1/2}).$$

Integrating in u , multiplying by the outer factor $2\sigma_0^M(0)$, and then multiplying by 3 gives the stated coefficient κ_{ij}^M . $\square$

Remark D.4 (Where the Gaussian structure enters). The first-variation identity

$$\partial_\varepsilon \kappa_3^{\mathbb{Q}}(R_\tau^\varepsilon)|_{\varepsilon=0} = 3 \mathbb{E}_{\mathbb{Q}}[X_\tau^2 Z_\tau]$$

is a specialization to the present Gaussian-Volterra $\mathbb{Q}$ -model. Gaussianity is absent from the abstract compression statement of Theorem 6.3; it enters only in the explicit closed-form cumulant law derived here.

Lemma D.5 (Market variance expansion). *Under Assumption 6.8,*

$$\text{Var}_{\mathbb{Q}}(R_\tau^0) = (\sigma_0^M(0))^2 \tau + o(\tau), \quad \tau \downarrow 0.$$

Proof. By Itô isometry,

$$\text{Var}_{\mathbb{Q}}(R_\tau^0) = \int_0^\tau (\sigma_0^M(u))^2 du.$$

Continuity of σ_0^M at 0 gives the claim. $\square$

Lemma D.6 (Projection of second chaos onto one Gaussian coordinate). *Let X be centered Gaussian with variance $v > 0$, and let $Z \in L^2(\mathbb{Q})$ be a centered second-Wiener-chaos random variable for the same Gaussian system. Then*

$$\mathbb{E}_{\mathbb{Q}}[Z | X] = \frac{\mathbb{E}_{\mathbb{Q}}[X^2 Z]}{2v^2} (X^2 - v).$$

Equivalently, if $\mathbb{E}_{\mathbb{Q}}[Z | X] = a(X^2 - v)$, then

$$a = \frac{\mathbb{E}_{\mathbb{Q}}[X^2 Z]}{2v^2} = \frac{\partial_\varepsilon \kappa_3^{\mathbb{Q}}(X + \varepsilon Z)|_{\varepsilon=0}}{6v^2}.$$

Consequently, for every $k \in \mathbb{R}$,

$$\mathbb{E}_{\mathbb{Q}}[Z \mathbf{1}_{\{X > k\}}] = \frac{\partial_\varepsilon \kappa_3^{\mathbb{Q}}(X + \varepsilon Z)|_{\varepsilon=0}}{6v} \frac{k}{\sqrt{v}} \phi\left(\frac{k}{\sqrt{v}}\right),$$

where ϕ is the standard normal density.

Proof. The closed subspace $L^2(\mathbb{Q}, \sigma(X))$ has orthogonal basis $\{\text{He}_m(X/\sqrt{v}) : m \geq 0\}$. Since Z lies in the second Wiener chaos, it is orthogonal to every Hermite direction except possibly $m = 2$. Therefore

$$\mathbb{E}_{\mathbb{Q}}[Z | X] = a \text{He}_2(X/\sqrt{v}) = a \left(\frac{X^2}{v} - 1 \right)$$

for some $a \in \mathbb{R}$. Equivalently,

$$\mathbb{E}_{\mathbb{Q}}[Z | X] = \frac{a}{v} (X^2 - v).$$

For a centered normal variable X with variance v ,

$$\mathbb{E}_{\mathbb{Q}}[(X^2 - v)^2] = 2v^2.$$

Taking inner products with $X^2 - v$ gives

$$\frac{a}{v} = \frac{\mathbb{E}_{\mathbb{Q}}[Z(X^2 - v)]}{2v^2} = \frac{\mathbb{E}_{\mathbb{Q}}[X^2 Z]}{2v^2},$$

which is the displayed conditional projection formula after renaming a/v as the coefficient of $X^2 - v$. Since X is centered Gaussian and $\mathbb{E}_{\mathbb{Q}}[Z] = 0$,

$$\partial_{\varepsilon} \kappa_3^{\mathbb{Q}}(X + \varepsilon Z)|_{\varepsilon=0} = 3\mathbb{E}_{\mathbb{Q}}[X^2 Z],$$

which proves the coefficient identity. Finally, if $Y = X/\sqrt{v}$, then the standard normal integration-by-parts identity gives

$$\mathbb{E}_{\mathbb{Q}}[(X^2 - v)\mathbf{1}_{\{X > k\}}] = v\mathbb{E}_{\mathbb{Q}}[(Y^2 - 1)\mathbf{1}_{\{Y > k/\sqrt{v}\}}] = v\frac{k}{\sqrt{v}}\phi\left(\frac{k}{\sqrt{v}}\right).$$

Combining the displays yields the claim. $\square$

Proposition D.7 (Gaussian-Volterra leverage perturbations satisfy the projection condition). *Under Assumption 6.8, for each sufficiently small $\tau > 0$, the pair (X_{τ}, Z_{τ}) satisfies the projection condition of Lemma D.6. If $\mathcal{E}_M^{\text{ret}} = \emptyset$, then this projection is zero.*

Proof. Decompose each $W^{\mathbb{Q},j}$ into its component correlated with $B^{\mathbb{Q}}$ and an orthogonal Gaussian component. The Brownian-correlated part of $Y_M(t)$ is a deterministic Volterra integral against $B^{\mathbb{Q}}$; integrating it again against $B^{\mathbb{Q}}$ places its contribution to Z_{τ} in the second Wiener chaos of the return Brownian motion. Conditional expectation onto $\sigma(X_{\tau})$, where X_{τ} is a first-chaos variable, keeps only the second Hermite direction generated by $X_{\tau}/\sqrt{\text{Var}_{\mathbb{Q}}(X_{\tau})}$. The orthogonal Gaussian components are independent of X_{τ} and have zero mean, so their $\sigma(X_{\tau})$ -projection vanishes. If $\mathcal{E}_M^{\text{ret}} = \emptyset$, all Brownian-correlated coefficients vanish, hence the whole projection is zero. $\square$

Theorem D.8 (Self-contained Bachelier skew-cumulant transfer). *Let X_{τ} be centered Gaussian with variance $v(\tau) > 0$, let $Z_{\tau} \in L^2(\mathbb{Q})$, and set $R_{\tau}^{\varepsilon} := X_{\tau} + \varepsilon Z_{\tau}$. Define*

$$C(\tau, k; \varepsilon) := \mathbb{E}_{\mathbb{Q}}[(X_{\tau} + \varepsilon Z_{\tau} - k)_+],$$

and let $\sigma_N(\tau, k; \varepsilon)$ be the Bachelier implied normal volatility defined by

$$C(\tau, k; \varepsilon) = \sigma_N(\tau, k; \varepsilon)\sqrt{\tau}\phi\left(\frac{k}{\sigma_N(\tau, k; \varepsilon)\sqrt{\tau}}\right) - k\Phi\left(-\frac{k}{\sigma_N(\tau, k; \varepsilon)\sqrt{\tau}}\right),$$

where Φ is the standard normal distribution function. Assume the projection condition of Lemma D.6 holds for (X_{τ}, Z_{τ}) . Then, for fixed small τ ,

$$\partial_{\varepsilon} C(\tau, k; \varepsilon)|_{\varepsilon=0} = \frac{\partial_{\varepsilon} \kappa_3^{\mathbb{Q}}(R_{\tau}^{\varepsilon})|_{\varepsilon=0}}{6v(\tau)} \frac{k}{\sqrt{v(\tau)}} \phi\left(\frac{k}{\sqrt{v(\tau)}}\right).$$

Consequently,

$$\partial_{\varepsilon} \sigma_N(\tau, k; \varepsilon)|_{\varepsilon=0} = \frac{\partial_{\varepsilon} \kappa_3^{\mathbb{Q}}(R_{\tau}^{\varepsilon})|_{\varepsilon=0}}{6v(\tau)^{3/2}\sqrt{\tau}} k,$$

and the ATM normal skew transfer is exact:

$$\partial_{\varepsilon} \partial_k \sigma_N(\tau, k; \varepsilon)|_{k=0, \varepsilon=0} = \frac{\partial_{\varepsilon} \kappa_3^{\mathbb{Q}}(R_{\tau}^{\varepsilon})|_{\varepsilon=0}}{6v(\tau)^{3/2}\sqrt{\tau}}.$$

Proof. Because the law of X_{τ} has no atom at k and $Z_{\tau} \in L^2(\mathbb{Q})$, differentiating the call payoff at $\varepsilon = 0$ gives

$$\partial_{\varepsilon} C(\tau, k; \varepsilon)|_{\varepsilon=0} = \mathbb{E}_{\mathbb{Q}}[Z_{\tau} \mathbf{1}_{\{X_{\tau} > k\}}].$$

Lemma D.6 gives the first displayed formula.

At $\varepsilon = 0$, the Bachelier implied normal volatility is

$$\sigma_N(\tau, k; 0) = \sqrt{v(\tau)/\tau},$$

because X_τ is Gaussian. The Bachelier vega at this point is

$$\partial_\sigma C_N(\tau, k; \sigma)|_{\sigma=\sqrt{v(\tau)/\tau}} = \sqrt{\tau} \phi\left(\frac{k}{\sqrt{v(\tau)}}\right).$$

Implicit differentiation of the defining equation for implied normal volatility therefore gives

$$\partial_\varepsilon \sigma_N(\tau, k; \varepsilon)|_{\varepsilon=0} = \frac{\partial_\varepsilon C(\tau, k; \varepsilon)|_{\varepsilon=0}}{\sqrt{\tau} \phi(k/\sqrt{v(\tau)})} = \frac{\partial_\varepsilon \kappa_3^\mathbb{Q}(R_\tau^\varepsilon)|_{\varepsilon=0}}{6v(\tau)^{3/2}\sqrt{\tau}} k.$$

Differentiating in k at zero proves the exact ATM skew transfer. $\square$

Proposition D.9 (The non-Gaussian Volterra leverage perturbation satisfies the projection condition). *Under Assumption A.6, for each sufficiently small $\tau > 0$, the perturbation return Z_τ satisfies the projection condition of Lemma D.6. The compound-Poisson and orthogonal Brownian components have zero contribution to $\mathbb{E}_\mathbb{Q}[Z_\tau | X_\tau]$; the Brownian-correlated Volterra component contributes only through the second Hermite direction $X_\tau^2 - v(\tau)$.*

Proof. The Brownian-correlated part of Y_t is a deterministic Volterra integral against B . Hence its stochastic integral against B is a second-Wiener-chaos random variable. Conditional expectation onto $\sigma(X_\tau)$, where X_τ is a first-chaos variable, keeps only the component in the second Hermite polynomial of $X_\tau/\sqrt{v(\tau)}$. The orthogonal Brownian components are independent of X_τ and have zero mean. Conditional on the path of a compensated jump-source Volterra integrand, the stochastic integral against B is jointly Gaussian with X_τ , and its conditional covariance is linear in that jump-source integrand. Taking expectation over the compensated jump source gives zero because the integrand has mean zero. Thus the jump and orthogonal Brownian parts have zero $\sigma(X_\tau)$ -projection, and the stated projection form follows. $\square$

Corollary D.10 (Self-contained Bachelier expansion for the primitive non-Gaussian Volterra class). *Under Assumption A.6, in the nonempty support case of Theorem A.8,*

$$\partial_\varepsilon \partial_k \sigma_N(\tau, k; \varepsilon)|_{k=0, \varepsilon=0} = \sum_{\beta \in \mathcal{B}_J} A_J(\beta) \tau^{\beta-1/2} + O(\tau^{\beta_0-1/2+\eta}),$$

with the same support and cancellation qualifications as in Theorem A.8.

Proof. Theorem A.8 gives the cumulant expansion with the displayed support, and Proposition D.9 is the admissibility condition that lets that cumulant coefficient enter the Bachelier skew transfer. Theorem D.8 then divides by the leading variance scale. Since

$$v(\tau) = \sigma_0(0)^2 \tau + O(\tau^{1+\eta}).$$

it follows that the division shifts each cumulant exponent by $-1/2$ and leaves the stated remainder order. $\square$

Corollary D.11 (Gaussian-Volterra/Bachelier verification of the cumulant-class transfer). *Under Assumption 6.8, if $\mathcal{E}_M^{\text{ret}} = \emptyset$, then*

$$\partial_\varepsilon \kappa_3^\mathbb{Q}(R_\tau^\varepsilon)|_{\varepsilon=0} = \partial_\varepsilon \mathfrak{S}_{M,N}(\tau; \varepsilon)|_{\varepsilon=0} = 0 \quad \text{for every sufficiently small } \tau > 0.$$

If instead $\mathcal{E}_M^{\text{ret}} \neq \emptyset$, so that $\beta_M < \infty$, then as $\tau \downarrow 0$,

$$\partial_\varepsilon \mathfrak{S}_{M,N}(\tau; \varepsilon)|_{\varepsilon=0} = \frac{1}{6(\sigma_0^M(0))^3 \tau^2} \partial_\varepsilon \kappa_3^\mathbb{Q}(R_\tau^\varepsilon)|_{\varepsilon=0} + o(\tau^{\beta_M-1/2}).$$

Equivalently,

$$\partial_\varepsilon \mathfrak{S}_{M,N}(\tau; \varepsilon)|_{\varepsilon=0} = \mathbb{E}_\mathbb{Q}[\Xi_{M,\tau}] + o(\tau^{\beta_M-1/2}),$$

where $\Xi_{M,\tau}$ is the Gaussian-Volterra first-variation pricing object from Definition D.1. In particular, the Gaussian-Volterra/Bachelier option block satisfies Assumption 6.1 with trivial public option sigma-field, identity summary functional, and, in the nontrivial case $\mathcal{E}_M^{\text{ret}} \neq \emptyset$, scale

$$a_{M,\tau} := \tau^{\beta_M-1/2}.$$

Moreover, with

$$G_M(\tau) := \frac{\text{Var}_{\mathbb{Q}}(R_\tau^0)^{3/2}}{(\sigma_0^M(0))^{3\tau^{3/2}}},$$

one has $G_M(\tau) \rightarrow 1$. Consequently, after the coefficient grouping used in Corollary D.61, the Gaussian-Volterra block verifies Assumption 6.5 whenever the grouped cumulant and quote-transfer remainders are lower order at the resulting selected exponent β_M^{comp} .

Proof. If $\mathcal{E}_M^{\text{ret}} = \emptyset$, then

$$\mathbb{E}_{\mathbb{Q}}[X_u Y_M(u)] = 0 \quad \text{for every } u \in [0, T],$$

by the correlated Itô-isometry formula used in the proof of Proposition D.3. Hence

$$\partial_\varepsilon \kappa_3^{\mathbb{Q}}(R_\tau^\varepsilon)|_{\varepsilon=0} = 0$$

for all small τ . Moreover, the nonzero channels in Y_M are orthogonal to $B^{\mathbb{Q}}$; joint Gaussianity gives independence from the return Brownian channel. Conditional on the volatility perturbation path, R_τ^ε is therefore centered Gaussian, hence symmetric. Equivalently, Proposition D.7 gives a zero second-Hermite projection, and Theorem D.8 gives vanishing of the ATM Bachelier skew derivative. Assume therefore that $\mathcal{E}_M^{\text{ret}} \neq \emptyset$. Apply Theorem D.8 to the additive Gaussian-Volterra market-index return R_τ^ε , and combine it with Lemma D.5. The first display is the resulting transfer law. Substituting the identity from Proposition D.3 gives

$$\frac{1}{6(\sigma_0^M(0))^{3\tau^2}} \partial_\varepsilon \kappa_3^{\mathbb{Q}}(R_\tau^\varepsilon)|_{\varepsilon=0} = \frac{1}{2(\sigma_0^M(0))^{3\tau^2}} \mathbb{E}_{\mathbb{Q}}[X_\tau^2 Z_\tau] = \mathbb{E}_{\mathbb{Q}}[\Xi_{M,\tau}],$$

which is the second display. $\square$

Proposition D.12 (First variation of the public skew summary as a compression of priced structural loading). *Under the assumptions of Corollary D.11, define*

$$\Gamma_{M,\tau}^{\mathbb{Q}} := C_{\mathcal{H}_{M,\tau}^{\text{opt}}}^{\mathbb{Q}}(L_{M,\tau}^{\mathbb{Q}}).$$

If $\alpha_{M,\tau}^{\mathbb{Q}} > 0$, then

$$\Gamma_{M,\tau}^{\mathbb{Q}} = \alpha_{M,\tau}^{\mathbb{Q}} C_{\mathcal{H}_{M,\tau}^{\text{opt}}}^{\mathbb{Q}}(\Lambda_{M,\tau}).$$

If $\mathcal{E}_M^{\text{ret}} = \emptyset$, then

$$\partial_\varepsilon \mathfrak{S}_{M,N}(\tau; \varepsilon)|_{\varepsilon=0} = \Gamma_{M,\tau}^{\mathbb{Q}} = 0 \quad \text{for every sufficiently small } \tau > 0.$$

If $\mathcal{E}_M^{\text{ret}} \neq \emptyset$, then, as $\tau \downarrow 0$,

$$\partial_\varepsilon \mathfrak{S}_{M,N}(\tau; \varepsilon)|_{\varepsilon=0} = \Gamma_{M,\tau}^{\mathbb{Q}} + r_{M,\tau}, \quad |r_{M,\tau}| = o(a_{M,\tau}),$$

with $a_{M,\tau} = \tau^{\beta_M - 1/2}$. Thus the public skew summary is the public compression of the option-facing priced structural loading up to a higher-order remainder at the Gaussian market scale.

Proof. Under Corollary D.11, the Gaussian-Volterra/Bachelier option family satisfies Assumption 6.1 with trivial public option sigma-field, identity summary functional, and option first-variation pricing object $\Xi_{M,\tau}$ from Definition D.1. The compression identity is therefore the abstract statement of Theorem 6.3 specialized to this model. Since the public option sigma-field is trivial, $\Gamma_{M,\tau}^{\mathbb{Q}}$ is the constant $\mathbb{E}_{\mathbb{Q}}[\Xi_{M,\tau}]$. The empty-channel zero follows from Corollary D.11; the nonempty remainder scale is the scale identified there. $\square$

Corollary D.13 (Option visibility is norm visibility at the quote scale). *Under the assumptions of Proposition D.12,*

$$\|C_{\mathcal{H}_{M,\tau}^{\text{opt}}}^{\mathbb{Q}}(L_{M,\tau}^{\mathbb{Q}})\|_{L^2(\mathbb{Q})} = |\Gamma_{M,\tau}^{\mathbb{Q}}|.$$

Because $\mathcal{H}_{M,\tau}^{\text{opt}}$ is trivial, the option-facing structural contribution induced by the synchronized perturbation is $\mathbb{Q}$ -visible at maturity τ exactly when $\Gamma_{M,\tau}^{\mathbb{Q}} \neq 0$, and it is screened at that quote scale exactly when $\Gamma_{M,\tau}^{\mathbb{Q}} = 0$ while $L_{M,\tau}^{\mathbb{Q}} \neq 0$.

Proof. The compressed loading $C_{\mathcal{H}_{M,\tau}^{\text{opt}}}^{\mathbb{Q}}(L_{M,\tau}^{\mathbb{Q}})$ is the constant random variable $\Gamma_{M,\tau}^{\mathbb{Q}}$. The $L^2(\mathbb{Q})$ -norm of a constant equals its absolute value, and Proposition 5.3 then gives the stated visibility criterion. $\square$

Theorem D.14 (Aggregate-factor approximation for the option-facing centered loading). *Assume Assumption 6.1, and suppose the aggregate factor benchmark is nontrivial, so*

$$\Lambda_M^* \in \mathcal{M}_{\mathbb{Q}}^{\circ}, \quad \|\Lambda_M^*\|_{L^2(\mathbb{Q})} = 1.$$

Fix a sufficiently small maturity $\tau > 0$, and suppose that the option first-variation pricing object admits a decomposition

$$\Xi_{M,\tau} = X_{M,\tau}^{\text{vis}} + a_{M,\tau}^* \Lambda_M^* + E_{M,\tau},$$

where

$$X_{M,\tau}^{\text{vis}} \in L^2(\mathbb{Q}, \mathcal{G}_T), \quad a_{M,\tau}^* \in \mathbb{R}, \quad E_{M,\tau} \in L^2(\mathbb{Q}, \mathcal{F}_T).$$

Define the centered projected remainder

$$\tilde{E}_{M,\tau}^{\mathbb{Q}} := C_{\mathcal{H}_{\omega}}^{\mathbb{Q}}(E_{M,\tau}) - C_{\mathcal{G}_T}^{\mathbb{Q}}(C_{\mathcal{H}_{\omega}}^{\mathbb{Q}}(E_{M,\tau})),$$

and set

$$\delta_{M,\tau}^* := \|\tilde{E}_{M,\tau}^{\mathbb{Q}}\|_{L^2(\mathbb{Q})}.$$

Then the centered option-facing priced structural loading satisfies

$$\tilde{L}_{M,\tau}^{\mathbb{Q}} := L_{M,\tau}^{\mathbb{Q}} - C_{\mathcal{G}_T}^{\mathbb{Q}}(L_{M,\tau}^{\mathbb{Q}}) = a_{M,\tau}^* \Lambda_M^* + \tilde{E}_{M,\tau}^{\mathbb{Q}}.$$

Consequently, in the decomposition of Proposition 4.6,

$$\tilde{L}_{M,\tau}^{\mathbb{Q}} = \alpha_{M,\tau}^{\mathbb{Q},*} \Lambda_M^* + R_{M,\tau}^{\mathbb{Q},*},$$

one has

$$|\alpha_{M,\tau}^{\mathbb{Q},*} - a_{M,\tau}^*| \leq \delta_{M,\tau}^*, \quad \|R_{M,\tau}^{\mathbb{Q},*}\|_{L^2(\mathbb{Q})} \leq \delta_{M,\tau}^*.$$

If, in addition,

$$a_{M,\tau}^* \neq 0 \quad \text{for all sufficiently small } \tau, \quad \delta_{M,\tau}^* = o(|a_{M,\tau}^*|),$$

then the centered option-facing loading is asymptotically aggregate-factor aligned:

$$\frac{\|R_{M,\tau}^{\mathbb{Q},*}\|_{L^2(\mathbb{Q})}}{|\alpha_{M,\tau}^{\mathbb{Q},*}|} \rightarrow 0 \quad \text{as } \tau \downarrow 0,$$

and

$$\left\| \frac{\tilde{L}_{M,\tau}^{\mathbb{Q}}}{\|\tilde{L}_{M,\tau}^{\mathbb{Q}}\|_{L^2(\mathbb{Q})}} - \frac{a_{M,\tau}^*}{|a_{M,\tau}^*|} \Lambda_M^* \right\|_{L^2(\mathbb{Q})} \rightarrow 0.$$

Proof. By linearity of conditional expectation and because $X_{M,\tau}^{\text{vis}} \in L^2(\mathbb{Q}, \mathcal{G}_T) \subseteq L^2(\mathbb{Q}, \mathcal{H}_{\omega})$,

$$L_{M,\tau}^{\mathbb{Q}} = X_{M,\tau}^{\text{vis}} + a_{M,\tau}^* \Lambda_M^* + C_{\mathcal{H}_{\omega}}^{\mathbb{Q}}(E_{M,\tau}).$$

Apply $C_{\mathcal{G}_T}^{\mathbb{Q}}$ and use $C_{\mathcal{G}_T}^{\mathbb{Q}}(\Lambda_M^*) = 0$ from Corollary 4.7:

$$C_{\mathcal{G}_T}^{\mathbb{Q}}(L_{M,\tau}^{\mathbb{Q}}) = X_{M,\tau}^{\text{vis}} + C_{\mathcal{G}_T}^{\mathbb{Q}}(C_{\mathcal{H}_{\omega}}^{\mathbb{Q}}(E_{M,\tau})).$$

Subtracting gives the first displayed identity for $\tilde{L}_{M,\tau}^{\mathbb{Q}}$. Now apply Proposition 4.6 to $\tilde{L}_{M,\tau}^{\mathbb{Q}}$. Since

$$\alpha_{M,\tau}^{\mathbb{Q},*} = \langle \tilde{L}_{M,\tau}^{\mathbb{Q}}, \Lambda_M^* \rangle_{L^2(\mathbb{Q})} = a_{M,\tau}^* + \langle \tilde{E}_{M,\tau}^{\mathbb{Q}}, \Lambda_M^* \rangle_{L^2(\mathbb{Q})},$$

Cauchy–Schwarz yields

$$|\alpha_{M,\tau}^{\mathbb{Q},*} - a_{M,\tau}^*| \leq \delta_{M,\tau}^*.$$

The orthogonal residual is

$$R_{M,\tau}^{\mathbb{Q},*} = \tilde{E}_{M,\tau}^{\mathbb{Q}} - \langle \tilde{E}_{M,\tau}^{\mathbb{Q}}, \Lambda_M^* \rangle_{L^2(\mathbb{Q})} \Lambda_M^*,$$

so

$$\|R_{M,\tau}^{\mathbb{Q},*}\|_{L^2(\mathbb{Q})} \leq \delta_{M,\tau}^*.$$

If $\delta_{M,\tau}^* = o(|a_{M,\tau}^*|)$, then the previous coefficient bound implies

$$|\alpha_{M,\tau}^{\mathbb{Q},*}| \geq |a_{M,\tau}^*| - \delta_{M,\tau}^* \sim |a_{M,\tau}^*|,$$

hence

$$\frac{\|R_{M,\tau}^{\mathbb{Q},*}\|_{L^2(\mathbb{Q})}}{|\alpha_{M,\tau}^{\mathbb{Q},*}|} \rightarrow 0.$$

In particular, $\alpha_{M,\tau}^{\mathbb{Q},*}$ has the same sign as $a_{M,\tau}^*$ for all sufficiently small τ . Finally,

$$\frac{\tilde{L}_{M,\tau}^{\mathbb{Q}}}{\|\tilde{L}_{M,\tau}^{\mathbb{Q}}\|_{L^2(\mathbb{Q})}} = \frac{\alpha_{M,\tau}^{\mathbb{Q},*}}{\sqrt{(\alpha_{M,\tau}^{\mathbb{Q},*})^2 + \|R_{M,\tau}^{\mathbb{Q},*}\|_{L^2(\mathbb{Q})}^2}} \Lambda_M^* + \frac{R_{M,\tau}^{\mathbb{Q},*}}{\|\tilde{L}_{M,\tau}^{\mathbb{Q}}\|_{L^2(\mathbb{Q})}},$$

where the first scalar coefficient converges to $a_{M,\tau}^*/|a_{M,\tau}^*|$ and the second term converges to 0 in $L^2(\mathbb{Q})$. This proves the normalized-direction claim. $\square$

Corollary D.15 (Exact option-factor alignment criterion). *Under the hypotheses of Theorem D.14, if*

$$C_{\mathcal{H}_\omega}^{\mathbb{Q}}(E_{M,\tau}) = C_{\mathcal{G}_T}^{\mathbb{Q}}(C_{\mathcal{H}_\omega}^{\mathbb{Q}}(E_{M,\tau})),$$

equivalently $\delta_{M,\tau}^ = 0$, then*

$$\tilde{L}_{M,\tau}^{\mathbb{Q}} = a_{M,\tau}^* \Lambda_M^*, \quad R_{M,\tau}^{\mathbb{Q},*} = 0.$$

In particular, the option-facing centered loading is exactly aggregate-factor aligned whenever the non-factor remainder contributes no centered $\mathcal{H}_\omega$ -content beyond the visible sigma-field.

Proof. If $\delta_{M,\tau}^* = 0$, then $\tilde{E}_{M,\tau}^{\mathbb{Q}} = 0$ in Theorem D.14. The factor decomposition in that theorem therefore reduces to the single span component $\tilde{L}_{M,\tau}^{\mathbb{Q}} = a_{M,\tau}^* \Lambda_M^*$, and the orthogonal residual $R_{M,\tau}^{\mathbb{Q},*}$ is zero by uniqueness of the Hilbert-space decomposition. $\square$

Remark D.16 (Meaning for public option summaries). Theorem D.14 is stated only in the nontrivial factor regime $\Lambda_M^* \neq 0$. When the centered aggregate benchmark collapses, there is no nonzero aggregate factor direction for the centered option-facing loading to align with, so the theorem's factor statement becomes vacuous.

Theorem D.14 does not claim direct public visibility of the aggregate factor. Corollary 4.7 still gives

$$C_{\mathcal{H}_{M,\tau}}^{\mathbb{Q}}(\Lambda_M^*) = 0 \quad \text{for } \mathcal{H}_{M,\tau}^{\text{opt}} \subseteq \mathcal{G}_T.$$

Its content is different. When the option first-variation object is generated by a visible component plus the aggregate factor benchmark up to a lower-order centered remainder, the centered option-facing loading itself points asymptotically in the aggregate factor direction. The quoted option summary still acts on the visible representative of that loading, but the nonlinear and hedging consequences of the option-facing claim inherit the same factor dominance.

Proposition D.17 (Gaussian reduction of option-factor alignment to the perturbation return). *Assume Assumption 6.8, and suppose the aggregate factor benchmark is nontrivial, so*

$$\Lambda_M^* \in \mathcal{M}_{\mathbb{Q}}^{\circ}, \quad \|\Lambda_M^*\|_{L^2(\mathbb{Q})} = 1.$$

Set

$$q_\tau := \frac{1}{2(\sigma_0^M(0))^3 \tau^2}, \quad v_\tau := \text{Var}_{\mathbb{Q}}(R_\tau^0).$$

Let

$$L_{Z_\tau}^{\mathbb{Q}} := C_{\mathcal{H}_\omega}^{\mathbb{Q}}(Z_\tau), \quad V_{Z_\tau}^{\mathbb{Q}} := C_{\mathcal{G}_T}^{\mathbb{Q}}(L_{Z_\tau}^{\mathbb{Q}}),$$

and write the canonical factor decomposition of the perturbation return loading as

$$L_{Z_\tau}^{\mathbb{Q}} = V_{Z_\tau}^{\mathbb{Q}} + b_{M,\tau}^{\mathbb{Q},*} \Lambda_M^* + R_{Z_\tau}^{\mathbb{Q},*},$$

where $R_{Z_\tau}^{\mathbb{Q},} \in \mathcal{M}_{\mathbb{Q}}^{\circ}$ and*

$$R_{Z_\tau}^{\mathbb{Q},*} \perp_{L^2(\mathbb{Q})} \Lambda_M^*.$$

Define the quadratic-fluctuation remainder

$$J_{M,\tau} := q_\tau(X_\tau^2 - v_\tau)Z_\tau,$$

and its centered structural projection

$$\tilde{J}_{M,\tau}^\mathbb{Q} := C_{\mathcal{H}_\omega}^\mathbb{Q}(J_{M,\tau}) - C_{\mathcal{G}_T}^\mathbb{Q}(C_{\mathcal{H}_\omega}^\mathbb{Q}(J_{M,\tau})).$$

Then the centered option-facing priced structural loading satisfies the exact identity

$$\tilde{L}_{M,\tau}^\mathbb{Q} = q_\tau v_\tau b_{M,\tau}^{\mathbb{Q},\star} \Lambda_M^\star + q_\tau v_\tau R_{Z_\tau}^{\mathbb{Q},\star} + \tilde{J}_{M,\tau}^\mathbb{Q}.$$

Consequently,

$$|\alpha_{M,\tau}^{\mathbb{Q},\star} - q_\tau v_\tau b_{M,\tau}^{\mathbb{Q},\star}| \leq \|\tilde{J}_{M,\tau}^\mathbb{Q}\|_{L^2(\mathbb{Q})},$$

and

$$\|R_{M,\tau}^{\mathbb{Q},\star}\|_{L^2(\mathbb{Q})} \leq q_\tau v_\tau \|R_{Z_\tau}^{\mathbb{Q},\star}\|_{L^2(\mathbb{Q})} + \|\tilde{J}_{M,\tau}^\mathbb{Q}\|_{L^2(\mathbb{Q})}.$$

If, in addition,

$$b_{M,\tau}^{\mathbb{Q},\star} \neq 0 \quad \text{for all sufficiently small } \tau,$$

and

$$q_\tau v_\tau \|R_{Z_\tau}^{\mathbb{Q},\star}\|_{L^2(\mathbb{Q})} + \|\tilde{J}_{M,\tau}^\mathbb{Q}\|_{L^2(\mathbb{Q})} = o(q_\tau v_\tau |b_{M,\tau}^{\mathbb{Q},\star}|),$$

then the centered option-facing loading is asymptotically aggregate-factor aligned.

Proof. By Definition D.1,

$$\Xi_{M,\tau} = q_\tau X_\tau^2 Z_\tau = q_\tau v_\tau Z_\tau + J_{M,\tau}.$$

Applying $C_{\mathcal{H}_\omega}^\mathbb{Q}$ and using the definition of $L_{Z_\tau}^\mathbb{Q}$ gives

$$L_{M,\tau}^\mathbb{Q} = q_\tau v_\tau L_{Z_\tau}^\mathbb{Q} + C_{\mathcal{H}_\omega}^\mathbb{Q}(J_{M,\tau}).$$

Apply $C_{\mathcal{G}_T}^\mathbb{Q}$ and subtract:

$$\tilde{L}_{M,\tau}^\mathbb{Q} := L_{M,\tau}^\mathbb{Q} - C_{\mathcal{G}_T}^\mathbb{Q}(L_{M,\tau}^\mathbb{Q}) = q_\tau v_\tau (L_{Z_\tau}^\mathbb{Q} - C_{\mathcal{G}_T}^\mathbb{Q}(L_{Z_\tau}^\mathbb{Q})) + \tilde{J}_{M,\tau}^\mathbb{Q}.$$

Using the centered part of the factor decomposition of $L_{Z_\tau}^\mathbb{Q}$,

$$L_{Z_\tau}^\mathbb{Q} - C_{\mathcal{G}_T}^\mathbb{Q}(L_{Z_\tau}^\mathbb{Q}) = b_{M,\tau}^{\mathbb{Q},\star} \Lambda_M^\star + R_{Z_\tau}^{\mathbb{Q},\star},$$

yields the exact identity for $\tilde{L}_{M,\tau}^\mathbb{Q}$. Now apply Proposition 4.6 to $\tilde{L}_{M,\tau}^\mathbb{Q}$. Since

$$\alpha_{M,\tau}^{\mathbb{Q},\star} = q_\tau v_\tau b_{M,\tau}^{\mathbb{Q},\star} + q_\tau v_\tau \langle R_{Z_\tau}^{\mathbb{Q},\star}, \Lambda_M^\star \rangle_{L^2(\mathbb{Q})} + \langle \tilde{J}_{M,\tau}^\mathbb{Q}, \Lambda_M^\star \rangle_{L^2(\mathbb{Q})},$$

and $R_{Z_\tau}^{\mathbb{Q},\star} \perp \Lambda_M^\star$, the coefficient error reduces to

$$\alpha_{M,\tau}^{\mathbb{Q},\star} - q_\tau v_\tau b_{M,\tau}^{\mathbb{Q},\star} = \langle \tilde{J}_{M,\tau}^\mathbb{Q}, \Lambda_M^\star \rangle_{L^2(\mathbb{Q})},$$

so Cauchy–Schwarz gives the coefficient bound. The residual bound follows by taking the orthogonal complement of $\text{span}\{\Lambda_M^\star\}$ in the displayed identity for $\tilde{L}_{M,\tau}^\mathbb{Q}$. If the last displayed smallness condition holds, then

$$|\alpha_{M,\tau}^{\mathbb{Q},\star} - q_\tau v_\tau b_{M,\tau}^{\mathbb{Q},\star}| = o(q_\tau v_\tau |b_{M,\tau}^{\mathbb{Q},\star}|),$$

and similarly

$$\|R_{M,\tau}^{\mathbb{Q},\star}\|_{L^2(\mathbb{Q})} = o(q_\tau v_\tau |b_{M,\tau}^{\mathbb{Q},\star}|).$$

Therefore

$$\frac{\|R_{M,\tau}^{\mathbb{Q},\star}\|_{L^2(\mathbb{Q})}}{|\alpha_{M,\tau}^{\mathbb{Q},\star}|} \rightarrow 0,$$

which is exactly aggregate-factor alignment in the sense of Theorem D.14. $\square$

Proposition D.18 (Moment control of the quadratic-fluctuation remainder). *Under Assumption 6.8,*

$$\|\tilde{J}_{M,\tau}^{\mathbb{Q}}\|_{L^2(\mathbb{Q})} \leq \|J_{M,\tau}\|_{L^2(\mathbb{Q})} \leq q_\tau \|X_\tau^2 - v_\tau\|_{L^4(\mathbb{Q})} \|Z_\tau\|_{L^4(\mathbb{Q})}.$$

Moreover, if $G \sim N(0, 1)$ and $c_G := \|G^2 - 1\|_{L^4}$, then

$$\|X_\tau^2 - v_\tau\|_{L^4(\mathbb{Q})} = c_G v_\tau,$$

so

$$\|\tilde{J}_{M,\tau}^{\mathbb{Q}}\|_{L^2(\mathbb{Q})} \leq c_G q_\tau v_\tau \|Z_\tau\|_{L^4(\mathbb{Q})}.$$

In particular, by Burkholder–Davis–Gundy there exists a universal constant $C_{\text{BDG}} > 0$ such that

$$\|\tilde{J}_{M,\tau}^{\mathbb{Q}}\|_{L^2(\mathbb{Q})} \leq C_{\text{BDG}} c_G q_\tau v_\tau \left(\mathbb{E}_{\mathbb{Q}} \left[\left(\int_0^\tau Y_M(u)^2 du \right)^2 \right] \right)^{1/4}.$$

Proof. Because $L^2(\mathbb{Q}, \mathcal{G}_T) \subseteq L^2(\mathbb{Q}, \mathcal{H}_\varpi)$, the term

$$\tilde{J}_{M,\tau}^{\mathbb{Q}} = C_{\mathcal{H}_\varpi}^{\mathbb{Q}}(J_{M,\tau}) - C_{\mathcal{G}_T}^{\mathbb{Q}}(C_{\mathcal{H}_\varpi}^{\mathbb{Q}}(J_{M,\tau}))$$

is the orthogonal remainder of $C_{\mathcal{H}_\varpi}^{\mathbb{Q}}(J_{M,\tau})$ relative to $L^2(\mathbb{Q}, \mathcal{G}_T)$. Therefore

$$\|\tilde{J}_{M,\tau}^{\mathbb{Q}}\|_{L^2(\mathbb{Q})} \leq \|C_{\mathcal{H}_\varpi}^{\mathbb{Q}}(J_{M,\tau})\|_{L^2(\mathbb{Q})} \leq \|J_{M,\tau}\|_{L^2(\mathbb{Q})}.$$

Now

$$J_{M,\tau} = q_\tau (X_\tau^2 - v_\tau) Z_\tau,$$

so Hölder gives

$$\|J_{M,\tau}\|_{L^2(\mathbb{Q})} \leq q_\tau \|X_\tau^2 - v_\tau\|_{L^4(\mathbb{Q})} \|Z_\tau\|_{L^4(\mathbb{Q})}.$$

Since $X_\tau/\sqrt{v_\tau} \sim N(0, 1)$,

$$X_\tau^2 - v_\tau = v_\tau \left(\left(\frac{X_\tau}{\sqrt{v_\tau}} \right)^2 - 1 \right),$$

which yields

$$\|X_\tau^2 - v_\tau\|_{L^4(\mathbb{Q})} = c_G v_\tau.$$

Finally, the L^4 Burkholder–Davis–Gundy inequality gives

$$\|Z_\tau\|_{L^4(\mathbb{Q})} \leq C_{\text{BDG}} \left(\mathbb{E}_{\mathbb{Q}} \left[\left(\int_0^\tau Y_M(u)^2 du \right)^2 \right] \right)^{1/4},$$

and substitution proves the last display. $\square$

Proposition D.19 (Kernel-energy control of the quadratic-fluctuation remainder). *Under Assumption 6.8, define the deterministic synchronized-perturbation kernel energy*

$$\mathfrak{D}_{M,\tau}^{\text{pert}} := \sum_{j=1}^N \int_0^\tau \int_0^u |G_j(u, s)|^2 ds du.$$

Then

$$\left(\mathbb{E}_{\mathbb{Q}} \left[\left(\int_0^\tau Y_M(u)^2 du \right)^2 \right] \right)^{1/4} \leq 3^{1/4} (\mathfrak{D}_{M,\tau}^{\text{pert}})^{1/2},$$

and therefore

$$\|\tilde{J}_{M,\tau}^{\mathbb{Q}}\|_{L^2(\mathbb{Q})} \leq 3^{1/4} C_{\text{BDG}} c_G q_\tau v_\tau (\mathfrak{D}_{M,\tau}^{\text{pert}})^{1/2}.$$

If, in addition,

$$\mathfrak{D}_{M,\tau}^{\text{pert}} = 0,$$

then

$$\tilde{J}_{M,\tau}^{\mathbb{Q}} = 0.$$

Proof. Set

$$\Sigma_{M,\tau}(u) := \mathbb{E}_{\mathbb{Q}}[Y_M(u)^2] = \sum_{j=1}^N \int_0^u |G_j(u, s)|^2 ds, \quad 0 \leq u \leq \tau,$$

where the second identity is Itô isometry and independence of the $\mathbb{Q}$ -Brownian coordinates. Because $Y_M(u)$ is centered Gaussian for every u , Wick's formula gives

$$\mathbb{E}_{\mathbb{Q}}[Y_M(u)^2 Y_M(v)^2] = \Sigma_{M,\tau}(u) \Sigma_{M,\tau}(v) + 2 \operatorname{Cov}_{\mathbb{Q}}(Y_M(u), Y_M(v))^2.$$

By Cauchy–Schwarz,

$$\operatorname{Cov}_{\mathbb{Q}}(Y_M(u), Y_M(v))^2 \leq \Sigma_{M,\tau}(u) \Sigma_{M,\tau}(v),$$

so

$$\mathbb{E}_{\mathbb{Q}}[Y_M(u)^2 Y_M(v)^2] \leq 3 \Sigma_{M,\tau}(u) \Sigma_{M,\tau}(v).$$

Integrating over $u, v \in [0, \tau]$ yields

$$\mathbb{E}_{\mathbb{Q}} \left[\left(\int_0^\tau Y_M(u)^2 du \right)^2 \right] \leq 3 \left(\int_0^\tau \Sigma_{M,\tau}(u) du \right)^2 = 3(\mathfrak{D}_{M,\tau}^{\text{pert}})^2,$$

which proves the first display after taking fourth roots. The second display then follows from the last bound in Proposition D.18. If $\mathfrak{D}_{M,\tau}^{\text{pert}} = 0$, the second display gives $\|\tilde{J}_{M,\tau}^{\mathbb{Q}}\|_{L^2(\mathbb{Q})} = 0$, hence $\tilde{J}_{M,\tau}^{\mathbb{Q}} = 0$. $\square$

Remark D.20 (What the Gaussian reduction theorem localizes). Proposition D.17 localizes the option-factor problem in the Gaussian-Volterra block. The two primitive ingredients are the factor decomposition of the perturbation return loading $L_{Z_\tau}^{\mathbb{Q}}$ and the centered size of the quadratic-fluctuation remainder $J_{M,\tau}$. Thus the localized task is to prove that the centered loading of Z_τ is itself aggregate-factor dominated. The Perron-projected kernel construction introduced below sharpens that statement one step further. It isolates a canonical Perron-generated proxy claim for the perturbation return and an explicit off-Perron kernel defect. Proposition D.18 shows that the quadratic remainder is already controlled by a fourth-moment perturbation-energy bound, so the remaining work is pushed into the proxy-compatibility error and the off-Perron kernel defect.

Proposition D.21 (Claim-level Perron approximation controls the perturbation-return loading). *Assume Assumption 6.8, and suppose the aggregate factor benchmark is nontrivial, so*

$$\Lambda_M^* \in \mathcal{M}_{\mathbb{Q}}^{\circ}, \quad \|\Lambda_M^*\|_{L^2(\mathbb{Q})} = 1.$$

Fix a sufficiently small maturity $\tau > 0$, and suppose the perturbation return admits a decomposition

$$Z_\tau = Z_{M,\tau}^{\varpi} + E_{Z,\tau},$$

where $Z_{M,\tau}^{\varpi} \in L^2(\mathbb{Q}, \mathcal{H}_{\varpi})$ and

$$Z_{M,\tau}^{\varpi} = Z_{M,\tau}^{\text{vis}} + b_{M,\tau}^{\varpi,*} \Lambda_M^*,$$

with

$$Z_{M,\tau}^{\text{vis}} \in L^2(\mathbb{Q}, \mathcal{G}_T), \quad b_{M,\tau}^{\varpi,*} \in \mathbb{R}, \quad E_{Z,\tau} \in L^2(\mathbb{Q}, \mathcal{F}_T).$$

Define

$$\tilde{E}_{Z,\tau}^{\mathbb{Q}} := C_{\mathcal{H}_{\varpi}}^{\mathbb{Q}}(E_{Z,\tau}) - C_{\mathcal{G}_T}^{\mathbb{Q}}(C_{\mathcal{H}_{\varpi}}^{\mathbb{Q}}(E_{Z,\tau})).$$

Then

$$L_{Z_\tau}^{\mathbb{Q}} = Z_{M,\tau}^{\text{vis}} + b_{M,\tau}^{\varpi,*} \Lambda_M^* + C_{\mathcal{H}_{\varpi}}^{\mathbb{Q}}(E_{Z,\tau}),$$

and

$$L_{Z_\tau}^{\mathbb{Q}} - C_{\mathcal{G}_T}^{\mathbb{Q}}(L_{Z_\tau}^{\mathbb{Q}}) = b_{M,\tau}^{\varpi,*} \Lambda_M^* + \tilde{E}_{Z,\tau}^{\mathbb{Q}}.$$

Consequently, in the canonical factor decomposition

$$L_{Z_\tau}^{\mathbb{Q}} = V_{Z_\tau}^{\mathbb{Q}} + b_{M,\tau}^{\varpi,*} \Lambda_M^* + R_{Z_\tau}^{\mathbb{Q},*},$$

one has

$$|b_{M,\tau}^{\mathbb{Q},*} - b_{M,\tau}^{\varpi,*}| \leq \|E_{Z,\tau}\|_{L^2(\mathbb{Q})}, \quad \|R_{Z_\tau}^{\mathbb{Q},*}\|_{L^2(\mathbb{Q})} \leq \|E_{Z,\tau}\|_{L^2(\mathbb{Q})}.$$

If, in addition,

$$b_{M,\tau}^{\varpi,\star} \neq 0 \quad \text{for all sufficiently small } \tau, \quad \|E_{Z,\tau}\|_{L^2(\mathbb{Q})} = o(|b_{M,\tau}^{\varpi,\star}|),$$

then the centered perturbation-return loading is asymptotically aggregate-factor aligned.

Proof. Because $Z_{M,\tau}^{\varpi}$ is $\mathcal{H}_\varpi$ -measurable,

$$L_{Z_\tau}^{\mathbb{Q}} = C_{\mathcal{H}_\varpi}^{\mathbb{Q}}(Z_\tau) = Z_{M,\tau}^{\varpi} + C_{\mathcal{H}_\varpi}^{\mathbb{Q}}(E_{Z,\tau}),$$

which gives the first displayed identity. Applying $C_{\mathcal{G}_T}^{\mathbb{Q}}$ and using

$$C_{\mathcal{G}_T}^{\mathbb{Q}}(\Lambda_M^\star) = 0$$

from Corollary 4.7 yields

$$C_{\mathcal{G}_T}^{\mathbb{Q}}(L_{Z_\tau}^{\mathbb{Q}}) = Z_{M,\tau}^{\text{vis}} + C_{\mathcal{G}_T}^{\mathbb{Q}}(C_{\mathcal{H}_\varpi}^{\mathbb{Q}}(E_{Z,\tau})),$$

so subtracting gives the centered identity with $\tilde{E}_{Z,\tau}^{\mathbb{Q}}$. Since

$$\|\tilde{E}_{Z,\tau}^{\mathbb{Q}}\|_{L^2(\mathbb{Q})} \leq \|E_{Z,\tau}\|_{L^2(\mathbb{Q})}$$

by the same orthogonal-remainder argument as in Proposition D.18, Proposition 4.6 gives

$$|b_{M,\tau}^{\mathbb{Q},\star} - b_{M,\tau}^{\varpi,\star}| \leq \|\tilde{E}_{Z,\tau}^{\mathbb{Q}}\|_{L^2(\mathbb{Q})}, \quad \|R_{Z_\tau}^{\mathbb{Q},\star}\|_{L^2(\mathbb{Q})} \leq \|\tilde{E}_{Z,\tau}^{\mathbb{Q}}\|_{L^2(\mathbb{Q})},$$

hence the displayed bounds. If $\|E_{Z,\tau}\|_{L^2(\mathbb{Q})} = o(|b_{M,\tau}^{\varpi,\star}|)$, then

$$\|R_{Z_\tau}^{\mathbb{Q},\star}\|_{L^2(\mathbb{Q})} = o(|b_{M,\tau}^{\varpi,\star}|),$$

and

$$|b_{M,\tau}^{\mathbb{Q},\star}| \geq |b_{M,\tau}^{\varpi,\star}| - \|E_{Z,\tau}\|_{L^2(\mathbb{Q})} \sim |b_{M,\tau}^{\varpi,\star}|,$$

which proves asymptotic aggregate-factor alignment of the centered perturbation-return loading. $\square$

Proposition D.22 (Off-Perron perturbation energy controls the perturbation-return factor error). *Assume the hypotheses of Proposition D.21. Suppose, moreover, that the claim-level Perron approximation error is realized as a stochastic integral remainder:*

$$E_{Z,\tau} = \int_0^\tau Y_{M,\tau}^\perp(u) dB_u^{\mathbb{Q}}$$

for some predictable process $Y_{M,\tau}^\perp$ satisfying

$$\mathbb{E}_{\mathbb{Q}} \left[\int_0^\tau (Y_{M,\tau}^\perp(u))^2 du \right] < \infty.$$

Then

$$\|E_{Z,\tau}\|_{L^2(\mathbb{Q})} = \left(\mathbb{E}_{\mathbb{Q}} \left[\int_0^\tau (Y_{M,\tau}^\perp(u))^2 du \right] \right)^{1/2}.$$

Consequently,

$$|b_{M,\tau}^{\mathbb{Q},\star} - b_{M,\tau}^{\varpi,\star}| \leq \left(\mathbb{E}_{\mathbb{Q}} \left[\int_0^\tau (Y_{M,\tau}^\perp(u))^2 du \right] \right)^{1/2},$$

and

$$\|R_{Z_\tau}^{\mathbb{Q},\star}\|_{L^2(\mathbb{Q})} \leq \left(\mathbb{E}_{\mathbb{Q}} \left[\int_0^\tau (Y_{M,\tau}^\perp(u))^2 du \right] \right)^{1/2}.$$

If, in addition,

$$b_{M,\tau}^{\varpi,\star} \neq 0 \quad \text{for all sufficiently small } \tau,$$

and

$$\mathbb{E}_{\mathbb{Q}} \left[\int_0^\tau (Y_{M,\tau}^\perp(u))^2 du \right] = o(|b_{M,\tau}^{\varpi,\star}|^2),$$

then the centered perturbation-return loading is asymptotically aggregate-factor aligned.

Proof. The identity for $\|E_{Z,\tau}\|_{L^2(\mathbb{Q})}$ is Itô isometry. Substituting that identity into Proposition D.21 yields the two displayed bounds. If the last displayed smallness condition holds, then

$$\|E_{Z,\tau}\|_{L^2(\mathbb{Q})} = o(|b_{M,\tau}^{\varpi,\star}|),$$

so the final implication of Proposition D.21 gives asymptotic aggregate-factor alignment of the centered perturbation-return loading. $\square$

Proposition D.23 (Channel-side decomposition of the synchronized perturbation implies perturbation-return factor control). *Assume Assumption 6.8, and suppose the aggregate factor benchmark is nontrivial. Fix a sufficiently small maturity $\tau > 0$, and suppose there exist predictable processes*

$$Y_{M,\tau}^{\varpi}, \quad Y_{M,\tau}^{\perp},$$

such that

$$Y_M(u) = Y_{M,\tau}^{\varpi}(u) + Y_{M,\tau}^{\perp}(u), \quad 0 \leq u \leq \tau.$$

Define

$$Z_{M,\tau}^{\varpi} := \int_0^\tau Y_{M,\tau}^{\varpi}(u) dB_u^{\mathbb{Q}}, \quad E_{Z,\tau} := \int_0^\tau Y_{M,\tau}^{\perp}(u) dB_u^{\mathbb{Q}}.$$

Assume further that

$$Z_{M,\tau}^{\varpi} = Z_{M,\tau}^{\text{vis}} + b_{M,\tau}^{\varpi,\star} \Lambda_M^{\star}$$

for some

$$Z_{M,\tau}^{\text{vis}} \in L^2(\mathbb{Q}, \mathcal{G}_T), \quad b_{M,\tau}^{\varpi,\star} \in \mathbb{R}.$$

Then the hypotheses of Proposition D.22 hold with this choice of $E_{Z,\tau}$, and

$$|b_{M,\tau}^{\mathbb{Q},\star} - b_{M,\tau}^{\varpi,\star}| \leq \left(\mathbb{E}_{\mathbb{Q}} \left[\int_0^\tau (Y_{M,\tau}^{\perp}(u))^2 du \right] \right)^{1/2},$$

and

$$\|R_{Z,\tau}^{\mathbb{Q},\star}\|_{L^2(\mathbb{Q})} \leq \left(\mathbb{E}_{\mathbb{Q}} \left[\int_0^\tau (Y_{M,\tau}^{\perp}(u))^2 du \right] \right)^{1/2}.$$

If, in addition,

$$b_{M,\tau}^{\varpi,\star} \neq 0 \quad \text{for all sufficiently small } \tau,$$

and

$$\mathbb{E}_{\mathbb{Q}} \left[\int_0^\tau (Y_{M,\tau}^{\perp}(u))^2 du \right] = o(|b_{M,\tau}^{\varpi,\star}|^2),$$

then the centered perturbation-return loading is asymptotically aggregate-factor aligned.

Proof. The definitions of $Z_{M,\tau}^{\varpi}$ and $E_{Z,\tau}$ give

$$Z_\tau = \int_0^\tau Y_M(u) dB_u^{\mathbb{Q}} = Z_{M,\tau}^{\varpi} + E_{Z,\tau}.$$

The assumed representation of $Z_{M,\tau}^{\varpi}$ is exactly the decomposition required in Proposition D.21. Proposition D.22 then applies with the present stochastic-integral remainder and gives the displayed bounds and the final alignment implication. $\square$

Remark D.24 (Descending Gaussian option chain). The Gaussian option block now follows a descending chain of reductions. First, Proposition D.17 reduces option-factor alignment to the perturbation return Z_τ and the quadratic-fluctuation remainder $\tilde{J}_{M,\tau}^{\mathbb{Q}}$. Next, Propositions D.21, D.22, and D.23 reduce the perturbation-return factor problem to the off-Perron energy of the synchronized perturbation process. The canonical Perron construction below then isolates a Perron-generated proxy claim and pushes the remaining structural uncertainty into the off-Perron kernel defect. The Gaussian subsection can therefore be read as one ladder:

option claim $\implies$ perturbation return $\implies$ Perron proxy claim

Perron proxy claim $\implies$ perturbation process $\implies$ kernel defect.

Corollary D.25 (Kernel-side off-Perron defect criterion for option-factor alignment). *Assume Assumption 6.8, and suppose that for each sufficiently small maturity $\tau > 0$ and each $1 \leq j \leq N$ the synchronized perturbation kernel admits a decomposition*

$$G_j(u, s) = G_{j,\tau}^\varpi(u, s) + G_{j,\tau}^\perp(u, s), \quad 0 \leq s \leq u \leq \tau,$$

with corresponding processes

$$Y_{M,\tau}^\varpi(u) := \sum_{j=1}^N \int_0^u G_{j,\tau}^\varpi(u, s) dW_s^{\mathbb{Q},j}, \quad Y_{M,\tau}^\perp(u) := \sum_{j=1}^N \int_0^u G_{j,\tau}^\perp(u, s) dW_s^{\mathbb{Q},j}.$$

Assume further that

$$\int_0^\tau Y_{M,\tau}^\varpi(u) dB_u^\mathbb{Q} = Z_{M,\tau}^{\text{vis}} + b_{M,\tau}^{\varpi,\star} \Lambda_M^\star$$

for some $Z_{M,\tau}^{\text{vis}} \in L^2(\mathbb{Q}, \mathcal{G}_T)$, with

$$b_{M,\tau}^{\varpi,\star} \neq 0 \quad \text{for all sufficiently small } \tau.$$

If

$$\sum_{j=1}^N \int_0^\tau \int_0^u |G_{j,\tau}^\perp(u, s)|^2 ds du = o(|b_{M,\tau}^{\varpi,\star}|^2),$$

and

$$\left(\mathbb{E}_\mathbb{Q} \left[\left(\int_0^\tau Y_M(u)^2 du \right)^2 \right] \right)^{1/4} = o(|b_{M,\tau}^{\varpi,\star}|),$$

then the centered option-facing loading is asymptotically aggregate-factor aligned.

Proof. Set

$$Y_{M,\tau}^\perp(u) := \sum_{j=1}^N \int_0^u G_{j,\tau}^\perp(u, s) dW_s^{\mathbb{Q},j}.$$

Because the $\mathbb{Q}$ -Brownian coordinates are independent, Itô isometry gives

$$\mathbb{E}_\mathbb{Q} \left[\int_0^\tau (Y_{M,\tau}^\perp(u))^2 du \right] = \sum_{j=1}^N \int_0^\tau \int_0^u |G_{j,\tau}^\perp(u, s)|^2 ds du.$$

Hence the kernel defect hypothesis implies the off-Perron perturbation-energy condition in Proposition D.23. The visible-plus-factor representation of $\int_0^\tau Y_{M,\tau}^\varpi(u) dB_u^\mathbb{Q}$ is exactly the structural hypothesis of that proposition. Therefore Proposition D.23 gives

$$\|R_{Z_\tau}^{\mathbb{Q},\star}\|_{L^2(\mathbb{Q})} = o(|b_{M,\tau}^{\varpi,\star}|),$$

and

$$|b_{M,\tau}^{\mathbb{Q},\star}| \sim |b_{M,\tau}^{\varpi,\star}|.$$

Proposition D.18 then yields

$$\|\tilde{J}_{M,\tau}^\mathbb{Q}\|_{L^2(\mathbb{Q})} = o(q_\tau v_\tau |b_{M,\tau}^{\mathbb{Q},\star}|).$$

Thus

$$q_\tau v_\tau \|R_{Z_\tau}^{\mathbb{Q},\star}\|_{L^2(\mathbb{Q})} + \|\tilde{J}_{M,\tau}^\mathbb{Q}\|_{L^2(\mathbb{Q})} = o(q_\tau v_\tau |b_{M,\tau}^{\mathbb{Q},\star}|),$$

and Proposition D.17 gives asymptotic aggregate-factor alignment of the centered option-facing loading. $\square$

Corollary D.26 (Kernel-side off-Perron defect and deterministic synchronized-perturbation kernel energy criterion for option-factor alignment). *Assume the hypotheses of Corollary D.25, except that instead of assuming*

$$\left(\mathbb{E}_\mathbb{Q} \left[\left(\int_0^\tau Y_M(u)^2 du \right)^2 \right] \right)^{1/4} = o(|b_{M,\tau}^{\varpi,\star}|),$$

assume

$$(\mathfrak{D}_{M,\tau}^{\text{pert}})^{1/2} = o(|b_{M,\tau}^{\varpi,\star}|).$$

Then the centered option-facing loading is asymptotically aggregate-factor aligned.

Proof. By Proposition D.19,

$$\left(\mathbb{E}_{\mathbb{Q}} \left[\left(\int_0^\tau Y_M(u)^2 du \right)^2 \right] \right)^{1/4} = O((\mathfrak{D}_{M,\tau}^{\text{pert}})^{1/2}),$$

so the new hypothesis implies the moment condition in Corollary D.25. The conclusion follows. $\square$

Definition D.27 (Canonical Perron-projected perturbation kernels and proxy claim). Under Definition 2.5 and Assumption 6.8, let

$$P_\varpi := vv^\top, \quad P_\perp := I_N - P_\varpi,$$

where I_N is the identity on $\mathbb{R}^N$. For $1 \leq j \leq N$ and $0 \leq s \leq u \leq T$, define the channel vector

$$K_j(u, s) := (\dot{\beta}_{ij} a_{ij}(u, s) (u - s)^{H_{ij}-1/2} \mathbf{1}_{\{(i,j) \in \mathcal{E}_M\}})_{i=1}^N \in \mathbb{R}^N.$$

Then

$$G_j(u, s) = b^\top K_j(u, s).$$

Define the canonical Perron-projected channel kernels

$$K_j^\varpi(u, s) := P_\varpi K_j(u, s), \quad K_j^\perp(u, s) := P_\perp K_j(u, s),$$

and the induced market kernels

$$G_j^\varpi(u, s) := b^\top K_j^\varpi(u, s), \quad G_j^\perp(u, s) := b^\top K_j^\perp(u, s).$$

For each maturity $\tau > 0$, define

$$\begin{aligned} Y_M^\varpi(u) &:= \sum_{j=1}^N \int_0^u G_j^\varpi(u, s) dW_s^{\mathbb{Q},j}, & Y_M^\perp(u) &:= \sum_{j=1}^N \int_0^u G_j^\perp(u, s) dW_s^{\mathbb{Q},j}, \\ Z_{M,\tau}^\varpi &:= \int_0^\tau Y_M^\varpi(u) dB_u^\mathbb{Q}, & E_{Z,\tau}^\perp &:= \int_0^\tau Y_M^\perp(u) dB_u^\mathbb{Q}. \end{aligned}$$

The canonical off-Perron kernel defect is

$$\mathfrak{D}_{M,\tau}^\perp := \sum_{j=1}^N \int_0^\tau \int_0^u |G_j^\perp(u, s)|^2 ds du.$$

Finally, define the centered Perron-generated proxy claim

$$\Upsilon_{M,\tau}^\varpi := C_{\mathcal{H}_\varpi}^\mathbb{Q}(Z_{M,\tau}^\varpi) - C_{\mathcal{G}_T}^\mathbb{Q}(C_{\mathcal{H}_\varpi}^\mathbb{Q}(Z_{M,\tau}^\varpi)) \in \mathcal{M}_\mathbb{Q}^\circ.$$

Remark D.28 (The Perron projection is algebraic). The projection $P_\varpi = vv^\top$ is the algebraic Perron projection associated with the left-right normalization $u^\top v = 1$. In general the projection is oblique. Thus $P_\perp = I_N - P_\varpi$ denotes the complementary Perron residual map; norms of $P_\perp$ -terms are residual norms, with Pythagorean orthogonal-energy interpretations requiring additional self-adjointness or a compatible-inner-product condition.

Proposition D.29 (Rank-one Perron reduction of synchronized perturbation kernels). *Under Definition D.27, define the scalar Perron profiles*

$$\psi_j(u, s) := u^\top K_j(u, s), \quad 1 \leq j \leq N, \quad 0 \leq s \leq u \leq T.$$

Then

$$K_j^\varpi(u, s) = v \psi_j(u, s), \quad K_j^\perp(u, s) = K_j(u, s) - v \psi_j(u, s),$$

and therefore

$$G_j^\varpi(u, s) = (b^\top v) \psi_j(u, s).$$

Thus, after Perron projection, the synchronized perturbation kernels are rank one in the asset index and the remaining structure is carried entirely by the scalar profiles ψ_j .

Proof. By definition,

$$K_j^\varpi(u, s) = P_\varpi K_j(u, s) = vu^\top K_j(u, s) = v\psi_j(u, s).$$

The identity for $K_j^\perp$ follows from $P_\perp = I_N - P_\varpi$, and then

$$G_j^\varpi(u, s) = b^\top K_j^\varpi(u, s) = (b^\top v)\psi_j(u, s).$$

□

Proposition D.30 (Canonical Perron proxy claim and kernel-defect control). *Under Definition D.27 and Assumption 6.8, for every sufficiently small maturity $\tau > 0$,*

$$Y_M(u) = Y_M^\varpi(u) + Y_M^\perp(u), \quad 0 \leq u \leq \tau,$$

and

$$Z_\tau = Z_{M,\tau}^\varpi + E_{Z,\tau}^\perp.$$

Moreover,

$$\|E_{Z,\tau}^\perp\|_{L^2(\mathbb{Q})}^2 = \mathfrak{D}_{M,\tau}^\perp,$$

$$\|\tilde{L}_{Z,\tau}^\mathbb{Q} - \Upsilon_{M,\tau}^\varpi\|_{L^2(\mathbb{Q})} \leq (\mathfrak{D}_{M,\tau}^\perp)^{1/2},$$

and, with q_τ and v_τ as in Proposition D.17,

$$\|\tilde{L}_{M,\tau}^\mathbb{Q} - q_\tau v_\tau \Upsilon_{M,\tau}^\varpi\|_{L^2(\mathbb{Q})} \leq q_\tau v_\tau (\mathfrak{D}_{M,\tau}^\perp)^{1/2} + \|\tilde{J}_{M,\tau}^\mathbb{Q}\|_{L^2(\mathbb{Q})}.$$

Proof. Because $P_\varpi + P_\perp = I_N$, one has

$$K_j(u, s) = K_j^\varpi(u, s) + K_j^\perp(u, s), \quad G_j(u, s) = G_j^\varpi(u, s) + G_j^\perp(u, s).$$

Substituting these identities into the stochastic-integral definitions yields

$$Y_M(u) = Y_M^\varpi(u) + Y_M^\perp(u), \quad Z_\tau = Z_{M,\tau}^\varpi + E_{Z,\tau}^\perp.$$

Since $E_{Z,\tau}^\perp$ is a stochastic integral against $B^\mathbb{Q}$, Itô isometry gives

$$\|E_{Z,\tau}^\perp\|_{L^2(\mathbb{Q})}^2 = \mathbb{E}_\mathbb{Q} \left[\int_0^\tau (Y_M^\perp(u))^2 du \right].$$

Using independence of the $\mathbb{Q}$ -Brownian coordinates,

$$\mathbb{E}_\mathbb{Q} \left[\int_0^\tau (Y_M^\perp(u))^2 du \right] = \sum_{j=1}^N \int_0^\tau \int_0^u |G_j^\perp(u, s)|^2 ds du = \mathfrak{D}_{M,\tau}^\perp.$$

Applying $C_{\mathcal{H}_\varpi}^\mathbb{Q}$ to $Z_\tau = Z_{M,\tau}^\varpi + E_{Z,\tau}^\perp$ gives

$$L_{Z,\tau}^\mathbb{Q} = C_{\mathcal{H}_\varpi}^\mathbb{Q}(Z_{M,\tau}^\varpi) + C_{\mathcal{H}_\varpi}^\mathbb{Q}(E_{Z,\tau}^\perp).$$

Subtracting the $\mathcal{G}_T$ -projection yields

$$\tilde{L}_{Z,\tau}^\mathbb{Q} = \Upsilon_{M,\tau}^\varpi + C_{\mathcal{H}_\varpi}^\mathbb{Q}(E_{Z,\tau}^\perp) - C_{\mathcal{G}_T}^\mathbb{Q}(C_{\mathcal{H}_\varpi}^\mathbb{Q}(E_{Z,\tau}^\perp)),$$

so

$$\|\tilde{L}_{Z,\tau}^\mathbb{Q} - \Upsilon_{M,\tau}^\varpi\|_{L^2(\mathbb{Q})} \leq \|E_{Z,\tau}^\perp\|_{L^2(\mathbb{Q})} = (\mathfrak{D}_{M,\tau}^\perp)^{1/2}.$$

Finally, the proof of Proposition D.17 gives the exact identity

$$\tilde{L}_{M,\tau}^\mathbb{Q} = q_\tau v_\tau \tilde{L}_{Z,\tau}^\mathbb{Q} + \tilde{J}_{M,\tau}^\mathbb{Q}.$$

Subtract $q_\tau v_\tau \Upsilon_{M,\tau}^\varpi$ and apply the previous bound. □

Theorem D.31 (Canonical Perron proxy compatibility implies option-factor alignment). *Assume the hypotheses of Proposition D.30, and suppose the aggregate factor benchmark is nontrivial. Fix a sufficiently small maturity $\tau > 0$, and suppose the centered Perron-generated proxy claim admits the factor decomposition*

$$\Upsilon_{M,\tau}^\varpi = c_{M,\tau}^\varpi \Lambda_M^\star + D_{M,\tau}^\varpi,$$

where

$$c_{M,\tau}^\varpi \in \mathbb{R}, \quad D_{M,\tau}^\varpi \in \mathcal{M}_\mathbb{Q}^\circ, \quad D_{M,\tau}^\varpi \perp_{L^2(\mathbb{Q})} \Lambda_M^\star.$$

Then

$$\left| \alpha_{M,\tau}^{\mathbb{Q},\star} - q_\tau v_\tau c_{M,\tau}^\varpi \right| \leq q_\tau v_\tau \|D_{M,\tau}^\varpi\|_{L^2(\mathbb{Q})} + q_\tau v_\tau (\mathfrak{D}_{M,\tau}^\perp)^{1/2} + \|\tilde{J}_{M,\tau}^\mathbb{Q}\|_{L^2(\mathbb{Q})},$$

and

$$\|R_{M,\tau}^{\mathbb{Q},\star}\|_{L^2(\mathbb{Q})} \leq q_\tau v_\tau \|D_{M,\tau}^\varpi\|_{L^2(\mathbb{Q})} + q_\tau v_\tau (\mathfrak{D}_{M,\tau}^\perp)^{1/2} + \|\tilde{J}_{M,\tau}^\mathbb{Q}\|_{L^2(\mathbb{Q})}.$$

If, moreover,

$$c_{M,\tau}^\varpi \neq 0 \quad \text{for all sufficiently small } \tau,$$

and

$$\|D_{M,\tau}^\varpi\|_{L^2(\mathbb{Q})} + (\mathfrak{D}_{M,\tau}^\perp)^{1/2} + \frac{\|\tilde{J}_{M,\tau}^\mathbb{Q}\|_{L^2(\mathbb{Q})}}{q_\tau v_\tau} = o(|c_{M,\tau}^\varpi|),$$

then the centered option-facing loading is asymptotically aggregate-factor aligned. If in addition

$$D_{M,\tau}^\varpi = 0, \quad \mathfrak{D}_{M,\tau}^\perp = 0, \quad \tilde{J}_{M,\tau}^\mathbb{Q} = 0,$$

then

$$\tilde{L}_{M,\tau}^\mathbb{Q} = q_\tau v_\tau c_{M,\tau}^\varpi \Lambda_M^\star, \quad R_{M,\tau}^{\mathbb{Q},\star} = 0,$$

so the centered option-facing loading is exactly aggregate-factor aligned.

Proof. Set

$$\mathcal{E}_{M,\tau} := \tilde{L}_{M,\tau}^\mathbb{Q} - q_\tau v_\tau \Upsilon_{M,\tau}^\varpi.$$

By Proposition D.30,

$$\|\mathcal{E}_{M,\tau}\|_{L^2(\mathbb{Q})} \leq q_\tau v_\tau (\mathfrak{D}_{M,\tau}^\perp)^{1/2} + \|\tilde{J}_{M,\tau}^\mathbb{Q}\|_{L^2(\mathbb{Q})}.$$

Substituting the factor decomposition of $\Upsilon_{M,\tau}^\varpi$ gives

$$\tilde{L}_{M,\tau}^\mathbb{Q} = q_\tau v_\tau c_{M,\tau}^\varpi \Lambda_M^\star + q_\tau v_\tau D_{M,\tau}^\varpi + \mathcal{E}_{M,\tau}.$$

Apply Proposition 4.6 to the centered claim $\tilde{L}_{M,\tau}^\mathbb{Q}$. Since

$$q_\tau v_\tau D_{M,\tau}^\varpi + \mathcal{E}_{M,\tau}$$

is the full perturbation away from the benchmark coefficient $q_\tau v_\tau c_{M,\tau}^\varpi$, the orthogonal projection onto $\text{span}\{\Lambda_M^\star\}$ yields

$$\left| \alpha_{M,\tau}^{\mathbb{Q},\star} - q_\tau v_\tau c_{M,\tau}^\varpi \right| \leq \|q_\tau v_\tau D_{M,\tau}^\varpi + \mathcal{E}_{M,\tau}\|_{L^2(\mathbb{Q})},$$

and similarly

$$\|R_{M,\tau}^{\mathbb{Q},\star}\|_{L^2(\mathbb{Q})} \leq \|q_\tau v_\tau D_{M,\tau}^\varpi + \mathcal{E}_{M,\tau}\|_{L^2(\mathbb{Q})}.$$

The displayed bounds follow by the triangle inequality and the bound on $\mathcal{E}_{M,\tau}$. If the smallness condition holds, then

$$\|q_\tau v_\tau D_{M,\tau}^\varpi + \mathcal{E}_{M,\tau}\|_{L^2(\mathbb{Q})} = o(q_\tau v_\tau |c_{M,\tau}^\varpi|),$$

so

$$|\alpha_{M,\tau}^{\mathbb{Q},\star}| \geq q_\tau v_\tau |c_{M,\tau}^\varpi| - \|q_\tau v_\tau D_{M,\tau}^\varpi + \mathcal{E}_{M,\tau}\|_{L^2(\mathbb{Q})} \sim q_\tau v_\tau |c_{M,\tau}^\varpi|.$$

Hence

$$\frac{\|R_{M,\tau}^{\mathbb{Q},\star}\|_{L^2(\mathbb{Q})}}{|\alpha_{M,\tau}^{\mathbb{Q},\star}|} \rightarrow 0,$$

which is exactly asymptotic aggregate-factor alignment. If

$$D_{M,\tau}^\varpi = 0, \quad \mathfrak{D}_{M,\tau}^\perp = 0, \quad \tilde{J}_{M,\tau}^\mathbb{Q} = 0,$$

then $\mathcal{E}_{M,\tau} = 0$. The preceding error inequality has zero right-hand side, so the centered loading equals its aggregate-factor component. $\square$

Theorem D.32 (Common-profile Perron benchmark kernels imply proxy compatibility). *Assume the hypotheses of Definition D.27, and suppose the aggregate factor benchmark is nontrivial. Let $\bar{G}_1, \dots, \bar{G}_N$ be deterministic benchmark kernels on $\{0 \leq s \leq u \leq T\}$, and define*

$$\bar{Y}_M(u) := \sum_{j=1}^N \int_0^u \bar{G}_j(u, s) dW_s^{\mathbb{Q},j}, \quad \bar{Z}_{M,\tau} := \int_0^\tau \bar{Y}_M(u) dB_u^{\mathbb{Q}}.$$

Write the centered benchmark-profile loading as

$$\tilde{L}_{\bar{Z}_{M,\tau}}^{\mathbb{Q}} := C_{\mathcal{H}^\varpi}^{\mathbb{Q}}(\bar{Z}_{M,\tau}) - C_{\mathcal{G}_T}^{\mathbb{Q}}(C_{\mathcal{H}^\varpi}^{\mathbb{Q}}(\bar{Z}_{M,\tau})) = \bar{c}_{M,\tau}^{\mathbb{Q},*} \Lambda_M^* + \bar{R}_{M,\tau}^{\mathbb{Q},*},$$

where

$$\bar{R}_{M,\tau}^{\mathbb{Q},*} \perp_{L^2(\mathbb{Q})} \Lambda_M^*.$$

Suppose further that the canonical Perron-projected market kernels admit the common-profile decomposition

$$G_j^\varpi(u, s) = a_{M,\tau}^\varpi \bar{G}_j(u, s) + G_{j,\tau}^{\text{comp}}(u, s), \quad 0 \leq s \leq u \leq \tau, \quad 1 \leq j \leq N,$$

for some scalar $a_{M,\tau}^\varpi \in \mathbb{R}$. Define the benchmark-compatibility defect

$$\mathfrak{D}_{M,\tau}^{\text{comp}} := \sum_{j=1}^N \int_0^\tau \int_0^u |G_{j,\tau}^{\text{comp}}(u, s)|^2 ds du.$$

Then, in the factor decomposition

$$\Upsilon_{M,\tau}^\varpi = c_{M,\tau}^\varpi \Lambda_M^* + D_{M,\tau}^\varpi$$

with

$$D_{M,\tau}^\varpi \perp_{L^2(\mathbb{Q})} \Lambda_M^*,$$

one has

$$|c_{M,\tau}^\varpi - a_{M,\tau}^\varpi \bar{c}_{M,\tau}^{\mathbb{Q},*}| \leq (\mathfrak{D}_{M,\tau}^{\text{comp}})^{1/2},$$

and

$$\|D_{M,\tau}^\varpi\|_{L^2(\mathbb{Q})} \leq |a_{M,\tau}^\varpi| \|\bar{R}_{M,\tau}^{\mathbb{Q},*}\|_{L^2(\mathbb{Q})} + (\mathfrak{D}_{M,\tau}^{\text{comp}})^{1/2}.$$

If, moreover,

$$a_{M,\tau}^\varpi \bar{c}_{M,\tau}^{\mathbb{Q},*} \neq 0 \quad \text{for all sufficiently small } \tau,$$

and

$$|a_{M,\tau}^\varpi| \|\bar{R}_{M,\tau}^{\mathbb{Q},*}\|_{L^2(\mathbb{Q})} + (\mathfrak{D}_{M,\tau}^{\text{comp}})^{1/2} = o(|a_{M,\tau}^\varpi \bar{c}_{M,\tau}^{\mathbb{Q},*}|),$$

then the proxy-compatibility error is asymptotically negligible:

$$\|D_{M,\tau}^\varpi\|_{L^2(\mathbb{Q})} = o(|a_{M,\tau}^\varpi \bar{c}_{M,\tau}^{\mathbb{Q},*}|),$$

and

$$c_{M,\tau}^\varpi = a_{M,\tau}^\varpi \bar{c}_{M,\tau}^{\mathbb{Q},*} + o(|a_{M,\tau}^\varpi \bar{c}_{M,\tau}^{\mathbb{Q},*}|).$$

If in addition

$$\bar{R}_{M,\tau}^{\mathbb{Q},*} = 0, \quad \mathfrak{D}_{M,\tau}^{\text{comp}} = 0,$$

then

$$\Upsilon_{M,\tau}^\varpi = a_{M,\tau}^\varpi \bar{c}_{M,\tau}^{\mathbb{Q},*} \Lambda_M^* \quad \text{exactly.}$$

Proof. Define

$$Y_{M,\tau}^{\text{comp}}(u) := \sum_{j=1}^N \int_0^u G_{j,\tau}^{\text{comp}}(u, s) dW_s^{\mathbb{Q},j}, \quad E_{M,\tau}^{\text{comp}} := \int_0^\tau Y_{M,\tau}^{\text{comp}}(u) dB_u^{\mathbb{Q}}.$$

By linearity of stochastic integration,

$$Z_{M,\tau}^\varpi = a_{M,\tau}^\varpi \bar{Z}_{M,\tau} + E_{M,\tau}^{\text{comp}}.$$

Itô isometry and independence of the $\mathbb{Q}$ -Brownian coordinates give

$$\|E_{M,\tau}^{\text{comp}}\|_{L^2(\mathbb{Q})}^2 = \mathfrak{D}_{M,\tau}^{\text{comp}}.$$

Applying $C_{\mathcal{H}_\varpi}^{\mathbb{Q}}$, subtracting the $\mathcal{G}_T$ -projection, and using the displayed definition of $\tilde{L}_{\tilde{Z}_{M,\tau}}^{\mathbb{Q}}$, one obtains

$$\Upsilon_{M,\tau}^\varpi = a_{M,\tau}^\varpi \tilde{L}_{\tilde{Z}_{M,\tau}}^{\mathbb{Q}} + \tilde{E}_{M,\tau}^{\text{comp},\mathbb{Q}}$$

and therefore

$$\Upsilon_{M,\tau}^\varpi = a_{M,\tau}^\varpi \bar{c}_{M,\tau}^{\mathbb{Q},\star} \Lambda_M^\star + a_{M,\tau}^\varpi \bar{R}_{M,\tau}^{\mathbb{Q},\star} + \tilde{E}_{M,\tau}^{\text{comp},\mathbb{Q}},$$

where

$$\tilde{E}_{M,\tau}^{\text{comp},\mathbb{Q}} := C_{\mathcal{H}_\varpi}^{\mathbb{Q}}(E_{M,\tau}^{\text{comp}}) - C_{\mathcal{G}_T}^{\mathbb{Q}}(C_{\mathcal{H}_\varpi}^{\mathbb{Q}}(E_{M,\tau}^{\text{comp}})).$$

As in Proposition D.21,

$$\|\tilde{E}_{M,\tau}^{\text{comp},\mathbb{Q}}\|_{L^2(\mathbb{Q})} \leq \|E_{M,\tau}^{\text{comp}}\|_{L^2(\mathbb{Q})} = (\mathfrak{D}_{M,\tau}^{\text{comp}})^{1/2}.$$

Taking the $L^2(\mathbb{Q})$ -inner product against $\Lambda_M^\star$, and using $\bar{R}_{M,\tau}^{\mathbb{Q},\star} \perp \Lambda_M^\star$, gives

$$\left| c_{M,\tau}^\varpi - a_{M,\tau}^\varpi \bar{c}_{M,\tau}^{\mathbb{Q},\star} \right| \leq \|\tilde{E}_{M,\tau}^{\text{comp},\mathbb{Q}}\|_{L^2(\mathbb{Q})}.$$

Subtracting the $\Lambda_M^\star$ -component from the displayed decomposition of $\Upsilon_{M,\tau}^\varpi$ gives

$$D_{M,\tau}^\varpi = a_{M,\tau}^\varpi \bar{R}_{M,\tau}^{\mathbb{Q},\star} + \tilde{E}_{M,\tau}^{\text{comp},\mathbb{Q}} - \langle \tilde{E}_{M,\tau}^{\text{comp},\mathbb{Q}}, \Lambda_M^\star \rangle_{L^2(\mathbb{Q})} \Lambda_M^\star,$$

hence

$$\|D_{M,\tau}^\varpi\|_{L^2(\mathbb{Q})} \leq |a_{M,\tau}^\varpi| \|\bar{R}_{M,\tau}^{\mathbb{Q},\star}\|_{L^2(\mathbb{Q})} + \|\tilde{E}_{M,\tau}^{\text{comp},\mathbb{Q}}\|_{L^2(\mathbb{Q})}.$$

Dividing the displayed bound by $|a_{M,\tau}^\varpi|$ gives the stated asymptotic estimates under the two smallness assumptions. If

$$\bar{R}_{M,\tau}^{\mathbb{Q},\star} = 0, \quad \mathfrak{D}_{M,\tau}^{\text{comp}} = 0,$$

then $E_{M,\tau}^{\text{comp}} = 0$ in $L^2(\mathbb{Q})$, so the exact identity follows. $\square$

Proposition D.33 (Common-profile proxy compatibility with a nondegenerate scalar amplitude forces benchmark-loading coefficient control). *Assume the hypotheses of Theorem D.32. Suppose there exist constants $\underline{a}_M, \bar{a}_M > 0$ such that*

$$\underline{a}_M \leq |a_{M,\tau}^\varpi| \leq \bar{a}_M \quad \text{for all sufficiently small } \tau,$$

and that

$$c_{M,\tau}^\varpi \neq 0 \quad \text{for all sufficiently small } \tau.$$

If

$$\|D_{M,\tau}^\varpi\|_{L^2(\mathbb{Q})} + (\mathfrak{D}_{M,\tau}^{\text{comp}})^{1/2} = o(|c_{M,\tau}^\varpi|),$$

then

$$\left| \bar{c}_{M,\tau}^{\mathbb{Q},\star} \right| \asymp |c_{M,\tau}^\varpi|,$$

and

$$\|\bar{R}_{M,\tau}^{\mathbb{Q},\star}\|_{L^2(\mathbb{Q})} = o(|c_{M,\tau}^\varpi|) = o(|\bar{c}_{M,\tau}^{\mathbb{Q},\star}|).$$

If, in addition,

$$\mathfrak{d}_{M,\tau}^{\varpi,\text{bench}} = o(|c_{M,\tau}^\varpi|),$$

then Proposition D.35 applies with

$$a_\tau := |c_{M,\tau}^\varpi|,$$

so

$$d_{M,\tau}^{\text{bench}} \neq 0 \quad \text{for all sufficiently small } \tau,$$

and

$$|d_{M,\tau}^{\text{bench}}| \sim |\bar{c}_{M,\tau}^{\mathbb{Q},\star}| \sim |c_{M,\tau}^\varpi|.$$

If, moreover,

$$D_{M,\tau}^\varpi = 0, \quad \mathfrak{D}_{M,\tau}^{\text{comp}} = 0,$$

then

$$\bar{R}_{M,\tau}^{\mathbb{Q},\star} = 0, \quad c_{M,\tau}^\varpi = a_{M,\tau}^\varpi \bar{c}_{M,\tau}^{\mathbb{Q},\star} \quad \text{for all sufficiently small } \tau.$$

Proof. Set

$$e_{M,\tau}^{\text{comp}} := (\mathfrak{D}_{M,\tau}^{\text{comp}})^{1/2}.$$

Theorem D.32 gives

$$\left| c_{M,\tau}^{\varpi} - a_{M,\tau}^{\varpi} \bar{c}_{M,\tau}^{\mathbb{Q},\star} \right| \leq e_{M,\tau}^{\text{comp}}.$$

Hence

$$|a_{M,\tau}^{\varpi}| \left| \bar{c}_{M,\tau}^{\mathbb{Q},\star} \right| \geq |c_{M,\tau}^{\varpi}| - e_{M,\tau}^{\text{comp}},$$

and

$$|a_{M,\tau}^{\varpi}| \left| \bar{c}_{M,\tau}^{\mathbb{Q},\star} \right| \leq |c_{M,\tau}^{\varpi}| + e_{M,\tau}^{\text{comp}}.$$

Because

$$e_{M,\tau}^{\text{comp}} = o\left(|c_{M,\tau}^{\varpi}|\right),$$

and $|a_{M,\tau}^{\varpi}| \asymp 1$, it follows that

$$\left| \bar{c}_{M,\tau}^{\mathbb{Q},\star} \right| \asymp |c_{M,\tau}^{\varpi}|.$$

In the proof of Theorem D.32, one has

$$D_{M,\tau}^{\varpi} = a_{M,\tau}^{\varpi} \bar{R}_{M,\tau}^{\mathbb{Q},\star} + \tilde{E}_{M,\tau}^{\text{comp},\mathbb{Q}} - \langle \tilde{E}_{M,\tau}^{\text{comp},\mathbb{Q}}, \Lambda_M^{\star} \rangle_{L^2(\mathbb{Q})} \Lambda_M^{\star},$$

with

$$\|\tilde{E}_{M,\tau}^{\text{comp},\mathbb{Q}}\|_{L^2(\mathbb{Q})} \leq e_{M,\tau}^{\text{comp}}.$$

Therefore

$$|a_{M,\tau}^{\varpi}| \|\bar{R}_{M,\tau}^{\mathbb{Q},\star}\|_{L^2(\mathbb{Q})} \leq \|D_{M,\tau}^{\varpi}\|_{L^2(\mathbb{Q})} + e_{M,\tau}^{\text{comp}} = o\left(|c_{M,\tau}^{\varpi}|\right),$$

so

$$\|\bar{R}_{M,\tau}^{\mathbb{Q},\star}\|_{L^2(\mathbb{Q})} = o\left(|c_{M,\tau}^{\varpi}|\right) = o\left(\left|\bar{c}_{M,\tau}^{\mathbb{Q},\star}\right|\right).$$

If, in addition,

$$\mathfrak{d}_{M,\tau}^{\varpi,\text{bench}} = o\left(|c_{M,\tau}^{\varpi}|\right),$$

then

$$\|\bar{R}_{M,\tau}^{\mathbb{Q},\star}\|_{L^2(\mathbb{Q})} = o(a_{\tau}), \quad \mathfrak{d}_{M,\tau}^{\varpi,\text{bench}} = o(a_{\tau}),$$

with $a_{\tau} = |c_{M,\tau}^{\varpi}|$, while the asymptotic equivalence

$$\left| \bar{c}_{M,\tau}^{\mathbb{Q},\star} \right| \asymp a_{\tau}$$

provides the required lower bound on $|\bar{c}_{M,\tau}^{\mathbb{Q},\star}|$. Proposition D.35 therefore gives the stated nonvanishing and asymptotic equivalences for $d_{M,\tau}^{\text{bench}}$. If, moreover,

$$D_{M,\tau}^{\varpi} = 0, \quad \mathfrak{D}_{M,\tau}^{\text{comp}} = 0,$$

then

$$\tilde{E}_{M,\tau}^{\text{comp},\mathbb{Q}} = 0,$$

so the displayed decomposition of $D_{M,\tau}^{\varpi}$ reduces to

$$0 = a_{M,\tau}^{\varpi} \bar{R}_{M,\tau}^{\mathbb{Q},\star}.$$

Since $a_{M,\tau}^{\varpi} \neq 0$, this gives

$$\bar{R}_{M,\tau}^{\mathbb{Q},\star} = 0.$$

The exact identity

$$c_{M,\tau}^{\varpi} = a_{M,\tau}^{\varpi} \bar{c}_{M,\tau}^{\mathbb{Q},\star}$$

is the exact branch of Theorem D.32. □

Corollary D.34 (Asymptotic benchmark-profile compatibility with the aggregate factor). *Under the notation of Theorem D.32, define the centered benchmark-profile direction by*

$$\Lambda_{M,\tau}^{\text{bench}} := \begin{cases} \tilde{L}_{\tilde{Z}_{M,\tau}}^{\mathbb{Q}} / \|\tilde{L}_{\tilde{Z}_{M,\tau}}^{\mathbb{Q}}\|_{L^2(\mathbb{Q})}, & \text{if } \tilde{L}_{\tilde{Z}_{M,\tau}}^{\mathbb{Q}} \neq 0, \\ 0, & \text{if } \tilde{L}_{\tilde{Z}_{M,\tau}}^{\mathbb{Q}} = 0. \end{cases}$$

If

$$\bar{c}_{M,\tau}^{\mathbb{Q},\star} \neq 0 \quad \text{for all sufficiently small } \tau, \quad \|\bar{R}_{M,\tau}^{\mathbb{Q},\star}\|_{L^2(\mathbb{Q})} = o(|\bar{c}_{M,\tau}^{\mathbb{Q},\star}|),$$

then

$$\left\| \Lambda_{M,\tau}^{\text{bench}} - \frac{\bar{c}_{M,\tau}^{\mathbb{Q},\star}}{|\bar{c}_{M,\tau}^{\mathbb{Q},\star}|} \Lambda_M^{\star} \right\|_{L^2(\mathbb{Q})} \rightarrow 0.$$

If, in addition, $\bar{R}_{M,\tau}^{\mathbb{Q},\star} = 0$, then

$$\Lambda_{M,\tau}^{\text{bench}} = \frac{\bar{c}_{M,\tau}^{\mathbb{Q},\star}}{|\bar{c}_{M,\tau}^{\mathbb{Q},\star}|} \Lambda_M^{\star} \quad \text{exactly}$$

whenever $\bar{c}_{M,\tau}^{\mathbb{Q},\star} \neq 0$.

Proof. The factor decomposition of $\tilde{L}_{\tilde{Z}_{M,\tau}}^{\mathbb{Q}}$ is

$$\tilde{L}_{\tilde{Z}_{M,\tau}}^{\mathbb{Q}} = \bar{c}_{M,\tau}^{\mathbb{Q},\star} \Lambda_M^{\star} + \bar{R}_{M,\tau}^{\mathbb{Q},\star}.$$

If $\|\bar{R}_{M,\tau}^{\mathbb{Q},\star}\|_{L^2(\mathbb{Q})} = o(|\bar{c}_{M,\tau}^{\mathbb{Q},\star}|)$, then

$$\|\tilde{L}_{\tilde{Z}_{M,\tau}}^{\mathbb{Q}}\|_{L^2(\mathbb{Q})} \sim |\bar{c}_{M,\tau}^{\mathbb{Q},\star}|,$$

and the normalized-direction claim follows exactly as in the proof of Theorem D.14. If $\bar{R}_{M,\tau}^{\mathbb{Q},\star} = 0$, the orthogonal residual in the benchmark-loading decomposition vanishes, so the exact identity follows from the displayed projection formula. $\square$

Proposition D.35 (Benchmark-loading coefficient lower bounds and residual smallness force benchmark/generator separation). *Under the notation of Theorem D.32, write*

$$\tilde{L}_{\tilde{Z}_{M,\tau}}^{\mathbb{Q}} = \bar{c}_{M,\tau}^{\mathbb{Q},\star} \Lambda_M^{\star} + \bar{R}_{M,\tau}^{\mathbb{Q},\star}, \quad \bar{R}_{M,\tau}^{\mathbb{Q},\star} \perp_{L^2(\mathbb{Q})} \Lambda_M^{\star}.$$

Suppose there exists a deterministic positive scale a_τ and a constant $c_M > 0$ such that

$$|\bar{c}_{M,\tau}^{\mathbb{Q},\star}| \geq c_M a_\tau \quad \text{for all sufficiently small } \tau,$$

and

$$\|\bar{R}_{M,\tau}^{\mathbb{Q},\star}\|_{L^2(\mathbb{Q})} = o(a_\tau).$$

Then

$$\left\| \tilde{L}_{\tilde{Z}_{M,\tau}}^{\mathbb{Q}} \right\|_{L^2(\mathbb{Q})} \geq c_M a_\tau \quad \text{for all sufficiently small } \tau,$$

and

$$\left\| \tilde{L}_{\tilde{Z}_{M,\tau}}^{\mathbb{Q}} \right\|_{L^2(\mathbb{Q})} \sim |\bar{c}_{M,\tau}^{\mathbb{Q},\star}|.$$

If, in addition,

$$\mathfrak{d}_{M,\tau}^{\varpi, \text{bench}} = o(a_\tau),$$

then the hypotheses of Proposition D.38 hold, and therefore

$$d_{M,\tau}^{\text{bench}} \neq 0 \quad \text{for all sufficiently small } \tau,$$

with

$$|d_{M,\tau}^{\text{bench}}| \sim \left\| \tilde{L}_{\tilde{Z}_{M,\tau}}^{\mathbb{Q}} \right\|_{L^2(\mathbb{Q})} \sim |\bar{c}_{M,\tau}^{\mathbb{Q},\star}|.$$

If, moreover,

$$\bar{R}_{M,\tau}^{\mathbb{Q},\star} = 0, \quad \mathfrak{d}_{M,\tau}^{\varpi,\text{bench}} = 0, \quad \bar{c}_{M,\tau}^{\mathbb{Q},\star} \neq 0 \quad \text{for all sufficiently small } \tau,$$

then

$$|d_{M,\tau}^{\text{bench}}| = |\bar{c}_{M,\tau}^{\mathbb{Q},\star}| \quad \text{for all sufficiently small } \tau.$$

Proof. By orthogonality of $\bar{R}_{M,\tau}^{\mathbb{Q},\star}$ and $\Lambda_M^\star$,

$$\left\| \tilde{L}_{\bar{Z}_{M,\tau}}^{\mathbb{Q}} \right\|_{L^2(\mathbb{Q})}^2 = |\bar{c}_{M,\tau}^{\mathbb{Q},\star}|^2 + \|\bar{R}_{M,\tau}^{\mathbb{Q},\star}\|_{L^2(\mathbb{Q})}^2.$$

Therefore

$$\left\| \tilde{L}_{\bar{Z}_{M,\tau}}^{\mathbb{Q}} \right\|_{L^2(\mathbb{Q})} \geq |\bar{c}_{M,\tau}^{\mathbb{Q},\star}| \geq c_M a_\tau,$$

which proves the displayed lower bound and, together with $\|\bar{R}_{M,\tau}^{\mathbb{Q},\star}\|_{L^2(\mathbb{Q})} = o(a_\tau)$, also gives

$$\left\| \tilde{L}_{\bar{Z}_{M,\tau}}^{\mathbb{Q}} \right\|_{L^2(\mathbb{Q})} \sim |\bar{c}_{M,\tau}^{\mathbb{Q},\star}|.$$

If, in addition,

$$\mathfrak{d}_{M,\tau}^{\varpi,\text{bench}} = o(a_\tau),$$

then

$$\mathfrak{d}_{M,\tau}^{\varpi,\text{bench}} = o\left(\left\| \tilde{L}_{\bar{Z}_{M,\tau}}^{\mathbb{Q}} \right\|_{L^2(\mathbb{Q})}\right),$$

so Proposition D.38 applies and gives the displayed nonvanishing and asymptotic equivalences for $d_{M,\tau}^{\text{bench}}$. If, moreover,

$$\bar{R}_{M,\tau}^{\mathbb{Q},\star} = 0, \quad \mathfrak{d}_{M,\tau}^{\varpi,\text{bench}} = 0, \quad \bar{c}_{M,\tau}^{\mathbb{Q},\star} \neq 0,$$

then

$$\left\| \tilde{L}_{\bar{Z}_{M,\tau}}^{\mathbb{Q}} \right\|_{L^2(\mathbb{Q})} = |\bar{c}_{M,\tau}^{\mathbb{Q},\star}|$$

and Proposition D.38 gives

$$|d_{M,\tau}^{\text{bench}}| = \left\| \tilde{L}_{\bar{Z}_{M,\tau}}^{\mathbb{Q}} \right\|_{L^2(\mathbb{Q})},$$

which proves the exact identity. $\square$

Definition D.36 (Benchmark-generated centered ϖ -profile generator). Under the notation of Theorem D.32, define the centered ϖ -shadow of the centered benchmark-profile loading by

$$\Theta_{M,\tau}^{\varpi,\text{bench}} := C_{\mathcal{L}_T}^{\mathbb{Q}}(\tilde{L}_{\bar{Z}_{M,\tau}}^{\mathbb{Q}}) - C_{\mathcal{G}_T}^{\mathbb{Q}}\left(C_{\mathcal{L}_T}^{\mathbb{Q}}(\tilde{L}_{\bar{Z}_{M,\tau}}^{\mathbb{Q}})\right) \in \mathcal{M}_{\mathbb{Q}}^\circ.$$

Define the benchmark-generated centered ϖ -profile generator by

$$\Gamma_{M,\tau}^{\varpi} := \begin{cases} \Theta_{M,\tau}^{\varpi,\text{bench}} / \|\Theta_{M,\tau}^{\varpi,\text{bench}}\|_{L^2(\mathbb{Q})}, & \text{if } \Theta_{M,\tau}^{\varpi,\text{bench}} \neq 0, \\ 0, & \text{if } \Theta_{M,\tau}^{\varpi,\text{bench}} = 0. \end{cases}$$

and the associated canonical benchmark decomposition

$$\tilde{L}_{\bar{Z}_{M,\tau}}^{\mathbb{Q}} = d_{M,\tau}^{\text{bench}} \Gamma_{M,\tau}^{\varpi} + E_{M,\tau}^{\text{bench},\varpi},$$

where

$$d_{M,\tau}^{\text{bench}} := \left\langle \tilde{L}_{\bar{Z}_{M,\tau}}^{\mathbb{Q}}, \Gamma_{M,\tau}^{\varpi} \right\rangle_{L^2(\mathbb{Q})}, \quad E_{M,\tau}^{\text{bench},\varpi} := \tilde{L}_{\bar{Z}_{M,\tau}}^{\mathbb{Q}} - d_{M,\tau}^{\text{bench}} \Gamma_{M,\tau}^{\varpi}.$$

Then

$$E_{M,\tau}^{\text{bench},\varpi} \in \mathcal{M}_{\mathbb{Q}}^\circ, \quad E_{M,\tau}^{\text{bench},\varpi} \perp_{L^2(\mathbb{Q})} \Gamma_{M,\tau}^{\varpi}.$$

Proposition D.37 (Canonical aggregate-factor benchmark decomposition relative to the benchmark-generated ϖ -profile generator). *Under the notation of Definition D.36, define the benchmark ϖ -generation defect by*

$$\mathfrak{d}_{M,\tau}^{\varpi,\text{bench}} := \left\| \tilde{L}_{\tilde{Z}_{M,\tau}}^{\mathbb{Q}} - \Theta_{M,\tau}^{\varpi,\text{bench}} \right\|_{L^2(\mathbb{Q})},$$

the canonical aggregate-factor coefficient relative to the benchmark-generated centered ϖ -profile generator by

$$g_{M,\tau}^{\varpi,*} := \left\langle \Lambda_M^*, \Gamma_{M,\tau}^{\varpi} \right\rangle_{L^2(\mathbb{Q})},$$

and the canonical aggregate-factor residual by

$$H_{M,\tau}^{\varpi,*} := \Lambda_M^* - g_{M,\tau}^{\varpi,*} \Gamma_{M,\tau}^{\varpi}.$$

Then

$$H_{M,\tau}^{\varpi,*} \in \mathcal{M}_{\mathbb{Q}}^{\circ}, \quad H_{M,\tau}^{\varpi,*} \perp_{L^2(\mathbb{Q})} \Gamma_{M,\tau}^{\varpi}.$$

Define the aggregate-factor ϖ -generation defect by

$$\mathfrak{d}_{M,\tau}^{\varpi,*} := \|H_{M,\tau}^{\varpi,*}\|_{L^2(\mathbb{Q})}.$$

Then

$$\|E_{M,\tau}^{\text{bench},\varpi}\|_{L^2(\mathbb{Q})} \leq \mathfrak{d}_{M,\tau}^{\varpi,\text{bench}}.$$

Moreover, under Proposition 10.11,

$$\tilde{L}_{Q_M}^{\mathbb{Q}} = \alpha_{Q_M}^{\mathbb{Q},*} \Lambda_M^*,$$

so the canonical market-variance coefficient and residual relative to $\Gamma_{M,\tau}^{\varpi}$ are

$$d_{M,\tau}^{\text{var}} := \left\langle \tilde{L}_{Q_M}^{\mathbb{Q}}, \Gamma_{M,\tau}^{\varpi} \right\rangle_{L^2(\mathbb{Q})} = \alpha_{Q_M}^{\mathbb{Q},*} g_{M,\tau}^{\varpi,*},$$

$$E_{M,\tau}^{\text{var},\varpi} := \tilde{L}_{Q_M}^{\mathbb{Q}} - d_{M,\tau}^{\text{var}} \Gamma_{M,\tau}^{\varpi} = \alpha_{Q_M}^{\mathbb{Q},*} H_{M,\tau}^{\varpi,*},$$

and therefore

$$\|E_{M,\tau}^{\text{var},\varpi}\|_{L^2(\mathbb{Q})} = \alpha_{Q_M}^{\mathbb{Q},*} \mathfrak{d}_{M,\tau}^{\varpi,*}.$$

If $g_{M,\tau}^{\varpi,} \neq 0$ and $\alpha_{Q_M}^{\mathbb{Q},*} \neq 0$, then*

$$\frac{\|E_{M,\tau}^{\text{var},\varpi}\|_{L^2(\mathbb{Q})}}{|d_{M,\tau}^{\text{var}}|} = \frac{\mathfrak{d}_{M,\tau}^{\varpi,*}}{|g_{M,\tau}^{\varpi,*}|}.$$

Proof. Because $\Theta_{M,\tau}^{\varpi,\text{bench}} \in \text{span}\{\Gamma_{M,\tau}^{\varpi}\}$ by Definition D.36, the orthogonal projection property gives

$$\|E_{M,\tau}^{\text{bench},\varpi}\|_{L^2(\mathbb{Q})} = \text{dist}\left(\tilde{L}_{\tilde{Z}_{M,\tau}}^{\mathbb{Q}}, \text{span}\{\Gamma_{M,\tau}^{\varpi}\}\right) \leq \mathfrak{d}_{M,\tau}^{\varpi,\text{bench}}.$$

The definitions of $g_{M,\tau}^{\varpi,*}$ and $H_{M,\tau}^{\varpi,*}$ make $H_{M,\tau}^{\varpi,*}$ the orthogonal residual of Λ_M^* relative to $\text{span}\{\Gamma_{M,\tau}^{\varpi}\}$, which yields the displayed orthogonality. The formulas for $d_{M,\tau}^{\text{var}}$ and $E_{M,\tau}^{\text{var},\varpi}$ follow by substituting $\tilde{L}_{Q_M}^{\mathbb{Q}} = \alpha_{Q_M}^{\mathbb{Q},*} \Lambda_M^*$ from Proposition 10.11 into the definitions of the coefficient and residual relative to $\Gamma_{M,\tau}^{\varpi}$. Orthogonality gives the Pythagorean norm identity, and the ratio identity follows by substituting $d_{M,\tau}^{\text{var}} = \alpha_{Q_M}^{\mathbb{Q},*} g_{M,\tau}^{\varpi,*}$ when both denominators are nonzero. $\square$

Proposition D.38 (Benchmark-profile / generator norm separation forces a nontrivial benchmark coefficient). *Assume Definition D.36 and Proposition D.37. If*

$$\mathfrak{d}_{M,\tau}^{\varpi,\text{bench}} < \left\| \tilde{L}_{\tilde{Z}_{M,\tau}}^{\mathbb{Q}} \right\|_{L^2(\mathbb{Q})} \quad \text{for all sufficiently small } \tau,$$

then

$$d_{M,\tau}^{\text{bench}} \neq 0 \quad \text{for all sufficiently small } \tau.$$

If, moreover,

$$\mathfrak{d}_{M,\tau}^{\varpi,\text{bench}} = o\left(\left\|\tilde{L}_{\tilde{Z}_{M,\tau}}^{\mathbb{Q}}\right\|_{L^2(\mathbb{Q})}\right),$$

then

$$\left|d_{M,\tau}^{\text{bench}}\right| \sim \left\|\tilde{L}_{\tilde{Z}_{M,\tau}}^{\mathbb{Q}}\right\|_{L^2(\mathbb{Q})}.$$

Proof. If $\Gamma_{M,\tau}^{\varpi} = 0$, then Definition D.36 gives

$$d_{M,\tau}^{\text{bench}} = 0, \quad E_{M,\tau}^{\text{bench},\varpi} = \tilde{L}_{\tilde{Z}_{M,\tau}}^{\mathbb{Q}}.$$

Proposition D.37 would then imply

$$\left\|\tilde{L}_{\tilde{Z}_{M,\tau}}^{\mathbb{Q}}\right\|_{L^2(\mathbb{Q})} = \|E_{M,\tau}^{\text{bench},\varpi}\|_{L^2(\mathbb{Q})} \leq \mathfrak{d}_{M,\tau}^{\varpi,\text{bench}},$$

contradicting the displayed strict inequality. Therefore

$$\Gamma_{M,\tau}^{\varpi} \neq 0 \quad \text{for all sufficiently small } \tau.$$

Since $\Gamma_{M,\tau}^{\varpi}$ is then a unit vector, the orthogonal decomposition in Definition D.36 yields

$$\left\|\tilde{L}_{\tilde{Z}_{M,\tau}}^{\mathbb{Q}}\right\|_{L^2(\mathbb{Q})}^2 = \left|d_{M,\tau}^{\text{bench}}\right|^2 + \|E_{M,\tau}^{\text{bench},\varpi}\|_{L^2(\mathbb{Q})}^2.$$

Together with Proposition D.37, this gives

$$\left|d_{M,\tau}^{\text{bench}}\right|^2 \geq \left\|\tilde{L}_{\tilde{Z}_{M,\tau}}^{\mathbb{Q}}\right\|_{L^2(\mathbb{Q})}^2 - (\mathfrak{d}_{M,\tau}^{\varpi,\text{bench}})^2.$$

The displayed strict inequality makes the lower bound positive, which gives the first claim. If, moreover,

$$\mathfrak{d}_{M,\tau}^{\varpi,\text{bench}} = o\left(\left\|\tilde{L}_{\tilde{Z}_{M,\tau}}^{\mathbb{Q}}\right\|_{L^2(\mathbb{Q})}\right),$$

then

$$\frac{\|E_{M,\tau}^{\text{bench},\varpi}\|_{L^2(\mathbb{Q})}}{\left\|\tilde{L}_{\tilde{Z}_{M,\tau}}^{\mathbb{Q}}\right\|_{L^2(\mathbb{Q})}} \leq \frac{\mathfrak{d}_{M,\tau}^{\varpi,\text{bench}}}{\left\|\tilde{L}_{\tilde{Z}_{M,\tau}}^{\mathbb{Q}}\right\|_{L^2(\mathbb{Q})}} \rightarrow 0,$$

and the orthogonal norm identity yields the asymptotic equivalence

$$\left|d_{M,\tau}^{\text{bench}}\right| \sim \left\|\tilde{L}_{\tilde{Z}_{M,\tau}}^{\mathbb{Q}}\right\|_{L^2(\mathbb{Q})}.$$

□

Proposition D.39 (Q_M -side ϖ -shadow control of the aggregate-factor defect). *Assume Proposition 10.11 and Definition D.36. If*

$$\alpha_{Q_M}^{\mathbb{Q},\star} \neq 0,$$

define the centered Q_M -side ϖ -shadow by

$$\Theta_{M,\tau}^{\varpi,Q} := C_{\tilde{\mathcal{L}}_T}^{\mathbb{Q}}(\tilde{L}_{Q_M}^{\mathbb{Q}}) - C_{\mathcal{G}_T}^{\mathbb{Q}}\left(C_{\tilde{\mathcal{L}}_T}^{\mathbb{Q}}(\tilde{L}_{Q_M}^{\mathbb{Q}})\right) \in \mathcal{M}_{\mathbb{Q}}^{\circ}$$

and the normalized Q_M -side ϖ -shadow defect by

$$\mathfrak{s}_{M,\tau}^{\varpi,Q} := \frac{\|\tilde{L}_{Q_M}^{\mathbb{Q}} - \Theta_{M,\tau}^{\varpi,Q}\|_{L^2(\mathbb{Q})} + \text{dist}\left(\Theta_{M,\tau}^{\varpi,Q}, \text{span}\{\Gamma_{M,\tau}^{\varpi}\}\right)}{\alpha_{Q_M}^{\mathbb{Q},\star}}.$$

Then

$$\mathfrak{d}_{M,\tau}^{\varpi,\star} \leq \mathfrak{s}_{M,\tau}^{\varpi,Q}.$$

Consequently, if

$$\|\tilde{L}_{Q_M}^{\mathbb{Q}} - \Theta_{M,\tau}^{\varpi,Q}\|_{L^2(\mathbb{Q})} = o\left(\alpha_{Q_M}^{\mathbb{Q},\star}\right), \quad \text{dist}\left(\Theta_{M,\tau}^{\varpi,Q}, \text{span}\{\Gamma_{M,\tau}^{\varpi}\}\right) = o\left(\alpha_{Q_M}^{\mathbb{Q},\star}\right),$$

then

$$\mathfrak{d}_{M,\tau}^{\varpi,\star} \rightarrow 0.$$

If both displayed terms vanish identically, then

$$\mathfrak{d}_{M,\tau}^{\varpi,\star} = 0.$$

Proof. Because

$$\tilde{L}_{Q_M}^{\mathbb{Q}} = \alpha_{Q_M}^{\mathbb{Q},\star} \Lambda_M^{\star}$$

by Proposition 10.11, one has

$$\alpha_{Q_M}^{\mathbb{Q},\star} \mathfrak{d}_{M,\tau}^{\varpi,\star} = \text{dist}\left(\tilde{L}_{Q_M}^{\mathbb{Q}}, \text{span}\{\Gamma_{M,\tau}^{\varpi}\}\right).$$

Applying the triangle inequality through $\Theta_{M,\tau}^{\varpi,Q}$ yields

$$\text{dist}\left(\tilde{L}_{Q_M}^{\mathbb{Q}}, \text{span}\{\Gamma_{M,\tau}^{\varpi}\}\right) \leq \|\tilde{L}_{Q_M}^{\mathbb{Q}} - \Theta_{M,\tau}^{\varpi,Q}\|_{L^2(\mathbb{Q})} + \text{dist}\left(\Theta_{M,\tau}^{\varpi,Q}, \text{span}\{\Gamma_{M,\tau}^{\varpi}\}\right),$$

which gives the first claim after dividing by $\alpha_{Q_M}^{\mathbb{Q},\star}$. If the two normalized terms on the right vanish, the quotient defect also vanishes. If both terms are zero, the distance to $\text{span}\{\Gamma_{M,\tau}^{\varpi}\}$ is zero, which is the exact statement. $\square$

Proposition D.40 (Explicit decomposition of the Q_M -side ϖ -shadow defect). *Assume Assumption 10.6, Proposition 10.9, and Definition D.36. Write*

$$\zeta(t) := (w^\top v)\varpi(t), \quad R_{\eta,M}^{\mathbb{Q}} := \int_0^T \mathbb{E}_{\mathbb{Q}}[\eta^{\mathbb{Q}}(t)^2 \mid \mathcal{H}_{\varpi}] dt.$$

Define the pure ϖ -benchmark

$$B_M^{\varpi,Q} := \int_0^T \zeta(t)^2 dt + C_{\mathcal{L}_T}^{\mathbb{Q}}(R_{\eta,M}^{\mathbb{Q}}) \in L^2(\mathbb{Q}, \mathcal{L}_T),$$

its centered image

$$\tilde{B}_M^{\varpi,Q} := B_M^{\varpi,Q} - C_{\mathcal{G}_T}^{\mathbb{Q}}(B_M^{\varpi,Q}) \in \mathcal{M}_{\mathbb{Q}}^{\circ},$$

the mixed ϖ -visibility remainder

$$S_M^{\varpi,Q} := 2 \int_0^T M^{\text{vis},\mathbb{Q}}(t) \zeta(t) dt + R_{\eta,M}^{\mathbb{Q}} - C_{\mathcal{L}_T}^{\mathbb{Q}}(R_{\eta,M}^{\mathbb{Q}}),$$

and the visible-shadow defect

$$\Delta_M^{\varpi,\text{vis}} := \left\| C_{\mathcal{L}_T}^{\mathbb{Q}}\left(C_{\mathcal{G}_T}^{\mathbb{Q}}(B_M^{\varpi,Q})\right) - C_{\mathcal{G}_T}^{\mathbb{Q}}\left(C_{\mathcal{L}_T}^{\mathbb{Q}}\left(C_{\mathcal{G}_T}^{\mathbb{Q}}(B_M^{\varpi,Q})\right)\right) \right\|_{L^2(\mathbb{Q})}.$$

Then

$$L_{Q_M}^{\mathbb{Q}} = \int_0^T (M^{\text{vis},\mathbb{Q}}(t))^2 dt + B_M^{\varpi,Q} + S_M^{\varpi,Q},$$

and therefore

$$\tilde{L}_{Q_M}^{\mathbb{Q}} = \tilde{B}_M^{\varpi,Q} + \tilde{S}_M^{\varpi,Q}, \quad \tilde{S}_M^{\varpi,Q} := S_M^{\varpi,Q} - C_{\mathcal{G}_T}^{\mathbb{Q}}(S_M^{\varpi,Q}).$$

If $\alpha_{Q_M}^{\mathbb{Q},\star} \neq 0$, then

$$\mathfrak{s}_{M,\tau}^{\varpi,Q} \leq \frac{5 \|\tilde{S}_M^{\varpi,Q}\|_{L^2(\mathbb{Q})} + 2 \Delta_M^{\varpi,\text{vis}} + \text{dist}\left(\tilde{B}_M^{\varpi,Q}, \text{span}\{\Gamma_{M,\tau}^{\varpi}\}\right)}{\alpha_{Q_M}^{\mathbb{Q},\star}}.$$

Consequently, if

$$\|\tilde{S}_M^{\varpi,Q}\|_{L^2(\mathbb{Q})} = o\left(\alpha_{Q_M}^{\mathbb{Q},\star}\right), \quad \Delta_M^{\varpi,\text{vis}} = o\left(\alpha_{Q_M}^{\mathbb{Q},\star}\right),$$

and

$$\text{dist}\left(\tilde{B}_M^{\varpi,Q}, \text{span}\{\Gamma_{M,\tau}^{\varpi}\}\right) = o\left(\alpha_{Q_M}^{\mathbb{Q},\star}\right),$$

then

$$\mathfrak{s}_{M,\tau}^{\varpi,Q} = o(1).$$

If all three displayed defects vanish identically, then

$$\mathfrak{s}_{M,\tau}^{\varpi,Q} = 0.$$

Proof. By Proposition 10.9,

$$L_{Q_M}^{\mathbb{Q}} = \int_0^T (M^{\text{vis},\mathbb{Q}}(t) + \zeta(t))^2 dt + R_{\eta,M}^{\mathbb{Q}}.$$

Expanding the square and regrouping the $\mathcal{L}_T$ -measurable part gives the first displayed decomposition of $L_{Q_M}^{\mathbb{Q}}$. Since $\int_0^T (M^{\text{vis},\mathbb{Q}}(t))^2 dt \in L^2(\mathbb{Q}, \mathcal{G}_T)$, subtracting $C_{\mathcal{G}_T}^{\mathbb{Q}}(L_{Q_M}^{\mathbb{Q}})$ yields

$$\tilde{L}_{Q_M}^{\mathbb{Q}} = \tilde{B}_M^{\varpi,Q} + \tilde{S}_M^{\varpi,Q}.$$

Write

$$\Pi_{\varpi}^{\mathbb{Q}}(Y) := C_{\mathcal{L}_T}^{\mathbb{Q}}(Y) - C_{\mathcal{G}_T}^{\mathbb{Q}}(C_{\mathcal{L}_T}^{\mathbb{Q}}(Y)).$$

Then $\Theta_{M,\tau}^{\varpi,Q} = \Pi_{\varpi}^{\mathbb{Q}}(\tilde{L}_{Q_M}^{\mathbb{Q}})$. Because $B_M^{\varpi,Q} \in L^2(\mathbb{Q}, \mathcal{L}_T)$,

$$\Pi_{\varpi}^{\mathbb{Q}}(B_M^{\varpi,Q}) = \tilde{B}_M^{\varpi,Q},$$

and hence

$$\Pi_{\varpi}^{\mathbb{Q}}(\tilde{B}_M^{\varpi,Q}) = \tilde{B}_M^{\varpi,Q} - \Pi_{\varpi}^{\mathbb{Q}}(C_{\mathcal{G}_T}^{\mathbb{Q}}(B_M^{\varpi,Q})).$$

Therefore

$$\Theta_{M,\tau}^{\varpi,Q} = \tilde{B}_M^{\varpi,Q} - \Pi_{\varpi}^{\mathbb{Q}}(C_{\mathcal{G}_T}^{\mathbb{Q}}(B_M^{\varpi,Q})) + \Pi_{\varpi}^{\mathbb{Q}}(\tilde{S}_M^{\varpi,Q}).$$

Since conditional expectation is an L^2 -contraction,

$$\|\Pi_{\varpi}^{\mathbb{Q}}(\tilde{S}_M^{\varpi,Q})\|_{L^2(\mathbb{Q})} \leq 2\|\tilde{S}_M^{\varpi,Q}\|_{L^2(\mathbb{Q})}$$

Using the displayed decomposition of $\Theta_{M,\tau}^{\varpi,Q}$, one gets

$$\|\tilde{L}_{Q_M}^{\mathbb{Q}} - \Theta_{M,\tau}^{\varpi,Q}\|_{L^2(\mathbb{Q})} \leq \|\tilde{S}_M^{\varpi,Q}\|_{L^2(\mathbb{Q})} + \Delta_M^{\varpi,\text{vis}} + \|\Pi_{\varpi}^{\mathbb{Q}}(\tilde{S}_M^{\varpi,Q})\|_{L^2(\mathbb{Q})} \leq 3\|\tilde{S}_M^{\varpi,Q}\|_{L^2(\mathbb{Q})} + \Delta_M^{\varpi,\text{vis}}.$$

Likewise,

$$\begin{aligned} \text{dist}(\Theta_{M,\tau}^{\varpi,Q}, \text{span}\{\Gamma_{M,\tau}^{\varpi}\}) &\leq \text{dist}(\tilde{B}_M^{\varpi,Q}, \text{span}\{\Gamma_{M,\tau}^{\varpi}\}) + \Delta_M^{\varpi,\text{vis}} + \|\Pi_{\varpi}^{\mathbb{Q}}(\tilde{S}_M^{\varpi,Q})\|_{L^2(\mathbb{Q})} \\ &\leq \text{dist}(\tilde{B}_M^{\varpi,Q}, \text{span}\{\Gamma_{M,\tau}^{\varpi}\}) + \Delta_M^{\varpi,\text{vis}} + 2\|\tilde{S}_M^{\varpi,Q}\|_{L^2(\mathbb{Q})}. \end{aligned}$$

Adding the two bounds and dividing by $\alpha_{Q_M}^{\mathbb{Q},\star}$ gives the stated estimate for $\mathfrak{s}_{M,\tau}^{\varpi,Q}$. Vanishing of each normalized term on the right gives the asymptotic claim; simultaneous zero defects give the exact claim. $\square$

Theorem D.41 (Common centered ϖ -profile generator implies benchmark-factor compatibility). *Under Definition D.36 and Proposition D.37, let*

$$\tilde{L}_{\bar{Z}_{M,\tau}}^{\mathbb{Q}} = \bar{c}_{M,\tau}^{\mathbb{Q},\star} \Lambda_M^{\star} + \bar{R}_{M,\tau}^{\mathbb{Q},\star}$$

be the canonical aggregate-factor decomposition of the centered benchmark-profile loading, with

$$\bar{R}_{M,\tau}^{\mathbb{Q},\star} \perp_{L^2(\mathbb{Q})} \Lambda_M^{\star}.$$

If

$$d_{M,\tau}^{\text{bench}} \neq 0, \quad g_{M,\tau}^{\varpi,\star} \neq 0 \quad \text{for all sufficiently small } \tau,$$

define

$$s_{M,\tau}^{\varpi,\star} := \frac{g_{M,\tau}^{\varpi,\star}}{|g_{M,\tau}^{\varpi,\star}|}, \quad \delta_{M,\tau}^{\varpi,\text{gen}} := \mathfrak{d}_{M,\tau}^{\varpi,\text{bench}} + |d_{M,\tau}^{\text{bench}}| \frac{\mathfrak{d}_{M,\tau}^{\varpi,\star}}{|g_{M,\tau}^{\varpi,\star}|}.$$

Then

$$|\bar{c}_{M,\tau}^{\mathbb{Q},\star} - s_{M,\tau}^{\varpi,\star} d_{M,\tau}^{\text{bench}}| \leq \delta_{M,\tau}^{\varpi,\text{gen}},$$

and

$$\|\bar{R}_{M,\tau}^{\mathbb{Q},\star}\|_{L^2(\mathbb{Q})} \leq \delta_{M,\tau}^{\varpi,\text{gen}}.$$

If, moreover,

$$\delta_{M,\tau}^{\varpi,\text{gen}} = o(|d_{M,\tau}^{\text{bench}}|),$$

then the benchmark-profile direction is asymptotically aggregate-factor compatible:

$$\left\| \Lambda_{M,\tau}^{\text{bench}} - \frac{d_{M,\tau}^{\text{bench}} g_{M,\tau}^{\varpi,\star}}{|d_{M,\tau}^{\text{bench}} g_{M,\tau}^{\varpi,\star}|} \Lambda_M^\star \right\|_{L^2(\mathbb{Q})} \rightarrow 0.$$

If, in addition,

$$\mathfrak{d}_{M,\tau}^{\varpi,\text{bench}} = 0, \quad \mathfrak{d}_{M,\tau}^{\varpi,\star} = 0,$$

then

$$\bar{c}_{M,\tau}^{\mathbb{Q},\star} = s_{M,\tau}^{\varpi,\star} d_{M,\tau}^{\text{bench}}, \quad \bar{R}_{M,\tau}^{\mathbb{Q},\star} = 0,$$

and therefore

$$\Lambda_{M,\tau}^{\text{bench}} = \frac{d_{M,\tau}^{\text{bench}} g_{M,\tau}^{\varpi,\star}}{|d_{M,\tau}^{\text{bench}} g_{M,\tau}^{\varpi,\star}|} \Lambda_M^\star \quad \text{exactly.}$$

Proof. Set

$$\eta_{M,\tau}^{\varpi,\star} := \left\| \Gamma_{M,\tau}^{\varpi} - s_{M,\tau}^{\varpi,\star} \Lambda_M^\star \right\|_{L^2(\mathbb{Q})}.$$

Because $\Gamma_{M,\tau}^{\varpi}$ and $\Lambda_M^\star$ are unit vectors under the present nondegeneracy assumptions and

$$\Lambda_M^\star = g_{M,\tau}^{\varpi,\star} \Gamma_{M,\tau}^{\varpi} + H_{M,\tau}^{\varpi,\star}, \quad H_{M,\tau}^{\varpi,\star} \perp_{L^2(\mathbb{Q})} \Gamma_{M,\tau}^{\varpi},$$

one has

$$1 = (g_{M,\tau}^{\varpi,\star})^2 + \|H_{M,\tau}^{\varpi,\star}\|_{L^2(\mathbb{Q})}^2.$$

Writing

$$s_{M,\tau}^{\varpi,\star} \Lambda_M^\star = |g_{M,\tau}^{\varpi,\star}| \Gamma_{M,\tau}^{\varpi} + s_{M,\tau}^{\varpi,\star} H_{M,\tau}^{\varpi,\star},$$

and using again the orthogonality of $H_{M,\tau}^{\varpi,\star}$ and $\Gamma_{M,\tau}^{\varpi}$, one gets

$$(\eta_{M,\tau}^{\varpi,\star})^2 = \left(1 - |g_{M,\tau}^{\varpi,\star}|\right)^2 + \|H_{M,\tau}^{\varpi,\star}\|_{L^2(\mathbb{Q})}^2 = 2 \left(1 - |g_{M,\tau}^{\varpi,\star}|\right) \leq \frac{\|H_{M,\tau}^{\varpi,\star}\|_{L^2(\mathbb{Q})}^2}{(g_{M,\tau}^{\varpi,\star})^2},$$

so

$$\eta_{M,\tau}^{\varpi,\star} \leq \frac{\mathfrak{d}_{M,\tau}^{\varpi,\star}}{|g_{M,\tau}^{\varpi,\star}|}.$$

By Definition D.36,

$$\tilde{L}_{M,\tau}^{\mathbb{Q}} = s_{M,\tau}^{\varpi,\star} d_{M,\tau}^{\text{bench}} \Lambda_M^\star + d_{M,\tau}^{\text{bench}} (\Gamma_{M,\tau}^{\varpi} - s_{M,\tau}^{\varpi,\star} \Lambda_M^\star) + E_{M,\tau}^{\text{bench},\varpi}.$$

Taking the $L^2(\mathbb{Q})$ -inner product against $\Lambda_M^\star$ gives

$$\left| \bar{c}_{M,\tau}^{\mathbb{Q},\star} - s_{M,\tau}^{\varpi,\star} d_{M,\tau}^{\text{bench}} \right| \leq |d_{M,\tau}^{\text{bench}}| \eta_{M,\tau}^{\varpi,\star} + \|E_{M,\tau}^{\text{bench},\varpi}\|_{L^2(\mathbb{Q})} \leq \delta_{M,\tau}^{\varpi,\text{gen}}.$$

Subtracting the $\Lambda_M^\star$ -component from the same decomposition yields

$$\|\bar{R}_{M,\tau}^{\mathbb{Q},\star}\|_{L^2(\mathbb{Q})} \leq |d_{M,\tau}^{\text{bench}}| \eta_{M,\tau}^{\varpi,\star} + \|E_{M,\tau}^{\text{bench},\varpi}\|_{L^2(\mathbb{Q})} \leq \delta_{M,\tau}^{\varpi,\text{gen}}.$$

If $\delta_{M,\tau}^{\varpi,\text{gen}} = o(|d_{M,\tau}^{\text{bench}}|)$, then

$$\bar{c}_{M,\tau}^{\mathbb{Q},\star} = s_{M,\tau}^{\varpi,\star} d_{M,\tau}^{\text{bench}} + o(|d_{M,\tau}^{\text{bench}}|),$$

and

$$\|\bar{R}_{M,\tau}^{\mathbb{Q},\star}\|_{L^2(\mathbb{Q})} = o(|d_{M,\tau}^{\text{bench}}|).$$

Since $|\bar{c}_{M,\tau}^{\mathbb{Q},\star}| \sim |d_{M,\tau}^{\text{bench}}|$, the displayed benchmark-direction compatibility is exactly Corollary D.34.

If

$$\mathfrak{d}_{M,\tau}^{\varpi,\text{bench}} = 0, \quad \mathfrak{d}_{M,\tau}^{\varpi,\star} = 0,$$

then $\eta_{M,\tau}^{\varpi,\star} = 0$, so the exact identities follow. $\square$

Corollary D.42 (Benchmark-factor compatibility from the Q_M -side ϖ -shadow). *Assume the hypotheses of Theorem D.41 and Propositions D.39 and D.40. If*

$$\alpha_{Q_M}^{\mathbb{Q},\star} \neq 0,$$

and if

$$d_{M,\tau}^{\text{bench}} \neq 0 \quad \text{for all sufficiently small } \tau,$$

and

$$\mathfrak{d}_{M,\tau}^{\varpi,\text{bench}} = o(|d_{M,\tau}^{\text{bench}}|), \quad \mathfrak{s}_{M,\tau}^{\varpi,Q} = o(1),$$

then the benchmark-profile direction is asymptotically aggregate-factor compatible:

$$\left\| \Lambda_{M,\tau}^{\text{bench}} - \frac{d_{M,\tau}^{\text{bench}} g_{M,\tau}^{\varpi,\star}}{|d_{M,\tau}^{\text{bench}} g_{M,\tau}^{\varpi,\star}|} \Lambda_M^\star \right\|_{L^2(\mathbb{Q})} \rightarrow 0.$$

In particular, it is enough that

$$\mathfrak{d}_{M,\tau}^{\varpi,\text{bench}} = o(|d_{M,\tau}^{\text{bench}}|)$$

and

$$\|S_M^{\varpi,Q}\|_{L^2(\mathbb{Q})} + \Delta_M^{\varpi,\text{vis}} + \text{dist}(\tilde{B}_M^{\varpi,Q}, \text{span}\{\Gamma_{M,\tau}^{\varpi}\}) = o(\alpha_{Q_M}^{\mathbb{Q},\star}).$$

If, in addition,

$$\mathfrak{d}_{M,\tau}^{\varpi,\text{bench}} = 0, \quad \mathfrak{s}_{M,\tau}^{\varpi,Q} = 0,$$

then the benchmark-profile direction is exactly aggregate-factor compatible. In particular, exact compatibility follows if

$$\mathfrak{d}_{M,\tau}^{\varpi,\text{bench}} = 0, \quad S_M^{\varpi,Q} = 0, \quad \Delta_M^{\varpi,\text{vis}} = 0, \quad \tilde{B}_M^{\varpi,Q} \in \text{span}\{\Gamma_{M,\tau}^{\varpi}\}.$$

Proof. Proposition D.39 gives

$$\mathfrak{d}_{M,\tau}^{\varpi,\star} \leq \mathfrak{s}_{M,\tau}^{\varpi,Q} = o(1),$$

so

$$1 = (g_{M,\tau}^{\varpi,\star})^2 + (\mathfrak{d}_{M,\tau}^{\varpi,\star})^2$$

implies $|g_{M,\tau}^{\varpi,\star}| \rightarrow 1$, hence $g_{M,\tau}^{\varpi,\star} \neq 0$ for all sufficiently small τ . Therefore

$$\delta_{M,\tau}^{\varpi,\text{gen}} = \mathfrak{d}_{M,\tau}^{\varpi,\text{bench}} + |d_{M,\tau}^{\text{bench}}| \frac{\mathfrak{d}_{M,\tau}^{\varpi,\star}}{|g_{M,\tau}^{\varpi,\star}|} = o(|d_{M,\tau}^{\text{bench}}|),$$

and Theorem D.41 yields the asymptotic claim. The explicit sufficient conditions follow from Proposition D.40. If both displayed defects vanish, then Proposition D.39 gives $\mathfrak{d}_{M,\tau}^{\varpi,\star} = 0$, and the exact branch of Theorem D.41 applies; the exact sufficient conditions are again those of Proposition D.40. $\square$

Proposition D.43 (Mixed ϖ -shadow remainder split). *Under the notation of Proposition D.40, define*

$$S_{M,\times}^{\varpi,Q} := 2 \int_0^T M^{\text{vis},\mathbb{Q}}(t) \zeta(t) dt, \quad S_{M,\eta}^{\varpi,Q} := R_{\eta,M}^{\mathbb{Q}} - C_{\mathcal{L}_T}^{\mathbb{Q}}(R_{\eta,M}^{\mathbb{Q}}).$$

Then

$$S_M^{\varpi,Q} = S_{M,\times}^{\varpi,Q} + S_{M,\eta}^{\varpi,Q}.$$

Consequently,

$$\|S_M^{\varpi,Q}\|_{L^2(\mathbb{Q})} \leq \|S_{M,\times}^{\varpi,Q}\|_{L^2(\mathbb{Q})} + \|S_{M,\eta}^{\varpi,Q}\|_{L^2(\mathbb{Q})}.$$

If

$$\|S_{M,\times}^{\varpi,Q}\|_{L^2(\mathbb{Q})} + \|S_{M,\eta}^{\varpi,Q}\|_{L^2(\mathbb{Q})} = o(\alpha_{Q_M}^{\mathbb{Q},\star}),$$

then

$$\|S_M^{\varpi,Q}\|_{L^2(\mathbb{Q})} = o(\alpha_{Q_M}^{\mathbb{Q},\star}).$$

If

$$S_{M,\times}^{\varpi,Q} = 0, \quad S_{M,\eta}^{\varpi,Q} = 0,$$

then

$$S_M^{\varpi,Q} = 0.$$

Proof. Substituting the definition of $S_M^{\varpi, Q}$ from Proposition D.40 gives the identity. Applying the triangle inequality to the three summands gives the displayed bounds; if the summands vanish, the sum is zero. $\square$

Proposition D.44 (Centered split of the mixed ϖ -shadow remainder). *Under the notation of Proposition D.43, define*

$$\tilde{S}_{M, \times}^{\varpi, Q} := S_{M, \times}^{\varpi, Q} - C_{\mathcal{G}_T}^{\mathbb{Q}}(S_{M, \times}^{\varpi, Q}), \quad \tilde{S}_{M, \eta}^{\varpi, Q} := S_{M, \eta}^{\varpi, Q} - C_{\mathcal{G}_T}^{\mathbb{Q}}(S_{M, \eta}^{\varpi, Q}).$$

Then

$$\tilde{S}_M^{\varpi, Q} = \tilde{S}_{M, \times}^{\varpi, Q} + \tilde{S}_{M, \eta}^{\varpi, Q},$$

and therefore

$$\|\tilde{S}_M^{\varpi, Q}\|_{L^2(\mathbb{Q})} \leq \|\tilde{S}_{M, \times}^{\varpi, Q}\|_{L^2(\mathbb{Q})} + \|\tilde{S}_{M, \eta}^{\varpi, Q}\|_{L^2(\mathbb{Q})}.$$

Moreover,

$$\|\tilde{S}_{M, \times}^{\varpi, Q}\|_{L^2(\mathbb{Q})} \leq 2\|S_{M, \times}^{\varpi, Q}\|_{L^2(\mathbb{Q})}, \quad \|\tilde{S}_{M, \eta}^{\varpi, Q}\|_{L^2(\mathbb{Q})} \leq 2\|S_{M, \eta}^{\varpi, Q}\|_{L^2(\mathbb{Q})}.$$

If

$$S_{M, \times}^{\varpi, Q} \in L^2(\mathbb{Q}, \mathcal{G}_T), \quad S_{M, \eta}^{\varpi, Q} \in L^2(\mathbb{Q}, \mathcal{G}_T),$$

then

$$\tilde{S}_{M, \times}^{\varpi, Q} = 0, \quad \tilde{S}_{M, \eta}^{\varpi, Q} = 0.$$

Proof. The identity follows from

$$\tilde{S}_M^{\varpi, Q} = S_M^{\varpi, Q} - C_{\mathcal{G}_T}^{\mathbb{Q}}(S_M^{\varpi, Q})$$

and the decomposition $S_M^{\varpi, Q} = S_{M, \times}^{\varpi, Q} + S_{M, \eta}^{\varpi, Q}$ from Proposition D.43. The norm bound is the triangle inequality. Since $C_{\mathcal{G}_T}^{\mathbb{Q}}$ is an L^2 -contraction,

$$\|Y - C_{\mathcal{G}_T}^{\mathbb{Q}}(Y)\|_{L^2(\mathbb{Q})} \leq \|Y\|_{L^2(\mathbb{Q})} + \|C_{\mathcal{G}_T}^{\mathbb{Q}}(Y)\|_{L^2(\mathbb{Q})} \leq 2\|Y\|_{L^2(\mathbb{Q})}$$

for every $Y \in L^2(\mathbb{Q}, \mathcal{F}_T)$, which yields the displayed bounds. If $Y \in L^2(\mathbb{Q}, \mathcal{G}_T)$, then $C_{\mathcal{G}_T}^{\mathbb{Q}}(Y) = Y$, so the centered residual is zero and the exact claim follows. $\square$

Proposition D.45 (Centered mixed term as visible–public projection interference). *Assume Assumption 10.6 and the notation of Propositions 10.9, D.43, and D.44. Define the public projection of the synchronized aggregate channel by*

$$\tilde{\zeta}^{\mathbb{Q}}(t) := \zeta(t) - C_{\mathcal{G}_T}^{\mathbb{Q}}(\zeta(t)), \quad 0 \leq t \leq T.$$

Then

$$\tilde{S}_{M, \times}^{\varpi, Q} = 2 \int_0^T M^{\text{vis}, \mathbb{Q}}(t) \tilde{\zeta}^{\mathbb{Q}}(t) dt.$$

Consequently,

$$\|\tilde{S}_{M, \times}^{\varpi, Q}\|_{L^2(\mathbb{Q})} \leq 2 \left\| \int_0^T (M^{\text{vis}, \mathbb{Q}}(t))^2 dt \right\|_{L^2(\mathbb{Q})}^{1/2} \left\| \int_0^T (\tilde{\zeta}^{\mathbb{Q}}(t))^2 dt \right\|_{L^2(\mathbb{Q})}^{1/2}.$$

If

$$\zeta(t) \in L^2(\mathbb{Q}, \mathcal{G}_T) \quad \text{for almost every } t \in [0, T],$$

then

$$S_{M, \times}^{\varpi, Q} \in L^2(\mathbb{Q}, \mathcal{G}_T), \quad \tilde{S}_{M, \times}^{\varpi, Q} = 0.$$

Proof. Because $M^{\text{vis}, \mathbb{Q}}(t)$ is $\mathcal{G}_t$ -measurable for every t , it is also $\mathcal{G}_T$ -measurable. Therefore

$$C_{\mathcal{G}_T}^{\mathbb{Q}}(M^{\text{vis}, \mathbb{Q}}(t)\zeta(t)) = M^{\text{vis}, \mathbb{Q}}(t) C_{\mathcal{G}_T}^{\mathbb{Q}}(\zeta(t))$$

for almost every t . Applying conditional Fubini gives

$$C_{\mathcal{G}_T}^{\mathbb{Q}}(S_{M, \times}^{\varpi, Q}) = 2 \int_0^T M^{\text{vis}, \mathbb{Q}}(t) C_{\mathcal{G}_T}^{\mathbb{Q}}(\zeta(t)) dt.$$

Subtracting this from

$$S_{M,\times}^{\varpi,Q} = 2 \int_0^T M^{\text{vis},Q}(t) \zeta(t) dt$$

yields the displayed identity for $\tilde{S}_{M,\times}^{\varpi,Q}$. The L^2 bound then follows from pathwise Cauchy–Schwarz:

$$\left| \int_0^T M^{\text{vis},Q}(t) \tilde{\zeta}^Q(t) dt \right|^2 \leq \left(\int_0^T (M^{\text{vis},Q}(t))^2 dt \right) \left(\int_0^T (\tilde{\zeta}^Q(t))^2 dt \right),$$

followed by Cauchy–Schwarz in expectation. If $\zeta(t)$ is $\mathcal{G}_T$ -measurable for almost every t , then $M^{\text{vis},Q}(t)\zeta(t)$ is $\mathcal{G}_T$ -measurable for almost every t , so $S_{M,\times}^{\varpi,Q} \in L^2(\mathbb{Q}, \mathcal{G}_T)$, and the centered term vanishes by definition. $\square$

Proposition D.46 (Visible-screening split of the pure ϖ -benchmark defect). *Under the notation of Proposition D.40, define*

$$B_{M,\zeta}^{\varpi,Q} := \int_0^T \zeta(t)^2 dt, \quad B_{M,\eta}^{\varpi,Q} := C_{\mathcal{L}_T}^Q(R_{\eta,M}^Q),$$

so that

$$B_M^{\varpi,Q} = B_{M,\zeta}^{\varpi,Q} + B_{M,\eta}^{\varpi,Q}.$$

Write

$$\Pi_{\varpi}^Q(Y) := C_{\mathcal{L}_T}^Q(Y) - C_{\mathcal{G}_T}^Q(C_{\mathcal{L}_T}^Q(Y)).$$

Define

$$\Delta_{M,\zeta}^{\varpi,\text{vis}} := \left\| \Pi_{\varpi}^Q \left(C_{\mathcal{G}_T}^Q(B_{M,\zeta}^{\varpi,Q}) \right) \right\|_{L^2(\mathbb{Q})}, \quad \Delta_{M,\eta}^{\varpi,\text{vis}} := \left\| \Pi_{\varpi}^Q \left(C_{\mathcal{G}_T}^Q(B_{M,\eta}^{\varpi,Q}) \right) \right\|_{L^2(\mathbb{Q})}.$$

Then

$$\Delta_M^{\varpi,\text{vis}} \leq \Delta_{M,\zeta}^{\varpi,\text{vis}} + \Delta_{M,\eta}^{\varpi,\text{vis}}.$$

Consequently, if

$$\Delta_{M,\zeta}^{\varpi,\text{vis}} + \Delta_{M,\eta}^{\varpi,\text{vis}} = o(\alpha_{Q_M}^{\mathbb{Q},\star}),$$

then

$$\Delta_M^{\varpi,\text{vis}} = o(\alpha_{Q_M}^{\mathbb{Q},\star}).$$

If

$$\Delta_{M,\zeta}^{\varpi,\text{vis}} = 0, \quad \Delta_{M,\eta}^{\varpi,\text{vis}} = 0,$$

then

$$\Delta_M^{\varpi,\text{vis}} = 0.$$

Proof. Because $B_M^{\varpi,Q} = B_{M,\zeta}^{\varpi,Q} + B_{M,\eta}^{\varpi,Q}$, linearity of Π_{ϖ}^Q gives

$$\Pi_{\varpi}^Q \left(C_{\mathcal{G}_T}^Q(B_M^{\varpi,Q}) \right) = \Pi_{\varpi}^Q \left(C_{\mathcal{G}_T}^Q(B_{M,\zeta}^{\varpi,Q}) \right) + \Pi_{\varpi}^Q \left(C_{\mathcal{G}_T}^Q(B_{M,\eta}^{\varpi,Q}) \right).$$

Taking $L^2(\mathbb{Q})$ -norms and applying the triangle inequality yields the displayed bound. If the normalized right-hand terms vanish, the left-hand defect vanishes asymptotically; if the component terms are zero, the centered pure benchmark defect is zero. $\square$

Proposition D.47 (Centered ϖ -benchmark split relative to the benchmark-generated generator). *Under the notation of Proposition D.46, define the centered pieces*

$$\tilde{B}_{M,\zeta}^{\varpi,Q} := B_{M,\zeta}^{\varpi,Q} - C_{\mathcal{G}_T}^Q(B_{M,\zeta}^{\varpi,Q}), \quad \tilde{B}_{M,\eta}^{\varpi,Q} := B_{M,\eta}^{\varpi,Q} - C_{\mathcal{G}_T}^Q(B_{M,\eta}^{\varpi,Q}),$$

so that

$$\tilde{B}_M^{\varpi,Q} = \tilde{B}_{M,\zeta}^{\varpi,Q} + \tilde{B}_{M,\eta}^{\varpi,Q}.$$

Define the canonical centered ϖ -benchmark coefficient and residual by

$$b_{M,\tau}^{\varpi,Q} := \left\langle \tilde{B}_M^{\varpi,Q}, \Gamma_{M,\tau}^{\varpi} \right\rangle_{L^2(\mathbb{Q})}, \quad E_{M,\tau}^{\varpi,B} := \tilde{B}_M^{\varpi,Q} - b_{M,\tau}^{\varpi,Q} \Gamma_{M,\tau}^{\varpi}.$$

Define the centered ϖ -benchmark generator defect by

$$\mathfrak{d}_{M,\tau}^{\varpi,B} := \text{dist} \left(\tilde{B}_{M,\zeta}^{\varpi,Q}, \text{span} \{ \Gamma_{M,\tau}^{\varpi} \} \right) + \text{dist} \left(\tilde{B}_{M,\eta}^{\varpi,Q}, \text{span} \{ \Gamma_{M,\tau}^{\varpi} \} \right).$$

Then

$$\text{dist}\left(\tilde{B}_M^{\varpi,Q}, \text{span}\{\Gamma_{M,\tau}^{\varpi}\}\right) \leq \mathfrak{d}_{M,\tau}^{\varpi,B}.$$

Moreover,

$$E_{M,\tau}^{\varpi,B} \in \mathcal{M}_{\mathbb{Q}}^{\circ}, \quad E_{M,\tau}^{\varpi,B} \perp_{L^2(\mathbb{Q})} \Gamma_{M,\tau}^{\varpi}, \quad \|E_{M,\tau}^{\varpi,B}\|_{L^2(\mathbb{Q})} = \text{dist}\left(\tilde{B}_M^{\varpi,Q}, \text{span}\{\Gamma_{M,\tau}^{\varpi}\}\right).$$

In particular, if

$$\mathfrak{d}_{M,\tau}^{\varpi,B} = o\left(\alpha_{Q_M}^{\mathbb{Q},\star}\right),$$

then

$$\text{dist}\left(\tilde{B}_M^{\varpi,Q}, \text{span}\{\Gamma_{M,\tau}^{\varpi}\}\right) = o\left(\alpha_{Q_M}^{\mathbb{Q},\star}\right).$$

If

$$\mathfrak{d}_{M,\tau}^{\varpi,B} = 0,$$

then

$$\tilde{B}_M^{\varpi,Q} \in \text{span}\{\Gamma_{M,\tau}^{\varpi}\}.$$

Proof. Centering the decomposition of $B_M^{\varpi,Q}$ gives the displayed split for $\tilde{B}_M^{\varpi,Q}$. Applying the triangle inequality for the distance to the subspace $\text{span}\{\Gamma_{M,\tau}^{\varpi}\}$ gives the displayed bound. If the two distance terms vanish asymptotically, the total distance does also; if both centered components lie in the span, so does their sum. $\square$

Proposition D.48 (Reduction of the centered ϖ -benchmark defect to the centered ϖ^2 benchmark). *Under the notation of Propositions D.46 and D.47, define the centered synchronized-square benchmark by*

$$Q_{\varpi}^{\mathbb{Q}} := \int_0^T \varpi(t)^2 dt, \quad \tilde{Q}_{\varpi}^{\mathbb{Q}} := Q_{\varpi}^{\mathbb{Q}} - C_{\mathcal{G}_T}^{\mathbb{Q}}(Q_{\varpi}^{\mathbb{Q}}).$$

Then

$$\tilde{B}_{M,\zeta}^{\varpi,Q} = (w^{\top}v)^2 \tilde{Q}_{\varpi}^{\mathbb{Q}},$$

and therefore

$$\mathfrak{d}_{M,\tau}^{\varpi,B} = |w^{\top}v|^2 \text{dist}\left(\tilde{Q}_{\varpi}^{\mathbb{Q}}, \text{span}\{\Gamma_{M,\tau}^{\varpi}\}\right) + \text{dist}\left(\tilde{B}_{M,\eta}^{\varpi,Q}, \text{span}\{\Gamma_{M,\tau}^{\varpi}\}\right).$$

In particular, if

$$\tilde{B}_{M,\eta}^{\varpi,Q} \in \text{span}\{\Gamma_{M,\tau}^{\varpi}\},$$

then

$$\mathfrak{d}_{M,\tau}^{\varpi,B} = |w^{\top}v|^2 \text{dist}\left(\tilde{Q}_{\varpi}^{\mathbb{Q}}, \text{span}\{\Gamma_{M,\tau}^{\varpi}\}\right).$$

If, more strongly,

$$B_{M,\eta}^{\varpi,Q} \in L^2(\mathbb{Q}, \mathcal{G}_T),$$

then

$$\tilde{B}_{M,\eta}^{\varpi,Q} = 0 \quad \text{and hence} \quad \mathfrak{d}_{M,\tau}^{\varpi,B} = |w^{\top}v|^2 \text{dist}\left(\tilde{Q}_{\varpi}^{\mathbb{Q}}, \text{span}\{\Gamma_{M,\tau}^{\varpi}\}\right).$$

Proof. Because $\zeta(t) = (w^{\top}v)\varpi(t)$,

$$B_{M,\zeta}^{\varpi,Q} = \int_0^T \zeta(t)^2 dt = (w^{\top}v)^2 \int_0^T \varpi(t)^2 dt = (w^{\top}v)^2 Q_{\varpi}^{\mathbb{Q}}.$$

Subtracting $C_{\mathcal{G}_T}^{\mathbb{Q}}(\cdot)$ gives

$$\tilde{B}_{M,\zeta}^{\varpi,Q} = (w^{\top}v)^2 \tilde{Q}_{\varpi}^{\mathbb{Q}}.$$

Since distance to a linear subspace is homogeneous,

$$\text{dist}\left(\tilde{B}_{M,\zeta}^{\varpi,Q}, \text{span}\{\Gamma_{M,\tau}^{\varpi}\}\right) = |w^{\top}v|^2 \text{dist}\left(\tilde{Q}_{\varpi}^{\mathbb{Q}}, \text{span}\{\Gamma_{M,\tau}^{\varpi}\}\right).$$

The displayed identity for $\mathfrak{d}_{M,\tau}^{\varpi,B}$ then follows from the definition in Proposition D.47. If $\tilde{B}_{M,\eta}^{\varpi,Q} \in \text{span}\{\Gamma_{M,\tau}^{\varpi}\}$, the second distance term vanishes, giving the first refinement. If $B_{M,\eta}^{\varpi,Q} \in L^2(\mathbb{Q}, \mathcal{G}_T)$, then

$$\tilde{B}_{M,\eta}^{\varpi,Q} = B_{M,\eta}^{\varpi,Q} - C_{\mathcal{G}_T}^{\mathbb{Q}}(B_{M,\eta}^{\varpi,Q}) = 0,$$

which gives the exact final identity. $\square$

Proposition D.49 (Centered ϖ^2 benchmark coefficient-residual decomposition). *Under the notation of Proposition D.48, define the centered ϖ^2 -benchmark coefficient and residual by*

$$b_{M,\tau}^{\varpi,2} := \left\langle \tilde{Q}_\varpi^\mathbb{Q}, \Gamma_{M,\tau}^\varpi \right\rangle_{L^2(\mathbb{Q})}, \quad E_{M,\tau}^{\varpi,2} := \tilde{Q}_\varpi^\mathbb{Q} - b_{M,\tau}^{\varpi,2} \Gamma_{M,\tau}^\varpi.$$

Then

$$E_{M,\tau}^{\varpi,2} \in \mathcal{M}_\mathbb{Q}^\circ, \quad E_{M,\tau}^{\varpi,2} \perp_{L^2(\mathbb{Q})} \Gamma_{M,\tau}^\varpi, \quad \|E_{M,\tau}^{\varpi,2}\|_{L^2(\mathbb{Q})} = \text{dist}\left(\tilde{Q}_\varpi^\mathbb{Q}, \text{span}\{\Gamma_{M,\tau}^\varpi\}\right).$$

If, in addition,

$$B_{M,\eta}^{\varpi,Q} \in L^2(\mathbb{Q}, \mathcal{G}_T),$$

then

$$\tilde{B}_M^{\varpi,Q} = (w^\top v)^2 \tilde{Q}_\varpi^\mathbb{Q}, \quad b_{M,\tau}^{\varpi,Q} = (w^\top v)^2 b_{M,\tau}^{\varpi,2}, \quad E_{M,\tau}^{\varpi,B} = (w^\top v)^2 E_{M,\tau}^{\varpi,2},$$

and

$$\mathfrak{d}_{M,\tau}^{\varpi,B} = |w^\top v|^2 \|E_{M,\tau}^{\varpi,2}\|_{L^2(\mathbb{Q})}.$$

Proof. Because $\Gamma_{M,\tau}^\varpi \in \mathcal{M}_\mathbb{Q}^\circ$, the centered residual $E_{M,\tau}^{\varpi,2}$ belongs to $\mathcal{M}_\mathbb{Q}^\circ$. The coefficient definition gives

$$\left\langle E_{M,\tau}^{\varpi,2}, \Gamma_{M,\tau}^\varpi \right\rangle_{L^2(\mathbb{Q})} = \left\langle \tilde{Q}_\varpi^\mathbb{Q}, \Gamma_{M,\tau}^\varpi \right\rangle_{L^2(\mathbb{Q})} - b_{M,\tau}^{\varpi,2} \|\Gamma_{M,\tau}^\varpi\|_{L^2(\mathbb{Q})}^2 = 0,$$

where the last identity is trivial if $\Gamma_{M,\tau}^\varpi = 0$. Hence $\|E_{M,\tau}^{\varpi,2}\|_{L^2(\mathbb{Q})}$ is the distance from $\tilde{Q}_\varpi^\mathbb{Q}$ to $\text{span}\{\Gamma_{M,\tau}^\varpi\}$. If $B_{M,\eta}^{\varpi,Q} \in L^2(\mathbb{Q}, \mathcal{G}_T)$, Proposition D.48 gives

$$\tilde{B}_M^{\varpi,Q} = \tilde{B}_{M,\zeta}^{\varpi,Q} = (w^\top v)^2 \tilde{Q}_\varpi^\mathbb{Q}.$$

Taking inner products against $\Gamma_{M,\tau}^\varpi$ yields

$$b_{M,\tau}^{\varpi,Q} = (w^\top v)^2 b_{M,\tau}^{\varpi,2},$$

and subtracting the aligned component gives

$$E_{M,\tau}^{\varpi,B} = \tilde{B}_M^{\varpi,Q} - b_{M,\tau}^{\varpi,Q} \Gamma_{M,\tau}^\varpi = (w^\top v)^2 E_{M,\tau}^{\varpi,2}.$$

Taking norms and using Proposition D.47 gives the final identity. $\square$

Assumption D.50 (Leading-generator dominance of the centered ϖ -benchmark). In the nondegenerate regime

$$b_{M,\tau}^{\varpi,Q} \neq 0,$$

the centered ϖ -benchmark generator defect from Proposition D.47 satisfies

$$\mathfrak{d}_{M,\tau}^{\varpi,B} = o(|b_{M,\tau}^{\varpi,Q}|) \quad \text{as } \tau \downarrow 0.$$

In the exact branch, assume instead

$$\mathfrak{d}_{M,\tau}^{\varpi,B} = 0 \quad \text{for all sufficiently small } \tau.$$

Corollary D.51 (Residual-variance screened reduction of the benchmark-generator defect). *Assume the hypotheses of Proposition D.49, and suppose*

$$B_{M,\eta}^{\varpi,Q} \in L^2(\mathbb{Q}, \mathcal{G}_T), \quad w^\top v \neq 0.$$

Assume further that

$$b_{M,\tau}^{\varpi,2} \neq 0 \quad \text{for all sufficiently small } \tau.$$

then Assumption D.50 follows from

$$\|E_{M,\tau}^{\varpi,2}\|_{L^2(\mathbb{Q})} = o(|b_{M,\tau}^{\varpi,2}|) \quad \text{as } \tau \downarrow 0.$$

In the exact branch, it is enough that

$$E_{M,\tau}^{\varpi,2} = 0 \quad \text{for all sufficiently small } \tau.$$

Proof. Under $B_{M,\eta}^{\varpi,Q} \in L^2(\mathbb{Q}, \mathcal{G}_T)$, Proposition D.49 gives

$$\mathfrak{d}_{M,\tau}^{\varpi,B} = |w^\top v|^2 \|E_{M,\tau}^{\varpi,2}\|_{L^2(\mathbb{Q})} \quad \text{and} \quad b_{M,\tau}^{\varpi,Q} = (w^\top v)^2 b_{M,\tau}^{\varpi,2}.$$

Since $w^\top v \neq 0$, the nondegenerate regimes $b_{M,\tau}^{\varpi,Q} \neq 0$ and $b_{M,\tau}^{\varpi,2} \neq 0$ coincide. Therefore

$$\|E_{M,\tau}^{\varpi,2}\|_{L^2(\mathbb{Q})} = o(|b_{M,\tau}^{\varpi,2}|) \implies \mathfrak{d}_{M,\tau}^{\varpi,B} = o(|b_{M,\tau}^{\varpi,Q}|),$$

which is exactly Assumption D.50. If $E_{M,\tau}^{\varpi,2} = 0$, then $\mathfrak{d}_{M,\tau}^{\varpi,B} = 0$, giving the exact branch. $\square$

D.1. Scalar Gaussian seam and Companion A source-family verification. The primitive Gaussian-Volterra source-family verification is long and mechanical. It is kept in Companion A rather than in the main body. This appendix records only the scalar seam used by the option-surface branch. Companion A verifies this seam from raw Volterra kernel primitives, Hilbert-Schmidt one-mode structure, leading-coordinate control, centered-square comparison, and natural-scale lower-order defect bounds.

Assumption D.52 (Gaussian scalar seam inputs). For the benchmark maturity family M , assume the primitive Gaussian-Volterra hypotheses and source-family estimates verified in Companion A. In scalar form, this means that there exist common channel weights $w = (w_j)_j$, a nonzero benchmark coefficient B_M^{quad} , and a natural scale $a_{M,\tau}^{\text{quad}} > 0$ such that the centered quadratic option block has the shrinking-triangle expansion

$$\frac{1}{a_{M,\tau}^{\text{quad}}} \left(\Gamma_{M,\tau}^{\mathbb{Q}} - B_M^{\text{quad}} \sum_j w_j Z_{j,\tau} \right) \longrightarrow 0 \quad \text{in } L^2(\mathbb{Q})$$

uniformly along the local maturity triangle. The benchmark coefficient is nonzero and the synchronized channel weights are the same weights selected by the Perron branch.

Proposition D.53 (Companion-verified scalar seam reduction). *Under Assumption D.52, the Gaussian-Volterra benchmark produces the scalar natural-scale seam needed by the local option-surface branch. Equivalently, the normalized centered quadratic profile has a nonzero common-channel limit in the Perron coordinate, and every singular-value, scalar-comparison, and option-block defect is lower order at that scale.*

Proof. Companion A proves the three primitive tasks. First, the synchronized-amplitude coefficient, the scalar quadratic comparison coefficient, and the benchmark-generated profile coefficient are nonzero. Second, the singular-value tail defect, scalar comparison defect, and combined option block defect are all $o(a_{M,\tau}^{\text{quad}})$. Third, the leading-coordinate map agrees with the common Perron channel weights. Substituting those three estimates into the displayed scalar expansion in Assumption D.52 gives the claimed seam. No additional theorem-side wrapper is used in the main body. $\square$

Remark D.54 (Companion source-family verification). The detailed source-family arguments that verify the primitive estimates live in the companion technical note. The appendix summary uses Assumption D.52, Proposition D.53, and the compact source-family verification below.

D.2. Gaussian-Volterra source-family verification. The companion Gaussian source-family note enters the compression argument through the compact assumption and verification proposition below, followed by Propositions D.58 and D.59. The branch-by-branch Volterra estimates remain in Companion A rather than in this article.

Assumption D.55 (Gaussian public source-family inputs). All derived results that invoke these inputs are read under the standing Gaussian-Volterra, Perron-proxy, benchmark-generator, scalar-comparison, and option-defect hypotheses needed to define the common-channel-weight structure. In particular, the following named objects and consequences are available:

- (i) the synchronized-amplitude coefficient, the scalar quadratic comparison coefficient, and the benchmark-generated profile coefficient are the three leading coefficients for the reduction;

- (ii) the singular-value tail defect, the scalar comparison defect, and the combined option-block defect are the three defect families whose lower-order estimates are measured at their natural scales;
- (iii) exact-branch implication templates use the corresponding exact vanishing of those defect families, while asymptotic-branch templates use the displayed lower-order estimates stated in the corollary itself or in the named source-family verification result being applied;
- (iv) the non-temporal, temporal singular-value, combined source-family, visible-origin, lower-cell, and denominator local-flatness branches are available only through their named verification or terminal estimates;
- (v) the non-temporal, temporal singular-value, and combined source-family verification results are available through Companion A whenever their own hypotheses are separately verified.

This assumption records the notation and finite dependency list used by the appendix. It does not by itself assert nonvanishing of the three leading coefficients, lower-order smallness of the three defects, or exact-branch zero defect identities. Those conditions must be stated in the result that uses these inputs, or supplied by the companion source-family verification.

Remark D.56 (Gaussian public source-family argument). The common-channel-weight reduction is used here as the input list for Proposition D.58. Its three coefficients are the synchronized amplitude coefficient, the scalar quadratic comparison coefficient, and the benchmark-generated profile coefficient. Its three defects are the singular-value tail defect, the scalar comparison defect, and the combined option-block defect. Once the stated coefficient lower bounds and lower-order defect estimates are supplied, the public Gaussian reduction follows.

The source-family verifications are proved in Companion A. Here the grouped argument is represented only by the compact verification proposition Proposition D.57.

On the model side, Companion A verifies the raw Volterra singular-value chain: kernel primitives, Hilbert–Schmidt approximation, one-mode and leading-coordinate structure, centered square lift, scalar comparison profile control, and the natural-scale scalar seam. The visible-origin/lower-cell branch and the denominator local-flatness branch remain part of that companion verification; in particular, the rescaled first-order fiber H^1 -flat profile and the L^2 -variation predecessor are terminal estimates recorded there.

Proposition D.57 (Companion Gaussian source-family verification). *Assume the source-family inputs in Assumption D.55 and the scalar seam of Proposition D.53. If a source-family argument in the companion technical note verifies the source separation, lower-order defect, exact-branch, and quote-transfer conditions required by those inputs, then it supplies the common-channel input required by the Gaussian public argument.*

Proof. The proposition is a dispatch statement from the companion verification to the present Gaussian public argument. The companion verification checks the primitive source-family conditions branch by branch. It proves the required nonzero source contribution, the natural-scale temporal estimate, the non-cancellation condition, the public-compression coefficient, and the quote-transfer remainder. These are the inputs listed in Assumption D.55, so the verified branch contributes the common-channel input asserted here. $\square$

Proposition D.58 (Companion-verified common-channel reduction). *Under Assumption D.55, the scalar seam of Proposition D.53, and a successful Companion A source-family verification in Proposition D.57, the Gaussian public argument has a nonzero common-channel coefficient at the local maturity anchor. In particular, the public option-facing coordinate is aligned with the Perron direction up to lower-order residuals at the stated maturity scale.*

Proof. The scalar seam supplies the normalized common-channel profile. Assumption D.55 fixes the public objects and nonzero coefficient convention. The companion verification supplies the branch-specific estimates and defect controls. The common-channel conclusion is the intersection of these three inputs. $\square$

Proposition D.59 (Gaussian option-factor consequence). *In the Gaussian-Volterra/Bachelier specialization, suppose Assumption D.55 holds and the common-channel reduction of Proposition D.58 has been verified by the companion source-family argument. Then the option-facing factor coefficient inherits the same Perron-aligned leading channel as the scalar seam. Consequently the local Bachelier surface coefficient is governed, to first order, by the compressed Perron stress coordinate, while branch-specific singular-value tails, scalar-comparison errors, and option-block defects are lower order.*

Proof. Proposition D.58 identifies the nonzero common-channel coefficient. The option-facing transfer map is linear at the local maturity anchor under the Gaussian/Bachelier source-family hypotheses and its remainder is part of the lower-order option-block defect. Applying that transfer to the scalar seam preserves the leading Perron-aligned coefficient and absorbs all controlled defects into the lower-order term. $\square$

Remark D.60 (Companion location of option-defect source arguments). The combined option-block defect, direct option-side maturity-scale hypotheses, and Gaussian public source-family checks are verified in Companion A. Only the consequence needed for quoted surfaces is kept here. Once the companion verification applies, the Bachelier/Gaussian surface coefficient is governed by the common-channel Perron factor.

Corollary D.61 (Gaussian-Volterra/Bachelier specialization of the non-Gaussian skew theorem). *Under the assumptions of Corollary D.11, let*

$$\Gamma_{M,\tau}^{\mathbb{Q}} := C_{\mathcal{H}_{M,\tau}^{\text{opt}}}^{\mathbb{Q}}(L_{M,\tau}^{\mathbb{Q}})$$

be the compressed option-facing priced structural loading from Proposition D.12. Because $\mathcal{H}_{M,\tau}^{\text{opt}}$ is trivial, $\Gamma_{M,\tau}^{\mathbb{Q}}$ is deterministic. Then, as $\tau \downarrow 0$,

$$\text{if } \mathcal{E}_M^{\text{ret}} = \emptyset, \quad \partial_\varepsilon \mathfrak{S}_{M,N}(\tau; \varepsilon)|_{\varepsilon=0} = \Gamma_{M,\tau}^{\mathbb{Q}} = 0;$$

If $\mathcal{E}_M^{\text{ret}} \neq \emptyset$ but $\text{supp}^{\text{comp}}(\Gamma_{M,\tau}^{\mathbb{Q}}) = \emptyset$, then the Gaussian coefficient system selects no nontrivial leading public power. In the nontrivial case

$$\mathcal{E}_M^{\text{ret}} \neq \emptyset \quad \text{and} \quad \text{supp}^{\text{comp}}(\Gamma_{M,\tau}^{\mathbb{Q}}) \neq \emptyset,$$

assume in addition that the option-compression remainder from Proposition D.12 is negligible at the selected cumulant scale:

$$r_{M,\tau} = o\left(\tau^{\beta_M^{\text{comp}}-1/2}\right).$$

Then

$$\partial_\varepsilon \mathfrak{S}_{M,N}(\tau; \varepsilon)|_{\varepsilon=0} = \Gamma_{M,\tau}^{\mathbb{Q}} + o\left(\tau^{\beta_M^{\text{comp}}-1/2}\right),$$

and

$$\Gamma_{M,\tau}^{\mathbb{Q}} = \sum_{\beta \in \text{supp}^{\text{comp}}(\Gamma_{M,\tau}^{\mathbb{Q}})} A_M(\beta) \tau^{\beta-1/2} + o\left(\tau^{\beta_M^{\text{comp}}-1/2}\right).$$

Equivalently, when $\alpha_{M,\tau}^{\mathbb{Q}} > 0$,

$$\alpha_{M,\tau}^{\mathbb{Q}} C_{\mathcal{H}_{M,\tau}^{\text{opt}}}^{\mathbb{Q}}(\Lambda_{M,\tau}) = \sum_{\beta \in \text{supp}^{\text{comp}}(\Gamma_{M,\tau}^{\mathbb{Q}})} A_M(\beta) \tau^{\beta-1/2} + o\left(\tau^{\beta_M^{\text{comp}}-1/2}\right).$$

Proof. For every exponent value β , define the aggregated compressed coefficient

$$A_M(\beta) := \sum_{(i,j) \in \mathcal{E}_M^{\text{ret}}: H_{ij}=\beta} c_{ij}^M.$$

Under Proposition D.3 and Corollary D.11, the Gaussian-Volterra block verifies Assumption 6.5 with these aggregated coefficients and $G_M(\tau) \rightarrow 1$. Together with the displayed scale-compatible option-remainder condition, the conclusion is Theorem 6.7 specialized to that closed-form coefficient system. Grouping channels with the same exponent produces the aggregated coefficients $A_M(\beta)$, so same-exponent cancellation is already absorbed into $\text{supp}^{\text{comp}}(\Gamma_{M,\tau}^{\mathbb{Q}})$. $\square$

Theorem D.62 (Market-scale skew visibility and screening for the perturbation-induced option-facing loading). *Under Corollary D.61, the perturbation-induced contribution carried by the option-facing priced structural loading falls into one of the following regimes for the visible summary family $\mathfrak{S}_{M,N}(\tau; \varepsilon)$, equivalently for the compressed option-facing priced structural loading*

$$\Gamma_{M,\tau}^{\mathbb{Q}} = \alpha_{M,\tau}^{\mathbb{Q}} C_{\mathcal{H}_{M,\tau}^{\text{opt}}}^{\mathbb{Q}}(\Lambda_{M,\tau})$$

whenever $\alpha_{M,\tau}^{\mathbb{Q}} > 0$:

- (i) if $\text{supp}^{\text{comp}}(\Gamma_{M,\tau}^{\mathbb{Q}}) = \emptyset$, then the present Gaussian coefficient system selects no nontrivial leading public power for the perturbation-induced loading;
- (ii) if $\beta_M^{\text{off}} < \beta_M^{\text{diag}}$ and

$$A_M^{\text{off}} = \sum_{i \neq j: H_{ij} = \beta_M^{\text{off}}} c_{ij}^M \neq 0,$$

then the perturbation-induced structural contribution is $\mathbb{Q}$ -visible at the leading scale

$$a_\tau = \tau^{\beta_M^{\text{off}} - 1/2};$$

- (iii) if $\beta_M^{\text{diag}} < \beta_M^{\text{off}}$ and $A_M^{\text{diag}} \neq 0$, then the public leading order is diagonal, and the off-diagonal contagion contribution is screened at the public scale

$$a_\tau = \tau^{\beta_M^{\text{diag}} - 1/2};$$

it can appear only at the lower-order scale $\tau^{\beta_M^{\text{off}} - 1/2}$ if the corresponding off-diagonal coefficient is nonzero;

- (iv) if $\beta_M^{\text{diag}} = \beta_M^{\text{off}} < \infty$, then leading-order visibility depends on the total coefficient on the tied exponent

$$A_M^{\text{tie}} := \sum_{(i,j) \in \mathcal{E}_M^{\text{ret}}: H_{ij} = \beta_M} c_{ij}^M.$$

Equivalently, the leading maturity power is β_M^{comp} , the smallest exponent in the aggregated compressed exponent support of $\alpha_{M,\tau}^{\mathbb{Q}} C_{\mathcal{H}_{M,\tau}^{\text{opt}}}^{\mathbb{Q}}(\Lambda_{M,\tau})$, and visibility at that power is decided by non-cancellation of the corresponding compressed coefficient. In the one-source/one-target limit, this reduces to the threshold logic of [2]. Since $\mathcal{H}_{M,\tau}^{\text{opt}}$ is trivial, this is equivalently the statement that $|\Gamma_{M,\tau}^{\mathbb{Q}}| > 0$ at the leading scale selected by the compressed priced structural loading.

Proof. If $\text{supp}^{\text{comp}}(\Gamma_{M,\tau}^{\mathbb{Q}}) = \emptyset$, this is exactly the first case. Assume therefore that the aggregated compressed support is nonempty. Corollary D.61 expresses the compressed priced structural loading $\Gamma_{M,\tau}^{\mathbb{Q}}$ as a finite sum of powers of τ grouped by exponent. The smallest active exponent in $\text{supp}^{\text{comp}}(\Gamma_{M,\tau}^{\mathbb{Q}})$, equivalently the smallest β with $A_M(\beta) \neq 0$, determines the public leading order. If the smallest active exponent is off-diagonal, visibility at that order requires non-cancellation of the aggregate coefficient A_M^{off} . If the smallest active exponent is diagonal and its aggregate coefficient is nonzero, then the public summary is dominated by diagonal channels, so the off-diagonal contagion contribution is screened at that leading scale. If channels with the same candidate leading exponent cancel, the effective public scale lifts to the next exponent in the aggregated compressed support. Restricting to a single incoming off-diagonal channel and a single diagonal benchmark recovers the threshold mechanism of [2]. $\square$

Corollary D.63 (Market-index option summaries are visible but generally incomplete). *The market-index ATM normal skew family is a legitimate visible $\mathbb{Q}$ -summary for the option-facing claim $\Xi_{M,\tau}$, but it need not span the full perturbation-induced structural contribution. In particular, under the hypotheses of Theorem D.62(iii), the option-facing structural loading $L_{M,\tau}^{\mathbb{Q}}$ is nonzero while its off-diagonal contagion contribution is screened at the public leading scale.*

Proof. Corollary D.61 gives the nontrivial option-facing priced loading, and Theorem D.62 identifies when the leading public scale is diagonal rather than off-diagonal. $\square$

APPENDIX E. BLACK AND IMPLIED-VOLATILITY CONVENTIONS

This section records that the visibility classification proved for the normal/Bachelier quote is unchanged by a smooth local passage to Black-implied-volatility conventions. The first result derives that passage directly from the option price surface by implicit differentiation; the second records the abstract conversion-map version used by derived visibility summaries. The local surface language follows the Black–Scholes/Black convention and local-volatility surface tradition [32, 33, 35].

Theorem E.1 (Black ATM skew conversion by implicit differentiation). *Assume $S_0^M > 0$. Let the same call price surface $C_M(\tau, k; \varepsilon)$ be quoted both in Bachelier volatility σ_N^M and in Black implied volatility σ_B^M , with strike $K = S_0^M + k$. Assume the baseline $\varepsilon = 0$ surface is regular near $k = 0$ and that the Black vega is nonzero. Suppose also that the perturbation is a pure skew perturbation at the money, in the sense that*

$$\partial_\varepsilon C_M(\tau, 0; \varepsilon)|_{\varepsilon=0} = 0.$$

Then

$$\partial_\varepsilon \partial_k \sigma_B^M(\tau, k; \varepsilon)|_{k=0, \varepsilon=0} = M_{B,N}(\tau) \partial_\varepsilon \partial_k \sigma_N^M(\tau, k; \varepsilon)|_{k=0, \varepsilon=0},$$

where

$$M_{B,N}(\tau) := \frac{\phi(0)}{S_0^M \phi(w_\tau/2)} = \frac{1}{S_0^M} \exp\left(\frac{w_\tau^2}{8}\right), \quad w_\tau := \sigma_B^M(\tau, 0; 0) \sqrt{\tau}.$$

Here ϕ is the standard normal density. In particular, if $w_\tau \rightarrow 0$ as $\tau \downarrow 0$, then

$$M_{B,N}(\tau) \rightarrow \frac{1}{S_0^M} > 0.$$

Thus Black quoting preserves the leading exponent support and all visibility/screening classifications carried by the Bachelier ATM skew derivative.

Proof. For fixed k , implicit differentiation of the Bachelier and Black pricing equations gives

$$\partial_\varepsilon \sigma_N^M(\tau, k; \varepsilon)|_{\varepsilon=0} = \frac{\partial_\varepsilon C_M(\tau, k; \varepsilon)|_{\varepsilon=0}}{V_N(\tau, k)}, \quad \partial_\varepsilon \sigma_B^M(\tau, k; \varepsilon)|_{\varepsilon=0} = \frac{\partial_\varepsilon C_M(\tau, k; \varepsilon)|_{\varepsilon=0}}{V_B(\tau, k)},$$

where, at the baseline,

$$V_N(\tau, k) = \sqrt{\tau} \phi\left(\frac{k}{\sigma_N^M(\tau, k; 0) \sqrt{\tau}}\right),$$

and

$$V_B(\tau, k) = S_0^M \sqrt{\tau} \phi(d_+(\tau, k)), \quad d_+(\tau, 0) = \frac{w_\tau}{2}.$$

Therefore

$$\partial_\varepsilon \sigma_B^M(\tau, k; \varepsilon)|_{\varepsilon=0} = Q_\tau(k) \partial_\varepsilon \sigma_N^M(\tau, k; \varepsilon)|_{\varepsilon=0},$$

with

$$Q_\tau(k) = \frac{V_N(\tau, k)}{V_B(\tau, k)}.$$

At $k = 0$,

$$Q_\tau(0) = \frac{\phi(0)}{S_0^M \phi(w_\tau/2)} = M_{B,N}(\tau).$$

The pure-skew condition implies

$$\partial_\varepsilon \sigma_N^M(\tau, 0; \varepsilon)|_{\varepsilon=0} = 0,$$

because the Bachelier vega is nonzero. Differentiating the relation $\partial_\varepsilon \sigma_B^M = Q_\tau \partial_\varepsilon \sigma_N^M$ in k at zero therefore removes the possible $Q'_\tau(0) \partial_\varepsilon \sigma_N^M(\tau, 0)|_{\varepsilon=0}$ term and yields the stated formula. Positivity of the Bachelier and Black vegas gives $M_{B,N}(\tau) > 0$; the assumed short-maturity limits of the quoting map give the stated positive limit. $\square$

Corollary E.2 (Black expansion for the primitive non-Gaussian Volterra class). *Consider an option price surface whose return perturbation satisfies Assumption A.6 and whose Bachelier/Black quoting maps satisfy the regularity hypotheses of Theorem E.1. Suppose the associated Bachelier quote satisfies the pure-skew condition at the money. Then, in the nonempty support case of Theorem A.8,*

$$\partial_\varepsilon \partial_k \sigma_B^M(\tau, k; \varepsilon)|_{k=0, \varepsilon=0} = M_{B,N}(\tau) \left(\sum_{\beta \in \mathcal{B}_J} A_J(\beta) \tau^{\beta-1/2} + O(\tau^{\beta_0-1/2+\eta}) \right).$$

Whenever the selected-support condition of Theorem A.8 holds, the Black expansion has the same selected exponent as the Bachelier expansion.

Proof. Corollary D.10 supplies the Bachelier slope expansion on the selected support. Theorem E.1 applies the local Black conversion with multiplier $G_B(\tau)$ and a lower-order Taylor remainder. Multiplication by $G_B(\tau)$ therefore preserves the selected exponent and multiplies the coefficient family by the displayed Black conversion factor. $\square$

Proposition E.3 (Black-implied visibility is equivalent at leading order). *Assume in addition that $S_0^M > 0$ and that the market quotes admit a local normal-to-Black ATM slope conversion map Φ_τ near the baseline, so that*

$$\mathfrak{S}_{M,B}(\tau; \varepsilon) = \Phi_\tau(\mathfrak{S}_{M,N}(\tau; \varepsilon)).$$

Assume that Φ_τ is differentiable at $\mathfrak{S}_{M,N}(\tau; 0)$ with derivative $G_M(\tau)$, where

$$G_M(\tau) \rightarrow G_M(0) > 0,$$

and that, in the perturbation regime considered here, its Taylor remainder satisfies

$$\Phi_\tau(\mathfrak{S}_{M,N}(\tau; \varepsilon)) - \Phi_\tau(\mathfrak{S}_{M,N}(\tau; 0)) - G_M(\tau)(\mathfrak{S}_{M,N}(\tau; \varepsilon) - \mathfrak{S}_{M,N}(\tau; 0)) = o(\varepsilon \tau^{\beta_M^{\text{comp}}-1/2}).$$

Assume also that

$$\text{supp}^{\text{comp}}(\Gamma_{M,\tau}^{\mathbb{Q}}) \neq \emptyset, \quad \text{so } \beta_M^{\text{comp}} < \infty.$$

Then

$$\mathfrak{S}_{M,B}(\tau; \varepsilon) - \mathfrak{S}_{M,B}(\tau; 0) = G_M(\tau) (\mathfrak{S}_{M,N}(\tau; \varepsilon) - \mathfrak{S}_{M,N}(\tau; 0)) + o(\varepsilon \tau^{\beta_M^{\text{comp}}-1/2}),$$

where $\mathfrak{S}_{M,B}$ denotes the public ATM Black-implied volatility slope summary. Consequently, Theorem D.62 remains unchanged under Black quoting conventions.

Proof. The displayed Taylor expansion for Φ_τ gives the leading-order relation between the Black and normal ATM slope increments. The limiting positivity of G_M preserves visibility, screening, and cancellation classifications. $\square$

Assumption E.4 (Regular local Black surface window). Let $\mathcal{D} \subset (0, T] \times \mathbb{R}$ be a compact maturity–moneyness window. Let $C^\varepsilon(\tau, k)$ be an arbitrage-regular price surface, and let $\sigma_B^\varepsilon(\tau, k)$ be the corresponding Black implied-volatility surface. Assume

$$C^\varepsilon = C^0 + \varepsilon \dot{C} + \rho^\varepsilon, \quad \|\rho^\varepsilon\|_{C^m(\mathcal{D})} = o(\varepsilon a_\tau),$$

for the leading scale a_τ being tracked. Assume the baseline Black surface is regular on $\mathcal{D}$:

$$V_B(\tau, k) := \partial_\sigma C_{BS}(\tau, k, \sigma_B^0(\tau, k)) > 0,$$

and the inverse Black map has uniformly bounded derivatives up to order $m+1$ on a neighborhood of $C^0(\mathcal{D})$. Equivalently, the window excludes local vega degeneracy.

Theorem E.5 (Local Black implied-volatility surface transfer). *Under Assumption E.4,*

$$\sigma_B^\varepsilon = \sigma_B^0 + \varepsilon \dot{\sigma}_B + o_{C^m(\mathcal{D})}(\varepsilon a_\tau),$$

where

$$\dot{\sigma}_B(\tau, k) = \frac{\dot{C}(\tau, k)}{V_B(\tau, k)}.$$

Consequently, for every continuous linear surface-jet functional $\ell : C^m(\mathcal{D}) \rightarrow \mathbb{R}$,

$$\ell(\sigma_B^\varepsilon - \sigma_B^0) = \varepsilon \ell\left(\frac{\dot{C}}{V_B}\right) + o(\varepsilon a_\tau).$$

If

$$\ell\left(\frac{\dot{C}}{V_B}\right) = c_\ell \tau^\beta + o(\tau^\beta), \quad c_\ell \neq 0,$$

then the Black surface jet has the same leading visibility exponent β . If this coefficient vanishes for a given ℓ , that Black jet does not select the exponent.

Proof. The Black price equation is

$$C^\varepsilon(\tau, k) = C_{BS}(\tau, k, \sigma_B^\varepsilon(\tau, k)).$$

The Banach-space implicit function theorem applies on the regular window because vega is positive and the inverse map has bounded derivatives. Differentiating the identity at $\varepsilon = 0$ gives

$$\dot{C}(\tau, k) = V_B(\tau, k) \dot{\sigma}_B(\tau, k),$$

which proves the formula for $\dot{\sigma}_B$. The remainder estimate follows from the Taylor remainder of the inverse map in $C^m(\mathcal{D})$. Applying ℓ gives the jet statement and the visibility conclusion. $\square$

Corollary E.6 (Normal-to-Black surface transfer). *Suppose, on the same regular window, the price perturbation is represented through a normal implied surface by*

$$C^\varepsilon = C_N(\sigma_N^\varepsilon), \quad \sigma_N^\varepsilon = \sigma_N^0 + \varepsilon \dot{\sigma}_N + o_{C^m(\mathcal{D})}(\varepsilon a_\tau).$$

Let

$$V_N(\tau, k) := \partial_\sigma C_N(\tau, k, \sigma_N^0(\tau, k)).$$

Then

$$\dot{\sigma}_B = \frac{V_N}{V_B} \dot{\sigma}_N.$$

Thus Black and normal visibility agree for every surface jet ℓ for which the linear operator

$$f \mapsto \ell\left(\frac{V_N}{V_B} f\right)$$

does not annihilate the leading normal-surface coefficient.

Proof. Differentiating $C_N(\sigma_N^\varepsilon)$ at $\varepsilon = 0$ gives $\dot{C} = V_N \dot{\sigma}_N$. Substitute this into Theorem E.5. $\square$

Remark E.7 (Relation to the ATM-skew proposition). Proposition E.3 is the one-dimensional jet $\ell(f) = \partial_k f(\tau, 0)$ applied to Theorem E.5, after the small-maturity vega ratio is expanded. The pure-skew hypothesis is needed only when one wants a single scalar exponent to be carried by that one jet. The surface theorem transfers the whole local jet vector on any regular pricing window.

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). *Latent Volatility Contagion in Rough Volatility Models: Spectral Thresholds and Detectability*. SSRN. <https://doi.org/10.2139/ssrn.6520718>.
- [3] J. Vidal Llauradó (2026). *Dynamic Observability of Latent Contagion: Sequential Revelation and Observable Screening*. SSRN. <https://doi.org/10.2139/ssrn.6539199>.
- [4] J. Vidal Llauradó (2026). *Detecting Latent Volatility Contagion: A Projected Score Estimator for Rough Volatility Models*. SSRN. <https://doi.org/10.2139/ssrn.6616118>.
- [5] J. M. Harrison and D. M. Kreps, *Martingales and arbitrage in multiperiod securities markets*, Journal of Economic Theory **20** (1979), no. 3, 381–408, doi:10.1016/0022-0531(79)90043-7.
- [6] 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.
- [7] 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.

- [8] F. Delbaen and W. Schachermayer, *A general version of the fundamental theorem of asset pricing*, *Mathematische Annalen* **300** (1994), 463–520, doi:10.1007/BF01450498.
- [9] I. Karatzas and S. E. Shreve, *Methods of Mathematical Finance*, Applications of Mathematics, vol. 39, Springer, 1998, doi:10.1007/978-1-4939-6845-9.
- [10] H. Föllmer and D. Sondermann, *Hedging of non-redundant contingent claims*, in W. Hildenbrand and A. Mas-Colell (eds.), *Contributions to Mathematical Economics*, North-Holland, 1986, 205–223.
- [11] 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.
- [12] M. Schweizer, *Mean-variance hedging for general claims*, *The Annals of Applied Probability* **2** (1992), no. 1, 171–179, doi:10.1214/aoap/1177005776.
- [13] M. Schweizer, *On the minimal martingale measure and the Föllmer–Schweizer decomposition*, *Stochastic Analysis and Applications* **13** (1995), no. 5, 573–599, doi:10.1080/07362999508809418.
- [14] T. Choulli, N. Vandaele, and M. Vanmaele, *The Föllmer–Schweizer decomposition: comparison and description*, *Stochastic Processes and their Applications* **120** (2010), no. 6, 853–872, doi:10.1016/j.spa.2010.02.004.
- [15] D. Williams, *Probability with Martingales*, Cambridge Mathematical Textbooks, Cambridge University Press, 1991, doi:10.1017/CBO9780511813658.
- [16] I. Karatzas and S. E. Shreve, *Brownian Motion and Stochastic Calculus*, 2nd ed., Graduate Texts in Mathematics, vol. 113, Springer, 1991, doi:10.1007/978-1-4612-0949-2.
- [17] P. Lakner, *Utility maximization with partial information*, *Stochastic Processes and their Applications* **56** (1995), no. 2, 247–273, doi:10.1016/0304-4149(94)00073-3.
- [18] D. C. Brody, L. P. Hughston, and A. Macrina, *Information-based asset pricing*, *International Journal of Theoretical and Applied Finance* **11** (2008), no. 1, 107–142, doi:10.1142/S0219024908004749.
- [19] J. Vidal Llauradó (2026). *Gaussian Seams in Risk-Neutral Compression: A Volterra Source-Family Companion*. Companion A.
- [20] C. Bayer, P. Friz, and J. Gatheral, *Pricing under rough volatility*, *Quantitative Finance* **16** (2016), no. 6, 887–904, doi:10.1080/14697688.2015.1099717.
- [21] J. Gatheral, T. Jaisson, and M. Rosenbaum, *Volatility is rough*, *Quantitative Finance* **18** (2018), no. 6, 933–949, doi:10.1080/14697688.2017.1393551.
- [22] E. Abi Jaber, M. Larsson, and S. Pulido, *Affine Volterra processes*, *Annals of Applied Probability* **29** (2019), no. 5, 3155–3200, doi:10.1214/19-AAP1477.
- [23] O. El Euch and M. Rosenbaum, *The characteristic function of rough Heston models*, *Mathematical Finance* **29** (2019), no. 1, 3–38, doi:10.1111/mafi.12173.
- [24] G. Di Nunno, K. Kubilius, Y. Mishura, and A. Yurchenko-Tytarenko, *From constant to rough: a survey of continuous volatility modeling*, *Mathematics* **11** (2023), no. 19, article 4201, doi:10.3390/math11194201.
- [25] E. Abi Jaber and S. Li, *Volatility models in practice: rough, path-dependent, or Markovian?*, *Mathematical Finance* **35** (2025), no. 4, 796–817, doi:10.1111/mafi.12463.
- [26] R. A. Horn and C. R. Johnson, *Matrix Analysis*, 2nd ed., Cambridge University Press, 2012, doi:10.1017/CBO9781139020411.
- [27] A. Berman and R. J. Plemmons, *Nonnegative Matrices in the Mathematical Sciences*, Classics in Applied Mathematics, vol. 9, SIAM, 1994, doi:10.1137/1.9781611971262.
- [28] G. Heber, A. Lunde, N. Shephard, and K. Sheppard, *Oxford-Man Institute’s realized library, version 0.1*, Oxford-Man Institute, University of Oxford, 2009, <https://realized.oxford-man.ox.ac.uk/>.
- [29] O. E. Barndorff-Nielsen, P. R. Hansen, A. Lunde, and N. Shephard, *Designing realized kernels to measure the ex post variation of equity prices in the presence of noise*, *Econometrica* **76** (2008), no. 6, 1481–1536, doi:10.3982/ECTA6495.
- [30] N. Shephard and K. Sheppard, *Realising the future: forecasting with high-frequency-based volatility (HEAVY) models*, *Journal of Applied Econometrics* **25** (2010), no. 2, 197–231, doi:10.1002/jae.1158.
- [31] L. Bachelier, *Théorie de la spéculation*, *Annales scientifiques de l’École Normale Supérieure* **17** (1900), 21–86, doi:10.24033/asens.476.
- [32] 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.
- [33] F. Black, *The pricing of commodity contracts*, *Journal of Financial Economics* **3** (1976), 167–179, doi:10.1016/0304-405X(76)90024-6.
- [34] D. T. Breeden and R. H. Litzenberger, *Prices of state-contingent claims implicit in option prices*, *The Journal of Business* **51** (1978), no. 4, 621–651, doi:10.1086/296025.
- [35] B. Dupire, *Pricing with a smile*, *Risk* **7** (1994), 18–20.
- [36] R. Jarrow and A. Rudd, *Approximate option valuation for arbitrary stochastic processes*, *Journal of Financial Economics* **10** (1982), no. 3, 347–369, doi:10.1016/0304-405X(82)90007-1.
- [37] C. J. Corrado and T. Su, *Skewness and kurtosis in S&P 500 index returns implied by option prices*, *The Journal of Financial Research* **19** (1996), no. 2, 175–192, doi:10.1111/j.1475-6803.1996.tb00592.x.
- [38] G. Bakshi, N. Kapadia, and D. Madan, *Stock return characteristics, skew laws, and the differential pricing of individual equity options*, *The Review of Financial Studies* **16** (2003), no. 1, 101–143, doi:10.1093/rfs/16.1.0101.

- [39] A. Neuberger, *The log contract*, The Journal of Portfolio Management **20** (1994), 74–80, doi:10.3905/jpm.1994.409478.
- [40] P. Carr and R. Lee, *Volatility derivatives*, Annual Review of Financial Economics **1** (2009), 319–339, doi:10.1146/annurev.financial.050808.114304.
- [41] K.-I. Sato, *Lévy Processes and Infinitely Divisible Distributions*, Cambridge Studies in Advanced Mathematics, vol. 68, Cambridge University Press, 1999.
- [42] R. Cont and P. Tankov, *Financial Modelling with Jump Processes*, Chapman & Hall/CRC, 2004, doi:10.1201/9780203485217.
- [43] W. K. Newey and K. D. West, *A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix*, Econometrica **55** (1987), no. 3, 703–708, doi:10.2307/1913610.
- [44] D. W. K. Andrews, *Heteroskedasticity and autocorrelation consistent covariance matrix estimation*, Econometrica **59** (1991), no. 3, 817–858, doi:10.2307/2938229.

Email address: `vidal.llaurado@gmail.com`