

PERRON STRESS PRICING UNDER PUBLIC COMPRESSION RESIDUAL PRICE INTERVALS, SYNCHRONIZED STRESS, AND VISIBLE HEDGING

J. VIDAL LLAURADÓ 

ABSTRACT. Public option surfaces can fix every admitted public payoff price while leaving latent market-stress claims unpriced. The missing object is the component of the priced Perron loading outside the public projection, beyond any additional surface coordinate. The paper studies this residual pricing fiber under a latent Perron stress sigma-field $\mathcal{K}$, a public sigma-field $\mathcal{H} \subseteq \mathcal{K}$, and a tradeable public subspace $\mathcal{S} \subseteq L^2(\mathbb{Q}, \mathcal{H})$. For a square-integrable claim X , the priced latent loading is $L_X^{\mathbb{Q}} = C_{\mathcal{K}}^{\mathbb{Q}}(X)$. Public calibration fixes $C_{\mathcal{H}}^{\mathbb{Q}}(L_X^{\mathbb{Q}})$. Visible hedging is a separate projection of X onto $\mathcal{S}$.

Public-equivalent pricing kernels preserve all public option prices and still move latent-stress claims along the public-annihilator fiber. This yields scalar and finite-dimensional residual price intervals, stress-charge selections, and the optimal visible-hedge decomposition. The hedge residual splits into the claim component outside the Perron stress field, the latent component not measurable with respect to public information, and the public component not spanned by tradeable instruments. Equivalent physical measures transfer the same residual into P&L-bias bounds and capital lower bounds. Strongly convex stress claims add a quadratic compression loss. A synchronized-stress component separates index-level variance exposure from the single-name carrier span: carrier enrichment can shrink residual price intervals, while tradeable carrier absorption remains a stronger condition than statistical spanning.

The observable layer contains estimands, not sample estimates. The finite objects are public projection shares, carrier-enrichment increments, residual hedge-risk floors, and the implied-versus-realized synchronization wedge. They specify what a data or listed-option implementation can measure. Contract-level execution remains outside the pricing geometry.

1. INTRODUCTION

An option surface is a calibration object before it is a spanning object. It fixes admitted public risk-neutral coordinates. It need not span the latent market-stress loading that enters a claim. A maturity slice, a local skew coefficient, or an index variance summary can be public and still leave a Perron market-variance component outside the tradeable public span. Synchronized stress is the sharp case: liquid index variance does not imply completeness of the single-name carrier span needed to absorb the synchronized component. The paper prices the residual left by this public compression.

The risk-neutral compression geometry of [1] supplies the baseline. For a square-integrable claim X , the priced latent loading is

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

where $\mathcal{K}$ is the latent Perron stress information field. Public option information $\mathcal{H} \subseteq \mathcal{K}$ fixes

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

not necessarily $L_X^{\mathbb{Q}}$. Visible hedging is governed by another map, the projection of X , or of $L_X^{\mathbb{Q}}$, onto a closed span $\mathcal{S}$ of tradeable public instruments. The three objects remain separate:

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

Date: May 2026.

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

Key words and phrases. public compression, Perron stress, synchronized stress, option-surface spanning, residual risk, visible hedging, expected shortfall, incomplete markets.

DOI: 10.2139/ssrn.6700598.

The first is latent loading. The second is public calibration. The third is visible hedging.

Public compression therefore creates a pricing fiber. Public option surfaces define equivalent pricing kernels that agree on public claims. Latent-stress claims can move inside that class when their residual loading is not spanned by public information. The corresponding object is a compression-consistent price interval. A point quote requires a separate selector: a stress premium, a state-price-density convention, or a capital charge. For synchronized-stress claims the same geometry separates the index public surface from the single-name carrier surfaces entering the index-versus-name variance decomposition. Carrier enrichment can reduce the residual interval. It becomes a stable hedge only after an additional tradeable-admissibility condition.

The hedge problem changes the space. If $\mathcal{S}$ is the closed $L^2(\mathbb{Q})$ -span of tradeable public instruments, the optimal visible hedge is the Hilbert projection $C_{\mathcal{S}}^{\mathbb{Q}}X$. Its residual decomposes into a claim component outside the latent stress field, a latent Perron residual not visible in public information, and a public component not spanned by tradeable instruments. If the Perron component retains variance outside $\mathcal{S}$, no visible hedge removes it. Under an equivalent physical measure $\mathbb{P}$, bounded density ratios transfer the $\mathbb{Q}$ -residual floor into physical hedge-risk bounds and compression-adjusted capital formulas.

The observable layer follows this projection order. A finite public shadow is obtained by projecting a fixed stress coordinate onto public option features. An index-only public span is compared with a carrier-enriched span. Residual hedge-risk floors are measured in the visible span. An implied-versus-realized synchronization wedge is formed from index and single-name total-variance coordinates. These are estimands of the theory, not latent-state recovery claims.

The latent-contagion, dynamic-observability, projected-score, market-scale, and risk-neutral compression constructions [2, 3, 4, 5, 1] isolate and verify the latent stress geometry. This paper uses that geometry to build price intervals, visible hedge decompositions, and finite observable estimands for the stress risk left outside calibrated option surfaces. A listed-option implementation starts at a different level. It treats the residual synchronized-stress wedge as a tradeable object and adds execution, liquidity, carrier absorption, and adaptedness constraints. The Hilbert geometry remains a pricing and hedging theory.

Relation to existing literature. Option-spanning theory gives the reference boundary. Classical option-pricing results show how derivative prices can complete or enlarge market contingencies [6, 7, 8, 9]. The present paper starts from the opposite case. The traded and observed option system leaves the latent stress claim outside its span. The pricing object is therefore the residual family of pricing kernels that preserve public option prices while changing latent-stress payoffs, rather than a fully specified parametric surface.

Empirical option-pricing work already separates surface fit from hedging and option-return performance. Deterministic implied-volatility functions can fit cross-sectional prices while performing poorly in hedging tests [10]; stochastic-volatility and jump models improve some pricing dimensions while leaving hedgeability as a separate property [11, 12]. Spanning tests based on option panels find priced smile factors beyond deterministic volatility functions [13]. The Hilbert formulation used here is the projection version of that boundary. Calibration fixes the public projection. Hedge error is governed by the component outside the visible span.

Option P&L and variance-risk-premium evidence supply the pricing pressure. Negative delta-hedged option gains, expected option-return patterns, and variance-risk premia show that volatility and tail risks are priced in ways not captured by static surface fit [14, 15, 16, 17, 18]. Andersen, Fusari, and Todorov [19] document a persistent wedge between market risks and their compensation in index options. Under public compression, that wedge is the separation between the visible part of stress and the residual that appears in upper-tail hedge losses and variance-pricing errors.

Implied-volatility surface dynamics and multifactor smirk models provide the finite public coordinates [20, 21]. These coordinates are used as public benchmarks, not as full latent states. The object priced below is the component of a latent Perron stress state spanned by such public coordinates, together with the residual risk that remains.

Unspanned stochastic volatility in fixed-income and commodity markets [22, 23] and incomplete-market pricing bounds such as good-deal restrictions [24] supply the incomplete-market language. The unspanned coordinate here is a Perron stress variable generated by rough Volterra contagion. The rough-volatility literature gives the temporal scale for such latent volatility states [25, 26]. The latent-contagion and compression papers derive the detectability thresholds, observable screening, projected-score construction, aggregate market Perron state, and risk-neutral compression map [2, 3, 4, 5, 1]. This paper takes that Perron object as given and studies its pricing, hedging, and capital consequences under finite public option information.

The no-arbitrage background is the martingale-measure tradition [27, 28, 29, 30]. The hedge projections used below belong to the Hilbert-space family of quadratic hedging and local-risk minimization [31, 32, 33, 34].

Main statements. Section 2 fixes the Hilbert-space setting. The public option system is a sigma-field $\mathcal{H}$, latent Perron stress information is a sigma-field $\mathcal{K}$, and visible trading instruments form a closed subspace $\mathcal{S} \subseteq L^2(\mathbb{Q}, \mathcal{H})$. Conditional expectation under a measure μ is denoted by $C_{\mathcal{A}}^{\mu}$.

Section 3 identifies the compression-consistent pricing kernels. A measure $\mathbb{Q}' \sim \mathbb{Q}$ is public-equivalent to $\mathbb{Q}$ exactly when its density $Z = d\mathbb{Q}'/d\mathbb{Q}$ satisfies

$$\mathbb{E}_{\mathbb{Q}}[Z \mid \mathcal{H}] = 1.$$

Then all $\mathcal{H}$ -measurable public payoffs have the same price under $\mathbb{Q}'$ and $\mathbb{Q}$. Theorem 3.3 constructs nontrivial local kernels by bounded latent residual tilts, while Theorem 3.7 gives the positive conditional exponential construction used for unbounded Gaussian residual factors. Theorem 2.7 gives the exact annihilator statement for public identification, and Theorem 3.10 applies it to compression-consistent pricing kernels.

Section 4 characterizes the Perron stress pricing model. The price interval generated by the public-equivalent kernel family is explicit for bounded residual tilts (Theorem 4.2), and Theorem 4.8 writes each tilted price as a public price plus a latent residual term. When the residual direction is a Perron stress coordinate, Theorem 4.11 gives the interval width in terms of the canonical Perron exposure coefficient and the unspanned public residual variance. Theorem 4.19 separates calibration-only surface intervals from dynamic-safe public-terminal intervals. Theorem 4.18 gives the model-wide support-function interval over bounded latent residual budgets, and Theorem 4.6 gives the corresponding finite-dimensional ellipsoidal interval. A selected compression price is obtained by adding a stress charge to the public price (Proposition 4.21), and Theorem 4.25 identifies the envelope case in which that charge is a linear public-equivalent endpoint.

Section 5 specializes the residual-pricing geometry to synchronized-stress claims. It defines index-versus-single-name synchronization loadings, separates index public spans from carrier-enriched public spans, and proves that carrier enrichment weakly shrinks the residual price interval. Tradeable carrier absorption remains a separate hedgeability condition.

Section 6 proves the visible-hedging results. The optimal visible hedge is $C_{\mathcal{S}}^{\mathbb{Q}}X$. If $L_X^{\mathbb{Q}} = C_{\mathcal{K}}^{\mathbb{Q}}X$, then

$$\|X - C_{\mathcal{S}}^{\mathbb{Q}}X\|_{2,\mathbb{Q}}^2 = \|X - L_X^{\mathbb{Q}}\|_{2,\mathbb{Q}}^2 + \|L_X^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}}L_X^{\mathbb{Q}}\|_{2,\mathbb{Q}}^2 + \|C_{\mathcal{H}}^{\mathbb{Q}}L_X^{\mathbb{Q}} - C_{\mathcal{S}}^{\mathbb{Q}}L_X^{\mathbb{Q}}\|_{2,\mathbb{Q}}^2.$$

The identity is Theorem 6.4. Theorem 6.3 characterizes the minimax visible hedge under residual-kernel uncertainty, and Theorem 6.6 gives a Perron coordinate lower bound whenever a latent stress factor has a nonzero residual outside the tradeable public span.

Section 7 transfers hedge residuals to the physical measure and defines compression-adjusted risk capital. It also bounds the physical drift of a $\mathbb{Q}$ -orthogonal hedge residual and gives explicit capital lower bounds. Section 8 gives a convexity inequality showing that public compression understates strongly convex stress losses by at least a quadratic residual term. Section 9 records the finite public-projection, carrier-enrichment, residual-risk, and synchronization-wedge estimands generated by the theory.

2. SETUP

2.1. Information fields and projections. Let $(\Omega, \mathcal{F}, \mathbb{Q})$ be a probability space. All random variables are identified up to $\mathbb{Q}$ -null sets unless a second measure is explicitly named. Let

$$\mathcal{H} \subseteq \mathcal{G} \subseteq \mathcal{K} \subseteq \mathcal{F}$$

be sub-sigma-fields. The fields carry the following roles.

- $\mathcal{H}$ is the public option and visible-market information used for calibration;
- $\mathcal{G}$ is the larger public terminal market information, when needed;
- $\mathcal{K}$ is the latent Perron stress information field;
- $\mathcal{F}$ is the full terminal information field of the claim.

In the market-scale construction of [5, 1], $\mathcal{K}$ is specialized to the one-factor Perron stress field

$$\mathcal{K} \equiv \mathcal{H}_\varpi := \mathcal{G}_T \vee \sigma(\varpi(t) : 0 \leq t \leq T), \quad \varpi(t) = u^\top x(t),$$

where u is the Perron direction of the screened contagion operator and x is the latent Volterra envelope. This is the scalar Perron specialization of a broader latent stress field. The field is fixed throughout the analysis. Pricing, hedging, and capital are studied after public option information has been compressed to $\mathcal{H}$.

For a sub-sigma-field $\mathcal{A} \subseteq \mathcal{F}$, let

$$C_{\mathcal{A}}^{\mathbb{Q}}(X) := \mathbb{E}_{\mathbb{Q}}[X \mid \mathcal{A}], \quad X \in L^2(\mathbb{Q}).$$

The operator is the orthogonal projection from $L^2(\mathbb{Q}, \mathcal{F})$ onto $L^2(\mathbb{Q}, \mathcal{A})$. If $\mathcal{S}$ is a closed linear subspace of $L^2(\mathbb{Q})$, $C_{\mathcal{S}}^{\mathbb{Q}}$ denotes the orthogonal projection onto $\mathcal{S}$.

Definition 2.1 (Latent loading, public shadow, and visible hedge). For $X \in L^2(\mathbb{Q})$, define

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

The random variable $L_X^{\mathbb{Q}}$ is the priced latent loading, $X_{\mathcal{H}}^{\mathbb{Q}}$ is its public calibration shadow, and $X_{\mathcal{S}}^{\mathbb{Q}}$ is the best visible $L^2(\mathbb{Q})$ -hedge in the closed tradeable span $\mathcal{S}$.

Assumption 2.2 (Visible tradeable span). The visible tradeable span $\mathcal{S}$ is a closed linear subspace of $L^2(\mathbb{Q}, \mathcal{H})$ containing the constants. Its elements are terminal gains from admissible public instruments after discounting and static normalization.

Remark 2.3. The closed subspace is the Hilbert-space abstraction underlying mean-variance hedging and local-risk-minimization arguments [31, 32, 33]. A finite implementation may take $\mathcal{S}$ to be the span of index returns, synthetic variance swaps, maturity-sliced option portfolios, or option-surface factor portfolios.

2.2. Public option systems. The construction is static at a fixed date and may be applied date by date in a panel. At a fixed date, public option information is summarized by $\mathcal{H}$ -measurable quantities

$$Y_1, \dots, Y_m \in L^\infty(\mathbb{Q}, \mathcal{H}), \quad p_i = \mathbb{E}_{\mathbb{Q}}[Y_i].$$

The Y_i are payoff or feature coordinates, while p_i are their observed discounted public prices at the calibration date. In a date-indexed panel, the observed price or feature coordinate is written $q_{d,\tau}^{(f)}$, such as implied-volatility level, total-variance level, or downside total-variance skew.

Definition 2.4 (Public calibration agreement). Two probability measures $\mathbb{Q}' \sim \mathbb{Q}$ agree on the public calibration system $\mathcal{P}_{\mathcal{H}} \subset L^1(\mathbb{Q}, \mathcal{H})$ if

$$\mathbb{E}_{\mathbb{Q}'}[Y] = \mathbb{E}_{\mathbb{Q}}[Y], \quad Y \in \mathcal{P}_{\mathcal{H}}.$$

They agree on all public payoffs if this holds for every bounded $\mathcal{H}$ -measurable Y .

Agreement on all bounded public payoffs is the mathematical primitive. A finite observed option system is its finite-dimensional restriction.

Conditional density normalization is the kernel condition for public calibration agreement.

Proposition 2.5 (Exact public-equivalence condition). *Let $Z \geq 0$ satisfy $\mathbb{E}_{\mathbb{Q}}[Z] = 1$, and define $d\mathbb{Q}^Z = Z d\mathbb{Q}$. The following are equivalent:*

$$\mathbb{E}_{\mathbb{Q}^Z}[Y] = \mathbb{E}_{\mathbb{Q}}[Y] \quad \text{for every bounded } \mathcal{H}\text{-measurable } Y,$$

and

$$\mathbb{E}_{\mathbb{Q}}[Z \mid \mathcal{H}] = 1.$$

If $Z > 0$ $\mathbb{Q}$ -a.s., then $\mathbb{Q}^Z \sim \mathbb{Q}$, giving the class of public-equivalent pricing kernels relative to $\mathcal{H}$.

Proof. If $\mathbb{E}_{\mathbb{Q}}[Z \mid \mathcal{H}] = 1$, then for every bounded $\mathcal{H}$ -measurable Y ,

$$\mathbb{E}_{\mathbb{Q}^Z}[Y] = \mathbb{E}_{\mathbb{Q}}[ZY] = \mathbb{E}_{\mathbb{Q}}[\mathbb{E}_{\mathbb{Q}}[Z \mid \mathcal{H}]Y] = \mathbb{E}_{\mathbb{Q}}[Y].$$

Conversely, if the equality of expectations holds for every bounded $\mathcal{H}$ -measurable Y , then

$$\mathbb{E}_{\mathbb{Q}}[(Z - 1)Y] = 0 \quad \text{for every bounded } \mathcal{H}\text{-measurable } Y.$$

That identity is the defining property of $\mathbb{E}_{\mathbb{Q}}[Z - 1 \mid \mathcal{H}] = 0$, and therefore $\mathbb{E}_{\mathbb{Q}}[Z \mid \mathcal{H}] = 1$. $\square$

The calibration condition above is a reduced Hilbert-feature formulation. It applies after the public option surface has been represented by $\mathcal{H}$ -measurable quoted coordinates. For terminal admitted option payoffs X_a , date- t price preservation is stronger. A density tilt Z must satisfy

$$\mathbb{E}_{\mathbb{Q}}[ZX_a \mid \mathcal{H}_t] = \mathbb{E}_{\mathbb{Q}}[X_a \mid \mathcal{H}_t]$$

for every admitted payoff. The distinction separates feature-level public calibration from payoff-module preservation.

Definition 2.6 (Public residual directions). The L^2 and bounded public residual directions are

$$\mathcal{R}_{\mathcal{H}}^2 := \{R \in L^2(\mathbb{Q}) : \mathbb{E}_{\mathbb{Q}}[R \mid \mathcal{H}] = 0\}, \quad \mathcal{R}_{\mathcal{H}}^{\infty} := \mathcal{R}_{\mathcal{H}}^2 \cap L^{\infty}(\mathbb{Q}).$$

To keep the annihilator notation uniform, write

$$\mathcal{N}_{\mathcal{H}}^{\infty} := \mathcal{R}_{\mathcal{H}}^{\infty}.$$

If $\mathcal{P}_m \subset L^2(\mathbb{Q})$ is a finite-dimensional public payoff space containing constants, its residual directions are

$$\mathcal{R}_m^2 := \{R \in L^2(\mathbb{Q}) : \mathbb{E}_{\mathbb{Q}}[RY] = 0 \text{ for every } Y \in \mathcal{P}_m\}, \quad \mathcal{R}_m^{\infty} := \mathcal{R}_m^2 \cap L^{\infty}(\mathbb{Q}),$$

and

$$\mathcal{N}_m^{\infty} := \mathcal{R}_m^{\infty}.$$

The public residual is the annihilator of the public payoff quotient.

Theorem 2.7 (Exact public residual duality). *For $X \in L^2(\mathbb{Q})$, set*

$$\Delta_{\mathcal{H}}X := X - C_{\mathcal{H}}^{\mathbb{Q}}X.$$

Then

$$\|\Delta_{\mathcal{H}}X\|_{2,\mathbb{Q}} = \sup \left\{ \mathbb{E}_{\mathbb{Q}}[XR] : R \in \mathcal{R}_{\mathcal{H}}^2, \|R\|_{2,\mathbb{Q}} \leq 1 \right\}.$$

The same value is obtained if the supremum is restricted to bounded residual directions

$$R \in \mathcal{R}_{\mathcal{H}}^{\infty}, \quad \|R\|_{2,\mathbb{Q}} \leq 1.$$

Consequently, the following are equivalent:

$$X = C_{\mathcal{H}}^{\mathbb{Q}}X \quad \text{in } L^2(\mathbb{Q}),$$

and

$$\mathbb{E}_{\mathbb{Q}}[XR] = 0 \quad \text{for every } R \in \mathcal{R}_{\mathcal{H}}^{\infty}.$$

More generally,

$$\mathbb{E}_{\mathbb{Q}}[XR] = \mathbb{E}_{\mathbb{Q}}[(\Delta_{\mathcal{H}}X)R], \quad R \in \mathcal{R}_{\mathcal{H}}^2.$$

Proof. For $R \in \mathcal{R}_{\mathcal{H}}^2$,

$$\mathbb{E}_{\mathbb{Q}}[(C_{\mathcal{H}}^{\mathbb{Q}}X)R] = \mathbb{E}_{\mathbb{Q}}[(C_{\mathcal{H}}^{\mathbb{Q}}X)\mathbb{E}_{\mathbb{Q}}[R \mid \mathcal{H}]] = 0.$$

Thus $\mathbb{E}_{\mathbb{Q}}[XR] = \mathbb{E}_{\mathbb{Q}}[(\Delta_{\mathcal{H}}X)R]$. By Cauchy–Schwarz,

$$\mathbb{E}_{\mathbb{Q}}[XR] \leq |\mathbb{E}_{\mathbb{Q}}[(\Delta_{\mathcal{H}}X)R]| \leq \|\Delta_{\mathcal{H}}X\|_{2,\mathbb{Q}} \|R\|_{2,\mathbb{Q}},$$

so the supremum over L^2 -residual directions is at most $\|\Delta_{\mathcal{H}}X\|_{2,\mathbb{Q}}$. If $\Delta_{\mathcal{H}}X \neq 0$, the choice

$$R^* = \frac{\Delta_{\mathcal{H}}X}{\|\Delta_{\mathcal{H}}X\|_{2,\mathbb{Q}}}$$

belongs to $\mathcal{R}_{\mathcal{H}}^2$, has unit L^2 -norm, and attains the upper bound. If $\Delta_{\mathcal{H}}X = 0$, both sides are zero.

Bounded residual directions must give the same support value. Let $\Delta = \Delta_{\mathcal{H}}X$, and suppose $\Delta \neq 0$. For $n \geq 1$, let

$$T_n := (-n) \vee (\Delta \wedge n), \quad R_n := T_n - \mathbb{E}_{\mathbb{Q}}[T_n \mid \mathcal{H}].$$

Then $R_n \in L^\infty(\mathbb{Q})$ and $\mathbb{E}_{\mathbb{Q}}[R_n \mid \mathcal{H}] = 0$, so $R_n \in \mathcal{R}_{\mathcal{H}}^\infty$. Since $T_n \rightarrow \Delta$ in $L^2(\mathbb{Q})$ and conditional expectation is an L^2 -contraction,

$$\mathbb{E}_{\mathbb{Q}}[T_n \mid \mathcal{H}] \rightarrow \mathbb{E}_{\mathbb{Q}}[\Delta \mid \mathcal{H}] = 0 \quad \text{in } L^2(\mathbb{Q}).$$

Thus $R_n \rightarrow \Delta$ in $L^2(\mathbb{Q})$, and $\|R_n\|_{2,\mathbb{Q}} \rightarrow \|\Delta\|_{2,\mathbb{Q}}$. For all large n , $R_n \neq 0$. With $\tilde{R}_n = R_n / \|R_n\|_{2,\mathbb{Q}}$,

$$\mathbb{E}_{\mathbb{Q}}[X\tilde{R}_n] = \frac{\mathbb{E}_{\mathbb{Q}}[\Delta R_n]}{\|R_n\|_{2,\mathbb{Q}}} \rightarrow \|\Delta\|_{2,\mathbb{Q}}.$$

This proves equality for bounded residual directions. The annihilator equivalence follows by applying the norm identity. $\square$

Finite calibration replaces the sigma-field quotient by a payoff-span quotient.

Corollary 2.8 (Finite-coordinate residual duality). *Let $\mathcal{P}_m \subset L^2(\mathbb{Q})$ be finite dimensional and contain constants, and let $C_m^{\mathbb{Q}}$ be the $L^2(\mathbb{Q})$ -projection onto $\mathcal{P}_m$. For every $X \in L^2(\mathbb{Q})$,*

$$\|X - C_m^{\mathbb{Q}}X\|_{2,\mathbb{Q}} = \sup \left\{ \mathbb{E}_{\mathbb{Q}}[XR] : R \in \mathcal{R}_m^2, \|R\|_{2,\mathbb{Q}} \leq 1 \right\}.$$

If $\mathcal{P}_m$ is spanned by bounded payoffs, the same value is obtained with $R \in \mathcal{R}_m^\infty$.

Proof. For $R \in \mathcal{R}_m^2$, the projection $C_m^{\mathbb{Q}}X$ belongs to $\mathcal{P}_m$, and therefore

$$\mathbb{E}_{\mathbb{Q}}[(C_m^{\mathbb{Q}}X)R] = 0.$$

Thus

$$\mathbb{E}_{\mathbb{Q}}[XR] = \mathbb{E}_{\mathbb{Q}}[(X - C_m^{\mathbb{Q}}X)R].$$

Cauchy–Schwarz gives the upper bound by $\|X - C_m^{\mathbb{Q}}X\|_{2,\mathbb{Q}}$. If $X - C_m^{\mathbb{Q}}X \neq 0$, the normalized residual

$$R^* = \frac{X - C_m^{\mathbb{Q}}X}{\|X - C_m^{\mathbb{Q}}X\|_{2,\mathbb{Q}}}$$

belongs to $\mathcal{R}_m^2$ and attains that bound. If $X = C_m^{\mathbb{Q}}X$, both sides are zero.

Assume now that $\mathcal{P}_m$ is spanned by bounded payoffs. Put $\Delta = X - C_m^{\mathbb{Q}}X$. If $\Delta \neq 0$, define

$$T_n = (-n) \vee (\Delta \wedge n), \quad R_n = T_n - C_m^{\mathbb{Q}}T_n.$$

The finite-dimensional projection $C_m^{\mathbb{Q}}T_n$ is a finite linear combination of bounded basis payoffs, hence $R_n \in L^\infty(\mathbb{Q})$. Also $R_n \perp \mathcal{P}_m$, so $R_n \in \mathcal{R}_m^\infty$. Since $T_n \rightarrow \Delta$ in $L^2(\mathbb{Q})$ and $C_m^{\mathbb{Q}}\Delta = 0$, one has $R_n \rightarrow \Delta$ in $L^2(\mathbb{Q})$. Normalizing the nonzero R_n and passing to the limit gives the same supremum over $\mathcal{R}_m^\infty$. $\square$

Nested fields separate the surface residual from the terminal-public residual.

Theorem 2.9 (Nested surface and public-terminal residuals). *Let $\mathcal{H} \subseteq \mathcal{G} \subseteq \mathcal{K}$, let $X \in L^2(\mathbb{Q})$, and set*

$$L := C_{\mathcal{K}}^{\mathbb{Q}} X.$$

Then

$$L - C_{\mathcal{H}}^{\mathbb{Q}} L = (L - C_{\mathcal{G}}^{\mathbb{Q}} L) + (C_{\mathcal{G}}^{\mathbb{Q}} L - C_{\mathcal{H}}^{\mathbb{Q}} L),$$

with orthogonal summands. Consequently,

$$\|L - C_{\mathcal{H}}^{\mathbb{Q}} L\|_{2,\mathbb{Q}}^2 = \|L - C_{\mathcal{G}}^{\mathbb{Q}} L\|_{2,\mathbb{Q}}^2 + \|C_{\mathcal{G}}^{\mathbb{Q}} L - C_{\mathcal{H}}^{\mathbb{Q}} L\|_{2,\mathbb{Q}}^2.$$

In addition,

$$\|L - C_{\mathcal{G}}^{\mathbb{Q}} L\|_{2,\mathbb{Q}} = \sup \left\{ \mathbb{E}_{\mathbb{Q}}[XR] : R \in L^2(\mathbb{Q}, \mathcal{K}), \mathbb{E}_{\mathbb{Q}}[R | \mathcal{G}] = 0, \|R\|_{2,\mathbb{Q}} \leq 1 \right\}.$$

Proof. The decomposition is the Pythagorean identity for the nested closed subspaces

$$L^2(\mathbb{Q}, \mathcal{H}) \subseteq L^2(\mathbb{Q}, \mathcal{G}) \subseteq L^2(\mathbb{Q}, \mathcal{K}).$$

The first summand is orthogonal to $L^2(\mathbb{Q}, \mathcal{G})$; the second belongs to $L^2(\mathbb{Q}, \mathcal{G})$ and is orthogonal to $L^2(\mathbb{Q}, \mathcal{H})$. This proves the squared-norm identity.

The dual formula is Theorem 2.7 applied on $(\Omega, \mathcal{K}, \mathbb{Q})$ with the sub-sigma-field $\mathcal{G}$. Since R is $\mathcal{K}$ -measurable,

$$\mathbb{E}_{\mathbb{Q}}[XR] = \mathbb{E}_{\mathbb{Q}}[LR].$$

□

The corresponding kernel restrictions separate surface-safe price freedom from dynamic-safe terminal-public price freedom.

Corollary 2.10 (Surface-safe and dynamic-safe price freedom). *Let $X \in L^2(\mathbb{Q})$ and set $L = C_{\mathcal{K}}^{\mathbb{Q}} X$. A density satisfying $\mathbb{E}_{\mathbb{Q}}[Z | \mathcal{H}] = 1$ preserves the calibrated surface field and can move the price contribution from the full surface residual $L - C_{\mathcal{H}}^{\mathbb{Q}} L$. A density satisfying the stronger condition $\mathbb{E}_{\mathbb{Q}}[Z | \mathcal{G}] = 1$ preserves every bounded terminal public payoff and can move only the price contribution from the public-terminal residual $L - C_{\mathcal{G}}^{\mathbb{Q}} L$. The component*

$$C_{\mathcal{G}}^{\mathbb{Q}} L - C_{\mathcal{H}}^{\mathbb{Q}} L$$

is therefore invisible to the calibrated surface field $\mathcal{H}$, but visible once the larger terminal public field $\mathcal{G}$ is admitted.

Proof. The tower property gives $\mathbb{E}_{\mathbb{Q}}[Z | \mathcal{H}] = \mathbb{E}_{\mathbb{Q}}[\mathbb{E}_{\mathbb{Q}}[Z | \mathcal{G}] | \mathcal{H}] = 1$ whenever $\mathbb{E}_{\mathbb{Q}}[Z | \mathcal{G}] = 1$. Hence $\mathcal{G}$ -safe densities preserve the calibrated surface and, by Proposition 2.5 applied to $\mathcal{G}$, preserve every bounded $\mathcal{G}$ -measurable terminal public payoff.

Let $R \in L^2(\mathbb{Q}, \mathcal{K})$ be a residual direction. If $\mathbb{E}_{\mathbb{Q}}[R | \mathcal{G}] = 0$, then

$$\mathbb{E}_{\mathbb{Q}}[XR] = \mathbb{E}_{\mathbb{Q}}[(L - C_{\mathcal{G}}^{\mathbb{Q}} L)R],$$

so any $\mathcal{G}$ -safe local tilt charges only the public-terminal residual. If only $\mathbb{E}_{\mathbb{Q}}[R | \mathcal{H}] = 0$ is imposed, the corresponding identity is

$$\mathbb{E}_{\mathbb{Q}}[XR] = \mathbb{E}_{\mathbb{Q}}[(L - C_{\mathcal{H}}^{\mathbb{Q}} L)R],$$

and the full surface residual is available. The orthogonal decomposition in Theorem 2.9 identifies $C_{\mathcal{G}}^{\mathbb{Q}} L - C_{\mathcal{H}}^{\mathbb{Q}} L$ as the part that is $\mathcal{H}$ -invisible but already $\mathcal{G}$ -measurable. Bounded residual directions approximate the corresponding nonzero components by Theorem 2.7 on the chosen conditioning field. □

3. COMPRESSION-CONSISTENT PRICING KERNELS

Definition 3.1 (Public-equivalent pricing kernel). Let $\mathbb{Q}' \sim \mathbb{Q}$ and set $Z = d\mathbb{Q}'/d\mathbb{Q}$. The measure $\mathbb{Q}'$ is *public-equivalent* to $\mathbb{Q}$ relative to $\mathcal{H}$ if

$$Z > 0, \quad \mathbb{E}_{\mathbb{Q}}[Z] = 1, \quad \mathbb{E}_{\mathbb{Q}}[Z \mid \mathcal{H}] = 1.$$

The family of such measures is denoted by $\mathcal{Q}_{\mathcal{H}}(\mathbb{Q})$. The corresponding density class is

$$\mathcal{Z}_{\mathcal{H}} := \{Z > 0 : \mathbb{E}_{\mathbb{Q}}[Z] = 1, \mathbb{E}_{\mathbb{Q}}[Z \mid \mathcal{H}] = 1\}.$$

The dynamic-safe public-terminal density class is

$$\mathcal{Z}_{\mathcal{G}} := \{Z > 0 : \mathbb{E}_{\mathbb{Q}}[Z] = 1, \mathbb{E}_{\mathbb{Q}}[Z \mid \mathcal{G}] = 1\}.$$

When $\mathcal{H} \subseteq \mathcal{G}$, the tower property gives

$$\mathcal{Z}_{\mathcal{G}} \subseteq \mathcal{Z}_{\mathcal{H}}.$$

Every dynamic-safe public-terminal kernel therefore preserves the calibrated surface, but a surface-equivalent kernel need not preserve all terminal public payoffs. If $\mathcal{P}_m$ is a finite public payoff space containing constants, its finite-calibration density class is

$$\mathcal{Z}_m := \{Z > 0 : \mathbb{E}_{\mathbb{Q}}[Z] = 1, \mathbb{E}_{\mathbb{Q}}[ZY] = \mathbb{E}_{\mathbb{Q}}[Y] \text{ for every } Y \in \mathcal{P}_m\}.$$

Then $\mathcal{Z}_{\mathcal{H}} \subseteq \mathcal{Z}_m$ whenever $\mathcal{P}_m \subset L^1(\mathbb{Q}, \mathcal{H})$, and the inclusion is generally strict for finite m .

Conditional normalization is the invariance condition for public prices.

Lemma 3.2 (Public prices are preserved). *If $\mathbb{Q}' \in \mathcal{Q}_{\mathcal{H}}(\mathbb{Q})$, then*

$$\mathbb{E}_{\mathbb{Q}'}[Y] = \mathbb{E}_{\mathbb{Q}}[Y]$$

for every bounded $\mathcal{H}$ -measurable Y . If $X \in L^1(\mathbb{Q}')$ and Y is bounded and $\mathcal{H}$ -measurable, then

$$\mathbb{E}_{\mathbb{Q}'}[XY] = \mathbb{E}_{\mathbb{Q}}[\mathbb{E}_{\mathbb{Q}}[ZX \mid \mathcal{H}]Y].$$

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

$$\mathbb{E}_{\mathbb{Q}'}[Y] = \mathbb{E}_{\mathbb{Q}}[ZY] = \mathbb{E}_{\mathbb{Q}}[\mathbb{E}_{\mathbb{Q}}[Z \mid \mathcal{H}]Y] = \mathbb{E}_{\mathbb{Q}}[Y].$$

For the mixed identity, $ZX \in L^1(\mathbb{Q})$ by the assumption $X \in L^1(\mathbb{Q}')$. Since Y is bounded and $\mathcal{H}$ -measurable,

$$\mathbb{E}_{\mathbb{Q}'}[XY] = \mathbb{E}_{\mathbb{Q}}[ZXY] = \mathbb{E}_{\mathbb{Q}}[\mathbb{E}_{\mathbb{Q}}[ZX \mid \mathcal{H}]Y],$$

which is the displayed formula. $\square$

Bounded residual directions generate local fibers of public-equivalent kernels.

Theorem 3.3 (Bounded local residual tilts). *Let $R \in L^\infty(\mathbb{Q})$ satisfy*

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

For any $\eta \in \mathbb{R}$ with $|\eta| \|R\|_\infty < 1$, define

$$Z_\eta := 1 + \eta R, \quad \frac{d\mathbb{Q}_\eta}{d\mathbb{Q}} := Z_\eta.$$

Then $\mathbb{Q}_\eta \in \mathcal{Q}_{\mathcal{H}}(\mathbb{Q})$. In addition, for every $X \in L^1(\mathbb{Q})$,

$$\mathbb{E}_{\mathbb{Q}_\eta}[X] = \mathbb{E}_{\mathbb{Q}}[X] + \eta \mathbb{E}_{\mathbb{Q}}[XR].$$

Proof. The bound on η gives $Z_\eta > 0$. Since $\mathbb{E}_{\mathbb{Q}}[R] = \mathbb{E}_{\mathbb{Q}}[\mathbb{E}_{\mathbb{Q}}[R \mid \mathcal{H}]] = 0$, one has $\mathbb{E}_{\mathbb{Q}}[Z_\eta] = 1$. Also

$$\mathbb{E}_{\mathbb{Q}}[Z_\eta \mid \mathcal{H}] = 1 + \eta \mathbb{E}_{\mathbb{Q}}[R \mid \mathcal{H}] = 1.$$

Thus $\mathbb{Q}_\eta \in \mathcal{Q}_{\mathcal{H}}(\mathbb{Q})$. Finally,

$$\mathbb{E}_{\mathbb{Q}_\eta}[X] = \mathbb{E}_{\mathbb{Q}}[Z_\eta X] = \mathbb{E}_{\mathbb{Q}}[X] + \eta \mathbb{E}_{\mathbb{Q}}[XR].$$

$\square$

Remark 3.4. The bounded linear tilt is a local residual-kernel device. It gives explicit intervals and the tangent calculus in Theorem 3.6. Unbounded Gaussian Perron factors use the positive conditional exponential kernels of Theorem 3.7.

Definition 3.5 (Public-equivalent tangent spaces). The public-equivalent tangent space at $\mathbb{Q}$ is

$$T_{\mathcal{H}} := \{R \in L^2(\mathbb{Q}) : \mathbb{E}_{\mathbb{Q}}[R \mid \mathcal{H}] = 0\}.$$

If $\mathcal{H} \subseteq \mathcal{K}$, the latent public-equivalent tangent space is

$$T_{\mathcal{H}, \mathcal{K}} := \{R \in L^2(\mathbb{Q}, \mathcal{K}) : \mathbb{E}_{\mathbb{Q}}[R \mid \mathcal{H}] = 0\}.$$

The tangent space of the public-equivalent kernel fiber is the public residual space.

Theorem 3.6 (Tangent characterization of public-equivalent kernels). *Let $(Z_{\varepsilon})_{\varepsilon > 0}$ be public-equivalent densities, so that*

$$Z_{\varepsilon} > 0, \quad \mathbb{E}_{\mathbb{Q}}[Z_{\varepsilon} \mid \mathcal{H}] = 1.$$

Assume

$$\frac{Z_{\varepsilon} - 1}{\varepsilon} \longrightarrow R \quad \text{in } L^2(\mathbb{Q}).$$

Then $R \in T_{\mathcal{H}}$. Conversely, if $R \in L^{\infty}(\mathbb{Q}) \cap T_{\mathcal{H}}$, then

$$Z_{\varepsilon} := \frac{\exp(\varepsilon R)}{\mathbb{E}_{\mathbb{Q}}[\exp(\varepsilon R) \mid \mathcal{H}]}$$

is public-equivalent and

$$\frac{Z_{\varepsilon} - 1}{\varepsilon} \longrightarrow R \quad \text{in } L^2(\mathbb{Q}).$$

If, in addition, $R \in T_{\mathcal{H}, \mathcal{K}}$ and $X \in L^2(\mathbb{Q})$, then

$$\frac{d}{d\varepsilon} \Big|_{\varepsilon=0} \mathbb{E}_{\mathbb{Q}}[Z_{\varepsilon} X] = \mathbb{E}_{\mathbb{Q}}[(C_{\mathcal{K}}^{\mathbb{Q}} X - C_{\mathcal{H}}^{\mathbb{Q}} C_{\mathcal{K}}^{\mathbb{Q}} X) R].$$

Proof. Because $\mathbb{E}_{\mathbb{Q}}[Z_{\varepsilon} \mid \mathcal{H}] = 1$,

$$\mathbb{E}_{\mathbb{Q}} \left[\frac{Z_{\varepsilon} - 1}{\varepsilon} \mid \mathcal{H} \right] = 0.$$

Conditional expectation is an L^2 -contraction, so passing to the L^2 -limit gives $\mathbb{E}_{\mathbb{Q}}[R \mid \mathcal{H}] = 0$.

For the converse, boundedness gives a uniform expansion

$$\exp(\varepsilon R) = 1 + \varepsilon R + \varepsilon^2 A_{\varepsilon}, \quad \sup_{|\varepsilon| \leq \varepsilon_0} \|A_{\varepsilon}\|_{\infty} < \infty$$

for every fixed small ε_0 . Since $\mathbb{E}_{\mathbb{Q}}[R \mid \mathcal{H}] = 0$,

$$D_{\varepsilon} := \mathbb{E}_{\mathbb{Q}}[\exp(\varepsilon R) \mid \mathcal{H}] = 1 + \varepsilon^2 \mathbb{E}_{\mathbb{Q}}[A_{\varepsilon} \mid \mathcal{H}].$$

For small $|\varepsilon|$, D_{ε} is bounded away from zero and $D_{\varepsilon}^{-1} = 1 + O(\varepsilon^2)$ in L^{∞} . Hence

$$Z_{\varepsilon} = \frac{1 + \varepsilon R + O(\varepsilon^2)}{1 + O(\varepsilon^2)} = 1 + \varepsilon R + O(\varepsilon^2) \quad \text{in } L^2(\mathbb{Q}),$$

so $(Z_{\varepsilon} - 1)/\varepsilon \rightarrow R$ in $L^2(\mathbb{Q})$. The conditional normalization gives $\mathbb{E}_{\mathbb{Q}}[Z_{\varepsilon} \mid \mathcal{H}] = 1$.

Finally,

$$\frac{\mathbb{E}_{\mathbb{Q}}[Z_{\varepsilon} X] - \mathbb{E}_{\mathbb{Q}}[X]}{\varepsilon} \rightarrow \mathbb{E}_{\mathbb{Q}}[XR].$$

If R is $\mathcal{K}$ -measurable and conditionally mean zero relative to $\mathcal{H}$, then

$$\mathbb{E}_{\mathbb{Q}}[XR] = \mathbb{E}_{\mathbb{Q}}[(C_{\mathcal{K}}^{\mathbb{Q}} X) R] = \mathbb{E}_{\mathbb{Q}}[(C_{\mathcal{K}}^{\mathbb{Q}} X - C_{\mathcal{H}}^{\mathbb{Q}} C_{\mathcal{K}}^{\mathbb{Q}} X) R].$$

□

Conditional exponential normalization gives positive kernels for unbounded residual factors.

Theorem 3.7 (Conditional exponential public-equivalent kernel). *Let $R \in L^0(\mathbb{Q}, \mathcal{K}; \mathbb{R}^d)$, let θ be a bounded $\mathcal{H}$ -measurable $\mathbb{R}^d$ -valued parameter, and assume*

$$0 < \mathbb{E}_{\mathbb{Q}}[\exp(\theta^\top R) \mid \mathcal{H}] < \infty \quad \mathbb{Q}\text{-a.s.}$$

Define

$$Z_\theta := \frac{\exp(\theta^\top R)}{\mathbb{E}_{\mathbb{Q}}[\exp(\theta^\top R) \mid \mathcal{H}]}.$$

Then $Z_\theta > 0$, $\mathbb{E}_{\mathbb{Q}}[Z_\theta \mid \mathcal{H}] = 1$, and every bounded public payoff price is preserved. If $X \in L^1(\mathbb{Q})$ and $XZ_\theta \in L^1(\mathbb{Q})$, then

$$\mathbb{E}_{\mathbb{Q}}[Z_\theta X] = \mathbb{E}_{\mathbb{Q}}[Z_\theta C_{\mathcal{K}}^{\mathbb{Q}} X].$$

If differentiation under the expectation is valid on an interval of parameters, then, with $Y = C_{\mathcal{K}}^{\mathbb{Q}} X$,

$$\left. \frac{d}{dt} \mathbb{E}_{\mathbb{Q}}[Z_{\theta+th} X] \right|_{t=0} = \mathbb{E}_{\mathbb{Q}}[Z_\theta Y h^\top \{R - \mathbb{E}_{\mathbb{Q}}[Z_\theta R \mid \mathcal{H}]\}]$$

for every bounded $\mathcal{H}$ -measurable direction h . In particular, at $\theta = 0$,

$$\left. \frac{d}{dt} \right|_{t=0} \mathbb{E}_{\mathbb{Q}}[Z_{th} X] = \mathbb{E}_{\mathbb{Q}}[h^\top \text{Cov}_{\mathbb{Q}}(Y, R \mid \mathcal{H})].$$

A sufficient condition for the displayed derivative in a fixed bounded direction h is the existence of $\delta > 0$ and an integrable random variable G such that, for $|t| \leq \delta$,

$$|Y| \exp((\theta + th)^\top R) (1 + |h^\top R|) \leq G \quad \mathbb{Q}\text{-a.s.},$$

with the corresponding conditional normalizing denominators finite and strictly positive.

Proof. The conditional normalization gives $\mathbb{E}_{\mathbb{Q}}[Z_\theta \mid \mathcal{H}] = 1$, and positivity is built into the exponential form. Proposition 2.5 then gives public payoff preservation. Since Z_θ is $\mathcal{K}$ -measurable,

$$\mathbb{E}_{\mathbb{Q}}[Z_\theta X] = \mathbb{E}_{\mathbb{Q}}[Z_\theta \mathbb{E}_{\mathbb{Q}}[X \mid \mathcal{K}]] = \mathbb{E}_{\mathbb{Q}}[Z_\theta C_{\mathcal{K}}^{\mathbb{Q}} X].$$

For the derivative, work under the stated differentiability hypothesis; the displayed domination condition is one sufficient way to obtain it. Set

$$D(t) := \mathbb{E}_{\mathbb{Q}}[\exp((\theta + th)^\top R) \mid \mathcal{H}].$$

Then

$$D'(0) = \mathbb{E}_{\mathbb{Q}}[h^\top R \exp(\theta^\top R) \mid \mathcal{H}],$$

and the conditional denominator remains finite and strictly positive. The logarithmic derivative is therefore

$$\left. \frac{d}{dt} \log Z_{\theta+th} \right|_{t=0} = h^\top R - \frac{D'(0)}{D(0)} = h^\top \{R - \mathbb{E}_{\mathbb{Q}}[Z_\theta R \mid \mathcal{H}]\}.$$

Dominated convergence justifies applying this derivative inside $\mathbb{E}_{\mathbb{Q}}[Z_{\theta+th} Y]$. Multiplication by $Z_\theta Y$ gives the displayed derivative formula. At $\theta = 0$, $Z_0 = 1$, and the derivative is

$$\mathbb{E}_{\mathbb{Q}}[h^\top Y (R - C_{\mathcal{H}}^{\mathbb{Q}} R)].$$

Because h is $\mathcal{H}$ -measurable, subtracting $C_{\mathcal{H}}^{\mathbb{Q}} Y$ does not change this moment. Thus

$$\mathbb{E}_{\mathbb{Q}}[h^\top Y (R - C_{\mathcal{H}}^{\mathbb{Q}} R)] = \mathbb{E}_{\mathbb{Q}}[h^\top (Y - C_{\mathcal{H}}^{\mathbb{Q}} Y) (R - C_{\mathcal{H}}^{\mathbb{Q}} R)],$$

which is the displayed conditional covariance expression. $\square$

Conditional Gaussianity turns the residual selector into a covariance charge.

Corollary 3.8 (Exact conditional-Gaussian Perron pricing formula). *Let $Y = C_{\mathcal{K}}^{\mathbb{Q}} X$ and let $R \in L^2(\mathbb{Q}, \mathcal{K}; \mathbb{R}^d)$. Suppose that, conditionally on $\mathcal{H}$, the vector (Y, R) is jointly Gaussian and has conditional covariance vector*

$$c_{\mathcal{H}}(X; R) := \text{Cov}_{\mathbb{Q}}(Y, R \mid \mathcal{H}) \in L^1(\mathbb{Q}, \mathcal{H}; \mathbb{R}^d).$$

For every bounded $\mathcal{H}$ -measurable $\theta \in \mathbb{R}^d$, the exponential residual kernel Z_θ satisfies

$$\mathbb{E}_{\mathbb{Q}}[Z_\theta X \mid \mathcal{H}] = C_{\mathcal{H}}^{\mathbb{Q}} Y + \theta^\top c_{\mathcal{H}}(X; R),$$

and therefore

$$\mathbb{E}_{\mathbb{Q}}[Z_{\theta}X] = \mathbb{E}_{\mathbb{Q}}[X] + \mathbb{E}_{\mathbb{Q}}[\theta^{\top} c_{\mathcal{H}}(X; R)].$$

If A is symmetric positive definite and the selector is constrained pointwise by $\theta^{\top} A \theta \leq \rho^2$, then the conditional-Gaussian residual interval is

$$\mathbb{E}_{\mathbb{Q}}[X] \pm \rho \mathbb{E}_{\mathbb{Q}}\left[\left(c_{\mathcal{H}}(X; R)^{\top} A^{-1} c_{\mathcal{H}}(X; R)\right)^{1/2}\right].$$

If θ is deterministic rather than $\mathcal{H}$ -measurable, the corresponding interval is

$$\mathbb{E}_{\mathbb{Q}}[X] \pm \rho (\bar{c}^{\top} A^{-1} \bar{c})^{1/2}, \quad \bar{c} := \mathbb{E}_{\mathbb{Q}}[c_{\mathcal{H}}(X; R)].$$

Proof. Conditionally on $\mathcal{H}$, the Gaussian moment-generating identity gives

$$\frac{\mathbb{E}_{\mathbb{Q}}[Y \exp(\theta^{\top} R) \mid \mathcal{H}]}{\mathbb{E}_{\mathbb{Q}}[\exp(\theta^{\top} R) \mid \mathcal{H}]} = C_{\mathcal{H}}^{\mathbb{Q}} Y + \theta^{\top} \text{Cov}_{\mathbb{Q}}(Y, R \mid \mathcal{H}).$$

The conditional normalizer in Z_{θ} is exactly the denominator in this identity. Hence

$$\mathbb{E}_{\mathbb{Q}}[Z_{\theta}Y \mid \mathcal{H}] = C_{\mathcal{H}}^{\mathbb{Q}} Y + \theta^{\top} c_{\mathcal{H}}(X; R).$$

Since Z_{θ} is $\mathcal{K}$ -measurable, X may be replaced by Y . Taking expectations gives the price formula. The support function of the ellipsoid is

$$\sup_{\theta^{\top} A \theta \leq \rho^2} \theta^{\top} c = \rho (c^{\top} A^{-1} c)^{1/2}.$$

Applying this identity pointwise to the $\mathcal{H}$ -measurable vector $c_{\mathcal{H}}(X; R)$ gives the stated interval. A measurable maximizer is obtained by taking

$$\theta = \rho \frac{A^{-1} c_{\mathcal{H}}(X; R)}{(c_{\mathcal{H}}(X; R)^{\top} A^{-1} c_{\mathcal{H}}(X; R))^{1/2}}$$

on the set where the denominator is positive and $\theta = 0$ otherwise; the opposite sign gives the lower endpoint. If θ is deterministic, the selector cannot vary across the public fiber, and the same support function is applied only to $\bar{c} = \mathbb{E}_{\mathbb{Q}}[c_{\mathcal{H}}(X; R)]$. $\square$

Public-terminal normalization imposes the same residual selection after conditioning on $\mathcal{G}$.

Corollary 3.9 (Public-terminal exponential kernels). *Let $\mathcal{H} \subseteq \mathcal{G} \subseteq \mathcal{K}$, let $R \in L^0(\mathbb{Q}, \mathcal{K}; \mathbb{R}^d)$, and let ϑ be a bounded $\mathcal{G}$ -measurable $\mathbb{R}^d$ -valued parameter. Assume*

$$0 < \mathbb{E}_{\mathbb{Q}}[\exp(\vartheta^{\top} R) \mid \mathcal{G}] < \infty \quad \mathbb{Q}\text{-a.s.}$$

Define

$$Z_{\vartheta}^{\mathcal{G}} := \frac{\exp(\vartheta^{\top} R)}{\mathbb{E}_{\mathbb{Q}}[\exp(\vartheta^{\top} R) \mid \mathcal{G}]}.$$

Then

$$\mathbb{E}_{\mathbb{Q}}[Z_{\vartheta}^{\mathcal{G}} \mid \mathcal{G}] = 1, \quad Z_{\vartheta}^{\mathcal{G}} \in \mathcal{Z}_{\mathcal{G}} \subseteq \mathcal{Z}_{\mathcal{H}}.$$

Hence $Z_{\vartheta}^{\mathcal{G}}$ preserves every bounded terminal public payoff and, in particular, every bounded calibrated surface payoff. If $(C_{\mathcal{K}}^{\mathbb{Q}} X, R)$ is conditionally Gaussian given $\mathcal{G}$, the conditional-Gaussian pricing formula of Corollary 3.8 holds with $\mathcal{H}$ replaced by $\mathcal{G}$.

Proof. The conditional normalization gives $\mathbb{E}_{\mathbb{Q}}[Z_{\vartheta}^{\mathcal{G}} \mid \mathcal{G}] = 1$, so $Z_{\vartheta}^{\mathcal{G}} \in \mathcal{Z}_{\mathcal{G}}$. The inclusion $\mathcal{Z}_{\mathcal{G}} \subseteq \mathcal{Z}_{\mathcal{H}}$ follows from the tower property. Payoff preservation follows from Proposition 2.5, first relative to $\mathcal{G}$ and then relative to $\mathcal{H}$. The Gaussian formula is exactly Corollary 3.8 applied with the conditioning field $\mathcal{G}$. $\square$

Public identification stops at the $L^2(\mathbb{Q}, \mathcal{H})$ quotient.

Theorem 3.10 (Public identification and latent non-identification). *Let $X \in L^2(\mathbb{Q})$. If all bounded $\mathcal{H}$ -measurable payoff tests against X are known, then they identify $C_{\mathcal{H}}^{\mathbb{Q}}X$ and identify no component of $X - C_{\mathcal{H}}^{\mathbb{Q}}X$. Equivalently, the public residual component is exactly the component that can be moved by conditionally mean-zero bounded tilts.*

For the kernel statement, let $R \in \mathcal{N}_{\mathcal{H}}^{\infty}$ and let $\mathbb{Q}_{\eta}$ be the tilt of Theorem 3.3. Then every bounded public payoff has the same price under all $\mathbb{Q}_{\eta}$, while

$$\mathbb{E}_{\mathbb{Q}_{\eta}}[X] = \mathbb{E}_{\mathbb{Q}}[X] + \eta \mathbb{E}_{\mathbb{Q}}[(X - C_{\mathcal{H}}^{\mathbb{Q}}X)R].$$

If $X \neq C_{\mathcal{H}}^{\mathbb{Q}}X$, there exist bounded public-equivalent tilts for which the price of X changes. In particular, public option information identifies compressed loadings such as $C_{\mathcal{H}}^{\mathbb{Q}}(L_X^{\mathbb{Q}})$, not the full latent loading $L_X^{\mathbb{Q}}$, unless an additional selection rule is imposed.

Proof. The Riesz representation theorem gives uniqueness of $C_{\mathcal{H}}^{\mathbb{Q}}X$ from the pairings

$$Y \mapsto \mathbb{E}_{\mathbb{Q}}[XY], \quad Y \in L^2(\mathbb{Q}, \mathcal{H}),$$

because

$$\mathbb{E}_{\mathbb{Q}}[XY] = \mathbb{E}_{\mathbb{Q}}[(C_{\mathcal{H}}^{\mathbb{Q}}X)Y] \quad \text{for } Y \in L^2(\mathbb{Q}, \mathcal{H}).$$

Bounded public tests are dense enough for the same identification. If $Y \in L^2(\mathbb{Q}, \mathcal{H})$, the truncations $Y_n = (-n) \vee (Y \wedge n)$ are bounded and $\mathcal{H}$ -measurable, and $Y_n \rightarrow Y$ in $L^2(\mathbb{Q})$. The pairings converge by Cauchy–Schwarz.

For the tilt statement, Lemma 3.2 preserves every bounded public payoff. The affine pricing identity of Theorem 3.3 and Theorem 2.7 give

$$\mathbb{E}_{\mathbb{Q}_{\eta}}[X] = \mathbb{E}_{\mathbb{Q}}[X] + \eta \mathbb{E}_{\mathbb{Q}}[XR] = \mathbb{E}_{\mathbb{Q}}[X] + \eta \mathbb{E}_{\mathbb{Q}}[(X - C_{\mathcal{H}}^{\mathbb{Q}}X)R].$$

If $X - C_{\mathcal{H}}^{\mathbb{Q}}X \neq 0$, Theorem 2.7 supplies $R_n \in \mathcal{N}_{\mathcal{H}}^{\infty}$ with $\mathbb{E}_{\mathbb{Q}}[XR_n] \neq 0$ for some n . For sufficiently small nonzero η , $1 + \eta R_n > 0$, and the resulting public-equivalent tilt changes the price of X . $\square$

Remark 3.11. The theorem is static. It does not assert that every tilt is generated by a particular continuous-time trading model. Its role is the pricing-kernel statement for public non-identification. If the market constrains only public $\mathcal{H}$ -measurable payoffs, latent residual directions remain free unless an additional selector is imposed.

Finite public calibration replaces the sigma-field quotient by the observed payoff span.

Theorem 3.12 (Finite public calibration and residual freedom). *Let $\mathcal{P}_m = \text{span}\{1, Y_1, \dots, Y_m\}$ be a finite-dimensional subspace of $L^{\infty}(\mathbb{Q}, \mathcal{H})$, and let $C_m^{\mathbb{Q}}$ be the $L^2(\mathbb{Q})$ -projection onto $\mathcal{P}_m$. For $X \in L^2(\mathbb{Q})$,*

$$X = C_m^{\mathbb{Q}}X \iff \mathbb{E}_{\mathbb{Q}}[XR] = 0 \text{ for every } R \in \mathcal{N}_m^{\infty}.$$

If $X - C_m^{\mathbb{Q}}X \neq 0$, then there is $R \in \mathcal{N}_m^{\infty}$ and $\eta_0 > 0$ such that, for $|\eta| < \eta_0$,

$$Z_{\eta} = 1 + \eta R$$

defines an equivalent measure that preserves the prices of every payoff in $\mathcal{P}_m$ but changes the price of X .

Proof. For $R \in \mathcal{N}_m^{\infty}$,

$$\mathbb{E}_{\mathbb{Q}}[(C_m^{\mathbb{Q}}X)R] = 0,$$

because $C_m^{\mathbb{Q}}X \in \mathcal{P}_m$. Hence

$$\mathbb{E}_{\mathbb{Q}}[XR] = \mathbb{E}_{\mathbb{Q}}[(X - C_m^{\mathbb{Q}}X)R].$$

If $X = C_m^{\mathbb{Q}}X$, all pairings vanish.

Conversely, set $\Delta := X - C_m^{\mathbb{Q}}X$. If $\Delta \neq 0$, define

$$T_n := (-n) \vee (\Delta \wedge n), \quad R_n := T_n - C_m^{\mathbb{Q}}T_n.$$

Since $\mathcal{P}_m$ is spanned by bounded functions, $C_m^{\mathbb{Q}}T_n$ is a finite linear combination of bounded functions; hence $R_n \in L^\infty(\mathbb{Q})$. It is orthogonal to every $Y \in \mathcal{P}_m$, so $R_n \in \mathcal{N}_m^\infty$. The projections $C_m^{\mathbb{Q}}T_n \rightarrow C_m^{\mathbb{Q}}\Delta = 0$ in $L^2(\mathbb{Q})$, while $T_n \rightarrow \Delta$ in $L^2(\mathbb{Q})$. Therefore $R_n \rightarrow \Delta$ in $L^2(\mathbb{Q})$, and

$$\mathbb{E}_{\mathbb{Q}}[XR_n] = \mathbb{E}_{\mathbb{Q}}[\Delta R_n] \rightarrow \|\Delta\|_{2,\mathbb{Q}}^2 > 0.$$

Choose n such that $\mathbb{E}_{\mathbb{Q}}[XR_n] \neq 0$, put $R = R_n$, and take $\eta_0 = \|R\|_\infty^{-1}$ if $R \neq 0$. Then $Z_\eta > 0$ for $|\eta| < \eta_0$, $\mathbb{E}_{\mathbb{Q}}[Z_\eta] = 1$ because constants lie in $\mathcal{P}_m$, and

$$\mathbb{E}_{\mathbb{Q}}[Z_\eta Y] = \mathbb{E}_{\mathbb{Q}}[Y] \quad \text{for every } Y \in \mathcal{P}_m.$$

The price of X changes linearly by $\eta \mathbb{E}_{\mathbb{Q}}[XR]$. □

4. THE PERRON STRESS PRICING MODEL

4.1. Price intervals. Let $R \in L^\infty(\mathbb{Q})$ be a public-invisible residual, meaning

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

For $0 < \bar{\eta}$ with $\bar{\eta} \|R\|_\infty < 1$, define

$$\mathcal{Q}_{\mathcal{H}}(\mathbb{Q}; R, \bar{\eta}) := \{\mathbb{Q}_\eta : |\eta| \leq \bar{\eta}\}.$$

A bounded residual direction generates a one-dimensional pricing fiber.

Definition 4.1 (Compression-consistent price interval). For $X \in L^1(\mathbb{Q})$, define

$$\Pi_{R,\bar{\eta}}^-(X) := \inf_{\mathbb{Q}' \in \mathcal{Q}_{\mathcal{H}}(\mathbb{Q}; R, \bar{\eta})} \mathbb{E}_{\mathbb{Q}'}[X], \quad \Pi_{R,\bar{\eta}}^+(X) := \sup_{\mathbb{Q}' \in \mathcal{Q}_{\mathcal{H}}(\mathbb{Q}; R, \bar{\eta})} \mathbb{E}_{\mathbb{Q}'}[X].$$

The scalar residual fiber has an exact support interval.

Theorem 4.2 (Explicit residual price interval). For $X \in L^1(\mathbb{Q})$,

$$\Pi_{R,\bar{\eta}}^-(X) = \mathbb{E}_{\mathbb{Q}}[X] - \bar{\eta} |\mathbb{E}_{\mathbb{Q}}[XR]|, \quad \Pi_{R,\bar{\eta}}^+(X) = \mathbb{E}_{\mathbb{Q}}[X] + \bar{\eta} |\mathbb{E}_{\mathbb{Q}}[XR]|.$$

The interval is nontrivial if and only if $\mathbb{E}_{\mathbb{Q}}[XR] \neq 0$.

Proof. By Theorem 3.3,

$$\mathbb{E}_{\mathbb{Q}_\eta}[X] = \mathbb{E}_{\mathbb{Q}}[X] + \eta \mathbb{E}_{\mathbb{Q}}[XR], \quad |\eta| \leq \bar{\eta}.$$

The infimum and supremum over a compact interval of a linear function are attained at the endpoints. If $\mathbb{E}_{\mathbb{Q}}[XR] \geq 0$, the infimum is attained at $\eta = -\bar{\eta}$ and the supremum at $\eta = \bar{\eta}$. If $\mathbb{E}_{\mathbb{Q}}[XR] < 0$, the endpoint roles are reversed. These two cases give the displayed formulas. □

Remark 4.3. A finite-dimensional public option calibration system gives the same construction after replacing $\mathbb{E}_{\mathbb{Q}}[R \mid \mathcal{H}] = 0$ by finitely many moment restrictions $\mathbb{E}_{\mathbb{Q}}[RY_i] = 0$, as in Theorem 3.12. The sigma-field version is stronger: it preserves every bounded public payoff and keeps the latent residual directions separate from the calibrated surface coordinates.

4.2. Finite-dimensional residual budgets. The scalar interval treats one residual direction. Several latent residual directions are kept by pricing over a bounded set of residual state-price perturbations.

Finite residual blocks require a positivity boundary for linear density tilts.

Lemma 4.4 (Ellipsoidal positivity bound). Let $R = (R_1, \dots, R_d)^\top$ with $R_i \in L^\infty(\mathbb{Q})$, and let A be symmetric positive definite. Define

$$B_A := \left\| (R^\top A^{-1} R)^{1/2} \right\|_\infty.$$

Then

$$\sup_{\vartheta^\top A \vartheta \leq \rho^2} \left\| \vartheta^\top R \right\|_\infty \leq \rho B_A.$$

In particular, if $\rho B_A < 1$, then $1 + \vartheta^\top R > 0$ for every $\vartheta^\top A \vartheta \leq \rho^2$.

Proof. For each ω ,

$$|\vartheta^\top R(\omega)| = |(A^{1/2}\vartheta)^\top (A^{-1/2}R(\omega))| \leq (\vartheta^\top A\vartheta)^{1/2} (R(\omega)^\top A^{-1}R(\omega))^{1/2}.$$

Taking the supremum over ω and then over the ellipsoid gives the bound. If $\rho B_A < 1$, every admissible ϑ satisfies $\|\vartheta^\top R\|_\infty < 1$, so $1 + \vartheta^\top R$ is strictly positive $\mathbb{Q}$ -a.s. $\square$

Definition 4.5 (Ellipsoidal residual kernel set). Let $R_1, \dots, R_d \in L^\infty(\mathbb{Q})$ satisfy

$$\mathbb{E}_{\mathbb{Q}}[R_i \mid \mathcal{H}] = 0, \quad i = 1, \dots, d.$$

Write $R = (R_1, \dots, R_d)^\top$. Let A be a symmetric positive definite $d \times d$ matrix and let $\rho > 0$. Assume

$$\sup_{\vartheta^\top A\vartheta \leq \rho^2} \|\vartheta^\top R\|_\infty < 1.$$

The ellipsoidal public-equivalent residual kernel set is

$$\mathcal{Q}_{\mathcal{H}}^{A,\rho}(\mathbb{Q}; R) := \left\{ \mathbb{Q}_\vartheta : \frac{d\mathbb{Q}_\vartheta}{d\mathbb{Q}} = 1 + \vartheta^\top R, \vartheta^\top A\vartheta \leq \rho^2 \right\}.$$

The ellipsoidal residual block prices through its support function.

Theorem 4.6 (Ellipsoidal residual price interval). *Under Definition 4.5, let $X \in L^1(\mathbb{Q})$ and define*

$$b_X := (\mathbb{E}_{\mathbb{Q}}[XR_1], \dots, \mathbb{E}_{\mathbb{Q}}[XR_d])^\top.$$

Then every measure in $\mathcal{Q}_{\mathcal{H}}^{A,\rho}(\mathbb{Q}; R)$ preserves bounded public payoff prices, and

$$\begin{aligned} \inf_{\mathbb{Q}' \in \mathcal{Q}_{\mathcal{H}}^{A,\rho}(\mathbb{Q}; R)} \mathbb{E}_{\mathbb{Q}'}[X] &= \mathbb{E}_{\mathbb{Q}}[X] - \rho (b_X^\top A^{-1}b_X)^{1/2}, \\ \sup_{\mathbb{Q}' \in \mathcal{Q}_{\mathcal{H}}^{A,\rho}(\mathbb{Q}; R)} \mathbb{E}_{\mathbb{Q}'}[X] &= \mathbb{E}_{\mathbb{Q}}[X] + \rho (b_X^\top A^{-1}b_X)^{1/2}. \end{aligned}$$

The interval collapses exactly when $b_X = 0$.

Proof. For each admissible ϑ , the density $1 + \vartheta^\top R$ is strictly positive by the assumption in Definition 4.5. Since constants are public and $\mathbb{E}_{\mathbb{Q}}[R_i \mid \mathcal{H}] = 0$,

$$\mathbb{E}_{\mathbb{Q}}[1 + \vartheta^\top R] = 1, \quad \mathbb{E}_{\mathbb{Q}}[1 + \vartheta^\top R \mid \mathcal{H}] = 1.$$

Thus $\mathbb{Q}_\vartheta \in \mathcal{Q}_{\mathcal{H}}(\mathbb{Q})$, and Lemma 3.2 gives preservation of bounded public payoff prices. For $X \in L^1(\mathbb{Q})$,

$$\mathbb{E}_{\mathbb{Q}_\vartheta}[X] = \mathbb{E}_{\mathbb{Q}}[X] + \vartheta^\top b_X.$$

It remains to optimize a linear functional over the ellipsoid $\{\vartheta : \vartheta^\top A\vartheta \leq \rho^2\}$. Write

$$\vartheta^\top b_X = (A^{1/2}\vartheta)^\top (A^{-1/2}b_X).$$

By Cauchy–Schwarz,

$$|\vartheta^\top b_X| \leq (\vartheta^\top A\vartheta)^{1/2} (b_X^\top A^{-1}b_X)^{1/2} \leq \rho (b_X^\top A^{-1}b_X)^{1/2}.$$

If $b_X \neq 0$, equality is attained by

$$\vartheta_\pm = \pm \rho \frac{A^{-1}b_X}{(b_X^\top A^{-1}b_X)^{1/2}}.$$

If $b_X = 0$, the linear term is zero for every admissible ϑ . The displayed infimum and supremum follow. $\square$

4.3. Pricing decomposition. The pricing identity separates public shadow, baseline scalar price, and residual contribution.

Definition 4.7 (Public shadow, baseline price, and residual price). For $X \in L^2(\mathbb{Q})$, define its public shadow and baseline scalar price by

$$S_X^{\text{pub}, \mathbb{Q}} := C_{\mathcal{H}}^{\mathbb{Q}}(L_X^{\mathbb{Q}}), \quad \Pi_0(X) := \mathbb{E}_{\mathbb{Q}}[X] = \mathbb{E}_{\mathbb{Q}}[S_X^{\text{pub}, \mathbb{Q}}].$$

If $R \in L^\infty(\mathbb{Q}, \mathcal{K})$ satisfies $\mathbb{E}_{\mathbb{Q}}[R \mid \mathcal{H}] = 0$, define the residual price contribution under the tilt parameter η by

$$\Pi_R^{\text{res}}(X; \eta) := \eta \left\langle L_X^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}} L_X^{\mathbb{Q}}, R \right\rangle_{L^2(\mathbb{Q})}.$$

The decomposition separates calibrated baseline price from residual selection.

Theorem 4.8 (Perron stress pricing decomposition). *Let $R \in L^\infty(\mathbb{Q}, \mathcal{K})$ satisfy $\mathbb{E}_{\mathbb{Q}}[R \mid \mathcal{H}] = 0$, and assume $|\eta| \|R\|_\infty < 1$. For every $X \in L^2(\mathbb{Q})$,*

$$\mathbb{E}_{\mathbb{Q}_\eta}[X] = \Pi_0(X) + \Pi_R^{\text{res}}(X; \eta).$$

Once a baseline calibrated measure $\mathbb{Q}$ has been selected, every public-equivalent tilt decomposes into the same baseline public term plus a latent residual term. The residual term is zero for every admissible R if and only if the latent loading $L_X^{\mathbb{Q}}$ is already public.

Proof. Since $\mathcal{H} \subseteq \mathcal{K}$,

$$\Pi_0(X) = \mathbb{E}_{\mathbb{Q}}[S_X^{\text{pub}, \mathbb{Q}}] = \mathbb{E}_{\mathbb{Q}}[L_X^{\mathbb{Q}}] = \mathbb{E}_{\mathbb{Q}}[X].$$

Theorem 3.3 gives

$$\mathbb{E}_{\mathbb{Q}_\eta}[X] = \mathbb{E}_{\mathbb{Q}}[X] + \eta \mathbb{E}_{\mathbb{Q}}[XR].$$

Because R is $\mathcal{K}$ -measurable,

$$\mathbb{E}_{\mathbb{Q}}[XR] = \mathbb{E}_{\mathbb{Q}}[L_X^{\mathbb{Q}} R].$$

Because $C_{\mathcal{H}}^{\mathbb{Q}} L_X^{\mathbb{Q}}$ is $\mathcal{H}$ -measurable and $\mathbb{E}_{\mathbb{Q}}[R \mid \mathcal{H}] = 0$,

$$\mathbb{E}_{\mathbb{Q}}[(C_{\mathcal{H}}^{\mathbb{Q}} L_X^{\mathbb{Q}})R] = 0.$$

Combining these identities gives

$$\mathbb{E}_{\mathbb{Q}_\eta}[X] = \Pi_0(X) + \eta \left\langle L_X^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}} L_X^{\mathbb{Q}}, R \right\rangle_{L^2(\mathbb{Q})}.$$

If $L_X^{\mathbb{Q}} = C_{\mathcal{H}}^{\mathbb{Q}} L_X^{\mathbb{Q}}$, the residual term is zero for every R . Conversely, if the residual term vanishes for every bounded $\mathcal{K}$ -measurable R with $\mathbb{E}_{\mathbb{Q}}[R \mid \mathcal{H}] = 0$, apply Theorem 2.7 on the probability space $(\Omega, \mathcal{K}, \mathbb{Q})$ to the random variable $L_X^{\mathbb{Q}}$. This yields $L_X^{\mathbb{Q}} = C_{\mathcal{H}}^{\mathbb{Q}} L_X^{\mathbb{Q}}$. $\square$

A $\mathcal{K}$ -measurable public-equivalent density leaves the latent loading fixed and moves only its public projection under the new measure.

Theorem 4.9 ($\mathcal{K}$ -measurable public-equivalent tilts). *Let $Z > 0$ satisfy $\mathbb{E}_{\mathbb{Q}}[Z] = 1$, $Z \in L^\infty(\mathbb{Q}, \mathcal{K})$, and*

$$\mathbb{E}_{\mathbb{Q}}[Z \mid \mathcal{H}] = 1.$$

Define $d\mathbb{Q}^Z = Z d\mathbb{Q}$. Then bounded public payoff prices are preserved. In addition, for every X such that the displayed conditional expectations are well defined,

$$C_{\mathcal{K}}^{\mathbb{Q}^Z} X = C_{\mathcal{K}}^{\mathbb{Q}} X,$$

and

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

In particular, if $Z_\eta = 1 + \eta R$, where

$$R \in L^\infty(\mathbb{Q}, \mathcal{K}), \quad \mathbb{E}_{\mathbb{Q}}[R \mid \mathcal{H}] = 0, \quad |\eta| \|R\|_\infty < 1,$$

then

$$C_{\mathcal{H}}^{\mathbb{Q}_\eta}(L_X^{\mathbb{Q}}) = C_{\mathcal{H}}^{\mathbb{Q}}(L_X^{\mathbb{Q}}) + \eta \mathbb{E}_{\mathbb{Q}}[(L_X^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}} L_X^{\mathbb{Q}})R \mid \mathcal{H}].$$

Proof. Public payoff preservation follows from Proposition 2.5. Since Z is $\mathcal{K}$ -measurable, Bayes' formula gives

$$C_{\mathcal{K}}^{\mathbb{Q}^Z} X = \frac{\mathbb{E}_{\mathbb{Q}}[ZX \mid \mathcal{K}]}{\mathbb{E}_{\mathbb{Q}}[Z \mid \mathcal{K}]} = \frac{Z C_{\mathcal{K}}^{\mathbb{Q}} X}{Z} = C_{\mathcal{K}}^{\mathbb{Q}} X.$$

The same Bayes formula at the public level gives, for $Y = C_{\mathcal{K}}^{\mathbb{Q}} X$,

$$C_{\mathcal{H}}^{\mathbb{Q}^Z} Y = \frac{\mathbb{E}_{\mathbb{Q}}[ZY \mid \mathcal{H}]}{\mathbb{E}_{\mathbb{Q}}[Z \mid \mathcal{H}]} = \mathbb{E}_{\mathbb{Q}}[ZY \mid \mathcal{H}],$$

because $\mathbb{E}_{\mathbb{Q}}[Z \mid \mathcal{H}] = 1$. For $Z_{\eta} = 1 + \eta R$, this becomes

$$C_{\mathcal{H}}^{\mathbb{Q}_{\eta}}(L_X^{\mathbb{Q}}) = C_{\mathcal{H}}^{\mathbb{Q}}(L_X^{\mathbb{Q}}) + \eta \mathbb{E}_{\mathbb{Q}}[L_X^{\mathbb{Q}} R \mid \mathcal{H}].$$

Since $C_{\mathcal{H}}^{\mathbb{Q}} L_X^{\mathbb{Q}}$ is $\mathcal{H}$ -measurable and $\mathbb{E}_{\mathbb{Q}}[R \mid \mathcal{H}] = 0$,

$$\mathbb{E}_{\mathbb{Q}}[L_X^{\mathbb{Q}} R \mid \mathcal{H}] = \mathbb{E}_{\mathbb{Q}}[(L_X^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}} L_X^{\mathbb{Q}}) R \mid \mathcal{H}].$$

□

4.4. Perron residual price intervals.

Definition 4.10 (Canonical public-residual Perron coefficient). Let $Z \in L^2(\mathbb{Q}, \mathcal{K})$, and define its public residual

$$Z_{\mathcal{H}}^{\perp} := Z - C_{\mathcal{H}}^{\mathbb{Q}} Z.$$

If $Z_{\mathcal{H}}^{\perp} \neq 0$, define the public-residual Perron coefficient of $X \in L^2(\mathbb{Q})$ by

$$a_{\mathcal{H}}(X; Z) := \frac{\langle L_X^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}} L_X^{\mathbb{Q}}, Z_{\mathcal{H}}^{\perp} \rangle_{L^2(\mathbb{Q})}}{\|Z_{\mathcal{H}}^{\perp}\|_{2, \mathbb{Q}}^2}.$$

This coefficient is defined for every square-integrable claim; no claim-specific decomposition is part of the definition.

A single Perron residual factor reduces the interval to its canonical exposure coefficient.

Theorem 4.11 (Explicit Perron residual interval). Assume $Z_{\mathcal{H}}^{\perp} \in L^{\infty}(\mathbb{Q})$, $Z_{\mathcal{H}}^{\perp} \neq 0$, and let $X \in L^2(\mathbb{Q})$. If

$$0 < \bar{\eta} \left\| Z_{\mathcal{H}}^{\perp} \right\|_{\infty} < 1,$$

then the public-equivalent kernels

$$\frac{d\mathbb{Q}_{\eta}}{d\mathbb{Q}} = 1 + \eta Z_{\mathcal{H}}^{\perp}, \quad |\eta| \leq \bar{\eta},$$

generate the price interval

$$\mathbb{E}_{\mathbb{Q}}[X] - \bar{\eta} |a_{\mathcal{H}}(X; Z)| \left\| Z_{\mathcal{H}}^{\perp} \right\|_{2, \mathbb{Q}}^2 \leq \mathbb{E}_{\mathbb{Q}_{\eta}}[X] \leq \mathbb{E}_{\mathbb{Q}}[X] + \bar{\eta} |a_{\mathcal{H}}(X; Z)| \left\| Z_{\mathcal{H}}^{\perp} \right\|_{2, \mathbb{Q}}^2.$$

The endpoints are attained at $\eta = \pm \bar{\eta}$, with the sign chosen according to $a_{\mathcal{H}}(X; Z)$. The interval is nontrivial if and only if $a_{\mathcal{H}}(X; Z) \neq 0$.

Proof. Since $Z_{\mathcal{H}}^{\perp} = Z - C_{\mathcal{H}}^{\mathbb{Q}} Z$, it satisfies $\mathbb{E}_{\mathbb{Q}}[Z_{\mathcal{H}}^{\perp} \mid \mathcal{H}] = 0$. The boundedness assumption and the restriction on $\bar{\eta}$ allow Theorem 3.3 to be applied with $R = Z_{\mathcal{H}}^{\perp}$. Hence

$$\mathbb{E}_{\mathbb{Q}_{\eta}}[X] = \mathbb{E}_{\mathbb{Q}}[X] + \eta \mathbb{E}_{\mathbb{Q}}[X Z_{\mathcal{H}}^{\perp}].$$

Because $Z_{\mathcal{H}}^{\perp}$ is $\mathcal{K}$ -measurable,

$$\mathbb{E}_{\mathbb{Q}}[X Z_{\mathcal{H}}^{\perp}] = \mathbb{E}_{\mathbb{Q}}[L_X^{\mathbb{Q}} Z_{\mathcal{H}}^{\perp}].$$

The public projection $C_{\mathcal{H}}^{\mathbb{Q}} L_X^{\mathbb{Q}}$ is orthogonal to $Z_{\mathcal{H}}^{\perp}$, so

$$\mathbb{E}_{\mathbb{Q}}[X Z_{\mathcal{H}}^{\perp}] = \langle L_X^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}} L_X^{\mathbb{Q}}, Z_{\mathcal{H}}^{\perp} \rangle_{L^2(\mathbb{Q})} = a_{\mathcal{H}}(X; Z) \left\| Z_{\mathcal{H}}^{\perp} \right\|_{2, \mathbb{Q}}^2.$$

Maximizing and minimizing the resulting affine function of η over $[-\bar{\eta}, \bar{\eta}]$ gives the interval and the endpoint statement. □

Remark 4.12. If one happens to have a decomposition

$$L_X^{\mathbb{Q}} = Y_X + a_X Z + B_X, \quad Y_X \in L^2(\mathbb{Q}, \mathcal{H}), \quad \langle B_X - C_{\mathcal{H}}^{\mathbb{Q}} B_X, Z_{\mathcal{H}}^{\perp} \rangle = 0,$$

then $a_{\mathcal{H}}(X; Z) = a_X$. The theorem instead uses the canonical projection coefficient, so this representation is optional.

Definition 4.13 (Normalized public-residual Perron factor). Let $P \in L^2(\mathbb{Q}, \mathcal{K})$ be a Perron stress coordinate and suppose

$$P_{\mathcal{H}}^{\perp} := P - C_{\mathcal{H}}^{\mathbb{Q}} P \neq 0.$$

Define

$$\hat{P}_{\mathcal{H}} := \frac{P_{\mathcal{H}}^{\perp}}{\|P_{\mathcal{H}}^{\perp}\|_{2, \mathbb{Q}}}, \quad \alpha_{\mathcal{H}}(X; P) := \langle C_{\mathcal{K}}^{\mathbb{Q}} X - C_{\mathcal{H}}^{\mathbb{Q}} C_{\mathcal{K}}^{\mathbb{Q}} X, \hat{P}_{\mathcal{H}} \rangle_{L^2(\mathbb{Q})}.$$

Normalization removes the scale of the Perron coordinate from the scalar interval.

Theorem 4.14 (Scale-free scalar Perron interval). Assume $\hat{P}_{\mathcal{H}} \in L^{\infty}(\mathbb{Q})$ and

$$0 < \rho \|\hat{P}_{\mathcal{H}}\|_{\infty} < 1.$$

The public-equivalent densities

$$Z_{\eta} = 1 + \eta \hat{P}_{\mathcal{H}}, \quad |\eta| \leq \rho,$$

generate the interval

$$\mathbb{E}_{\mathbb{Q}}[X] \pm \rho |\alpha_{\mathcal{H}}(X; P)|.$$

The interval is invariant under replacing P by cP for any nonzero scalar c , up to the sign convention for $\hat{P}_{\mathcal{H}}$ when $c < 0$.

Proof. The factor $\hat{P}_{\mathcal{H}}$ has zero conditional mean relative to $\mathcal{H}$ and unit L^2 -norm. The boundedness condition gives positivity of Z_{η} . Hence

$$\mathbb{E}_{\mathbb{Q}}[Z_{\eta} X] = \mathbb{E}_{\mathbb{Q}}[X] + \eta \langle C_{\mathcal{K}}^{\mathbb{Q}} X - C_{\mathcal{H}}^{\mathbb{Q}} C_{\mathcal{K}}^{\mathbb{Q}} X, \hat{P}_{\mathcal{H}} \rangle_{L^2(\mathbb{Q})}.$$

Optimizing over $|\eta| \leq \rho$ gives the interval. Rescaling P by $c \neq 0$ cancels in the normalization, except for a sign change when $c < 0$, and the absolute value removes that sign from the interval width. $\square$

Definition 4.15 (Whitened Perron block). Let $P = (P_1, \dots, P_r)^{\top} \in L^2(\mathbb{Q}, \mathcal{K}; \mathbb{R}^r)$, and set

$$P_{\mathcal{H}}^{\perp} := P - C_{\mathcal{H}}^{\mathbb{Q}} P, \quad G_{\mathcal{H}}(P) := \mathbb{E}_{\mathbb{Q}}[P_{\mathcal{H}}^{\perp} (P_{\mathcal{H}}^{\perp})^{\top}].$$

Assume $G_{\mathcal{H}}(P)$ is positive definite. Define the whitened residual block

$$W_{\mathcal{H}} := G_{\mathcal{H}}(P)^{-1/2} P_{\mathcal{H}}^{\perp}$$

and the whitened exposure vector

$$\alpha_{\mathcal{H}}^{(r)}(X; P) := \mathbb{E}_{\mathbb{Q}}[W_{\mathcal{H}} (C_{\mathcal{K}}^{\mathbb{Q}} X - C_{\mathcal{H}}^{\mathbb{Q}} C_{\mathcal{K}}^{\mathbb{Q}} X)].$$

Whitening makes the vector exposure independent of the chosen Perron basis.

Proposition 4.16 (Coordinate-invariant vector exposure). If B is an invertible $r \times r$ matrix and $\tilde{P} = BP$, then

$$\|\alpha_{\mathcal{H}}^{(r)}(X; \tilde{P})\| = \|\alpha_{\mathcal{H}}^{(r)}(X; P)\|.$$

The norm $\|\cdot\|$ is the Euclidean norm on $\mathbb{R}^r$. Consequently, a spherical residual budget in whitened coordinates reports an interval width that is independent of the chosen basis for the Perron block.

Proof. Let $R = P - C_{\mathcal{H}}^{\mathbb{Q}}P$. Then the residual block of $\tilde{P}$ is BR , and its Gram matrix is $BG_{\mathcal{H}}(P)B^{\top}$. Hence

$$W_{\mathcal{H}}(\tilde{P}) = (BG_{\mathcal{H}}(P)B^{\top})^{-1/2}BR = OW_{\mathcal{H}}(P), \quad O := (BG_{\mathcal{H}}(P)B^{\top})^{-1/2}BG_{\mathcal{H}}(P)^{1/2}.$$

The matrix O is orthogonal because

$$OO^{\top} = (BG_{\mathcal{H}}(P)B^{\top})^{-1/2}BG_{\mathcal{H}}(P)B^{\top}(BG_{\mathcal{H}}(P)B^{\top})^{-1/2} = I_r.$$

The exposure vector is therefore multiplied by O , and its Euclidean norm is unchanged. $\square$

Definition 4.17 (Model-wide residual budget). For $B < \infty$, define the bounded public-invisible latent residual budget

$$\mathcal{R}_{\mathcal{H},\mathcal{K}}(B) := \left\{ R \in L^{\infty}(\mathbb{Q}, \mathcal{K}) : \mathbb{E}_{\mathbb{Q}}[R \mid \mathcal{H}] = 0, \|R\|_{2,\mathbb{Q}} \leq 1, \|R\|_{\infty} \leq B \right\}.$$

For $0 < \rho$ with $\rho B < 1$, define

$$\mathcal{Z}_{\rho,B} := \{1 + \eta R : R \in \mathcal{R}_{\mathcal{H},\mathcal{K}}(B), |\eta| \leq \rho\},$$

and the residual support function

$$s_B(X) := \sup_{R \in \mathcal{R}_{\mathcal{H},\mathcal{K}}(B)} \mathbb{E}_{\mathbb{Q}}[XR].$$

The bounded residual budget yields the support interval for the latent fiber.

Theorem 4.18 (Residual-budget pricing interval). *Let $X \in L^2(\mathbb{Q})$, and assume $B < \infty$, $\rho > 0$, and $\rho B < 1$. Every $Z \in \mathcal{Z}_{\rho,B}$ is a public-equivalent pricing density. In addition,*

$$\inf_{Z \in \mathcal{Z}_{\rho,B}} \mathbb{E}_{\mathbb{Q}}[ZX] = \mathbb{E}_{\mathbb{Q}}[X] - \rho s_B(X), \quad \sup_{Z \in \mathcal{Z}_{\rho,B}} \mathbb{E}_{\mathbb{Q}}[ZX] = \mathbb{E}_{\mathbb{Q}}[X] + \rho s_B(X).$$

If $\Delta_{\mathcal{H}}(X) = L_X^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}}L_X^{\mathbb{Q}}$, then

$$0 \leq s_B(X) \leq \|\Delta_{\mathcal{H}}(X)\|_{2,\mathbb{Q}}, \quad s_B(X) \uparrow \|\Delta_{\mathcal{H}}(X)\|_{2,\mathbb{Q}} \quad \text{as } B \uparrow \infty.$$

If $\Delta_{\mathcal{H}}(X) \neq 0$ and

$$\frac{\Delta_{\mathcal{H}}(X)}{\|\Delta_{\mathcal{H}}(X)\|_{2,\mathbb{Q}}} \in \mathcal{R}_{\mathcal{H},\mathcal{K}}(B),$$

then $s_B(X) = \|\Delta_{\mathcal{H}}(X)\|_{2,\mathbb{Q}}$ already at that finite budget.

Proof. For $Z = 1 + \eta R \in \mathcal{Z}_{\rho,B}$, the condition $\rho B < 1$ gives $Z > 0$. Since $\mathbb{E}_{\mathbb{Q}}[R \mid \mathcal{H}] = 0$, one has $\mathbb{E}_{\mathbb{Q}}[Z] = 1$ and $\mathbb{E}_{\mathbb{Q}}[Z \mid \mathcal{H}] = 1$. Thus Z is public-equivalent.

For fixed R ,

$$\mathbb{E}_{\mathbb{Q}}[(1 + \eta R)X] = \mathbb{E}_{\mathbb{Q}}[X] + \eta \mathbb{E}_{\mathbb{Q}}[XR].$$

The residual set $\mathcal{R}_{\mathcal{H},\mathcal{K}}(B)$ is symmetric: with R it contains $-R$. Hence

$$\sup_{R \in \mathcal{R}_{\mathcal{H},\mathcal{K}}(B), |\eta| \leq \rho} \eta \mathbb{E}_{\mathbb{Q}}[XR] = \rho \sup_{R \in \mathcal{R}_{\mathcal{H},\mathcal{K}}(B)} \mathbb{E}_{\mathbb{Q}}[XR] = \rho s_B(X),$$

and the infimum is the negative of the same support value. This gives the two endpoint formulas.

Because R is $\mathcal{K}$ -measurable and conditionally mean zero relative to $\mathcal{H}$,

$$\mathbb{E}_{\mathbb{Q}}[XR] = \mathbb{E}_{\mathbb{Q}}[(L_X^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}}L_X^{\mathbb{Q}})R].$$

Cauchy–Schwarz and $\|R\|_{2,\mathbb{Q}} \leq 1$ give

$$s_B(X) \leq \|\Delta_{\mathcal{H}}(X)\|_{2,\mathbb{Q}}.$$

Nonnegativity follows because $0 \in \mathcal{R}_{\mathcal{H},\mathcal{K}}(B)$.

To prove the limiting statement, first note that $\Delta_{\mathcal{H}}(X) = 0$ makes both sides vanish. Assume therefore that $\Delta_{\mathcal{H}}(X) \neq 0$. Let $\Delta = \Delta_{\mathcal{H}}(X)$ and define

$$T_n = (-n) \vee (\Delta \wedge n), \quad R_n = T_n - \mathbb{E}_{\mathbb{Q}}[T_n \mid \mathcal{H}].$$

Then $R_n \in L^\infty(\mathbb{Q}, \mathcal{K})$, $\mathbb{E}_\mathbb{Q}[R_n \mid \mathcal{H}] = 0$, and $R_n \rightarrow \Delta$ in $L^2(\mathbb{Q})$. For large n , set $\tilde{R}_n = R_n / \|R_n\|_{2,\mathbb{Q}}$. Each $\tilde{R}_n$ belongs to $\mathcal{R}_{\mathcal{H},\mathcal{K}}(B_n)$ for some finite B_n , and

$$\mathbb{E}_\mathbb{Q}[X \tilde{R}_n] \rightarrow \|\Delta\|_{2,\mathbb{Q}}.$$

Fix $\varepsilon > 0$ and choose n such that $\mathbb{E}_\mathbb{Q}[X \tilde{R}_n] > \|\Delta\|_{2,\mathbb{Q}} - \varepsilon$. For every $B \geq B_n$, monotonicity of the residual budgets gives $s_B(X) \geq \|\Delta\|_{2,\mathbb{Q}} - \varepsilon$. Together with the Cauchy–Schwarz upper bound, this proves $s_B(X) \uparrow \|\Delta\|_{2,\mathbb{Q}}$. If the normalized residual itself lies in $\mathcal{R}_{\mathcal{H},\mathcal{K}}(B)$, it attains the Cauchy–Schwarz upper bound. $\square$

The support function also separates surface-equivalent and dynamic-safe budgets.

Corollary 4.19 (Surface and dynamic-safe residual-budget intervals). *Let $\mathcal{H} \subseteq \mathcal{G} \subseteq \mathcal{K}$, let $X \in L^2(\mathbb{Q})$, and let $B < \infty$, $\rho > 0$ with $\rho B < 1$. For $\mathcal{A} \in \{\mathcal{H}, \mathcal{G}\}$, define*

$$\mathcal{R}_{\mathcal{A},\mathcal{K}}(B) := \{R \in L^\infty(\mathbb{Q}, \mathcal{K}) : \mathbb{E}_\mathbb{Q}[R \mid \mathcal{A}] = 0, \|R\|_{2,\mathbb{Q}} \leq 1, \|R\|_\infty \leq B\},$$

$$s_{\mathcal{A},B}(X) := \sup_{R \in \mathcal{R}_{\mathcal{A},\mathcal{K}}(B)} \mathbb{E}_\mathbb{Q}[XR].$$

Then the surface-equivalent residual-budget interval is

$$\mathbb{E}_\mathbb{Q}[X] \pm \rho s_{\mathcal{H},B}(X),$$

while the dynamic-safe public-terminal interval is

$$\mathbb{E}_\mathbb{Q}[X] \pm \rho s_{\mathcal{G},B}(X).$$

In addition,

$$0 \leq s_{\mathcal{G},B}(X) \leq s_{\mathcal{H},B}(X),$$

and, as $B \uparrow \infty$,

$$s_{\mathcal{H},B}(X) \uparrow \|L_X^\mathbb{Q} - C_{\mathcal{H}}^\mathbb{Q} L_X^\mathbb{Q}\|_{2,\mathbb{Q}}, \quad s_{\mathcal{G},B}(X) \uparrow \|L_X^\mathbb{Q} - C_{\mathcal{G}}^\mathbb{Q} L_X^\mathbb{Q}\|_{2,\mathbb{Q}}.$$

Proof. Apply Theorem 4.18 twice, first with $\mathcal{A} = \mathcal{H}$ and then with $\mathcal{A} = \mathcal{G}$. Since $\mathcal{H} \subseteq \mathcal{G}$, the condition $\mathbb{E}_\mathbb{Q}[R \mid \mathcal{G}] = 0$ implies $\mathbb{E}_\mathbb{Q}[R \mid \mathcal{H}] = 0$, so $\mathcal{R}_{\mathcal{G},\mathcal{K}}(B) \subseteq \mathcal{R}_{\mathcal{H},\mathcal{K}}(B)$ and $s_{\mathcal{G},B}(X) \leq s_{\mathcal{H},B}(X)$. The limiting identities are the norm-duality limits from Theorem 4.18 for the two sigma-fields. $\square$

4.5. Static Hilbert selected stress-charge prices. A point quote requires a convention for unspanned latent stress. The static Hilbert selected price below is a convex residual selector.

Definition 4.20 (Compression residual and static Hilbert selected price). For $X \in L^2(\mathbb{Q})$, define the public compression residual

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

For $\gamma \geq 0$, define the static Hilbert selected Perron stress price

$$\Pi_{\text{Hilb}}^\gamma(X) := \mathbb{E}_\mathbb{Q}[X] + \gamma \|\Delta_{\mathcal{H}}(X)\|_{2,\mathbb{Q}}.$$

The static Hilbert selected price separates cash additivity from the nonlinear residual charge.

Proposition 4.21 (Basic properties of the static Hilbert selected price). *The map $\Pi_{\text{Hilb}}^\gamma : L^2(\mathbb{Q}) \rightarrow \mathbb{R}$ is cash additive, convex, and Lipschitz:*

$$\Pi_{\text{Hilb}}^\gamma(X + m) = \Pi_{\text{Hilb}}^\gamma(X) + m, \quad m \in \mathbb{R},$$

and

$$|\Pi_{\text{Hilb}}^\gamma(X) - \Pi_{\text{Hilb}}^\gamma(Y)| \leq (1 + \gamma) \|X - Y\|_{2,\mathbb{Q}}.$$

If $L_X^\mathbb{Q}$ is already public, then $\Pi_{\text{Hilb}}^\gamma(X) = \mathbb{E}_\mathbb{Q}[X]$.

Proof. Cash additivity follows because $L_{X+m}^{\mathbb{Q}} = L_X^{\mathbb{Q}} + m$ and the public residual of a constant is zero. Convexity follows from linearity of conditional expectation and the triangle inequality for the L^2 -norm. For Lipschitz continuity,

$$|\mathbb{E}_{\mathbb{Q}}[X - Y]| \leq \|X - Y\|_{2,\mathbb{Q}},$$

and

$$\|\Delta_{\mathcal{H}}(X) - \Delta_{\mathcal{H}}(Y)\|_{2,\mathbb{Q}} = \|\Delta_{\mathcal{H}}(X - Y)\|_{2,\mathbb{Q}} \leq \|X - Y\|_{2,\mathbb{Q}},$$

because $\Delta_{\mathcal{H}}$ is the difference of two nested orthogonal projections and hence an orthogonal projection onto $L^2(\mathbb{Q}, \mathcal{K}) \ominus L^2(\mathbb{Q}, \mathcal{H})$. If $L_X^{\mathbb{Q}}$ is already public, then $\Delta_{\mathcal{H}}(X) = 0$, so the stress-charge term vanishes and $\Pi_{\text{Hilb}}^{\gamma}(X) = \mathbb{E}_{\mathbb{Q}}[X]$. $\square$

Remark 4.22. $\Pi_{\text{Hilb}}^{\gamma}$ is a selected quote or risk-adjusted valuation rule after linear public calibration. Linear no-arbitrage pricing remains represented by the public-equivalent interval in Theorem 4.2. The static Hilbert selected price is the upper endpoint for a claim-specific residual family when the boundedness and positivity conditions of Theorem 4.25 hold. Outside those conditions, the construction is a nonlinear residual selector rather than the endpoint of a common linear pricing family.

Norm charges fail monotonicity unless an additional order-preserving selector is imposed.

Proposition 4.23 (Norm stress charges are not monotone in general). *Let $\gamma > 0$. Suppose $\mathcal{H}$ is trivial, $\mathcal{K} = \mathcal{F}$, and the probability space contains an event A with*

$$0 < \mathbb{Q}(A) < \frac{\gamma^2}{1 + \gamma^2}.$$

Then $\Pi_{\text{Hilb}}^{\gamma}(X) = \mathbb{E}_{\mathbb{Q}}[X] + \gamma \|X - \mathbb{E}_{\mathbb{Q}}[X]\|_{2,\mathbb{Q}}$ is not monotone: there are bounded claims $X \leq Y$ such that

$$\Pi_{\text{Hilb}}^{\gamma}(X) > \Pi_{\text{Hilb}}^{\gamma}(Y).$$

Proof. Let $p = \mathbb{Q}(A)$, choose $a > 0$, and set $Y = c$ and

$$X = c - a\mathbf{1}_A.$$

Then $X \leq Y$. Since $\Pi_{\text{Hilb}}^{\gamma}(Y) = c$,

$$\Pi_{\text{Hilb}}^{\gamma}(X) = c - ap + \gamma a \sqrt{p(1-p)}.$$

The strict inequality $p < \gamma^2/(1 + \gamma^2)$ is equivalent to $\gamma \sqrt{p(1-p)} > p$, and hence $\Pi_{\text{Hilb}}^{\gamma}(X) > c = \Pi_{\text{Hilb}}^{\gamma}(Y)$. $\square$

Remark 4.24. The proposition marks the interpretation boundary. The norm charge is a convex residual valuation or capitalized quote after public calibration. Monotone selected nonlinear pricing requires a monotone residual selector, such as the conditional entropic construction in Appendix C, or a coherent tail-risk functional.

The stress charge becomes a linear-envelope endpoint only in the claim's own bounded residual direction.

Theorem 4.25 (Stress charge as a residual-envelope endpoint). *Let $X \in L^2(\mathbb{Q})$, and set*

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

Assume $\Delta_{\mathcal{H}}(X) \in L^{\infty}(\mathbb{Q})$ and $\Delta_{\mathcal{H}}(X) \neq 0$. Define

$$R_X := \frac{\Delta_{\mathcal{H}}(X)}{\|\Delta_{\mathcal{H}}(X)\|_{2,\mathbb{Q}}}.$$

If $0 < \gamma \|R_X\|_{\infty} < 1$, then $R_X \in \mathcal{N}_{\mathcal{H}}^{\infty}$, and the upper endpoint of the residual price interval generated by

$$\frac{d\mathbb{Q}_{\eta}}{d\mathbb{Q}} = 1 + \eta R_X, \quad |\eta| \leq \gamma,$$

is exactly

$$\Pi_{\text{Hilb}}^\gamma(X) = \mathbb{E}_{\mathbb{Q}}[X] + \gamma \|\Delta_{\mathcal{H}}(X)\|_{2,\mathbb{Q}}.$$

The lower endpoint is

$$\mathbb{E}_{\mathbb{Q}}[X] - \gamma \|\Delta_{\mathcal{H}}(X)\|_{2,\mathbb{Q}}.$$

Proof. Since $C_{\mathcal{H}}^{\mathbb{Q}} L_X^{\mathbb{Q}}$ is the projection of $L_X^{\mathbb{Q}}$ onto $L^2(\mathbb{Q}, \mathcal{H})$,

$$\mathbb{E}_{\mathbb{Q}}[\Delta_{\mathcal{H}}(X) \mid \mathcal{H}] = 0.$$

Thus $R_X \in \mathcal{N}_{\mathcal{H}}^\infty$. The assumption $\gamma \|R_X\|_\infty < 1$ makes $1 + \eta R_X$ strictly positive for every $|\eta| \leq \gamma$, so Theorem 4.2 applies. In addition,

$$\mathbb{E}_{\mathbb{Q}}[X R_X] = \mathbb{E}_{\mathbb{Q}}[L_X^{\mathbb{Q}} R_X] = \mathbb{E}_{\mathbb{Q}}[(L_X^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}} L_X^{\mathbb{Q}}) R_X] = \|\Delta_{\mathcal{H}}(X)\|_{2,\mathbb{Q}}.$$

The price interval of Theorem 4.2 is therefore

$$\mathbb{E}_{\mathbb{Q}}[X] \pm \gamma \|\Delta_{\mathcal{H}}(X)\|_{2,\mathbb{Q}}.$$

Its upper endpoint is $\Pi_{\text{Hilb}}^\gamma(X)$ by Definition 4.20. $\square$

Remark 4.26. The theorem is the finite-budget equality case of Theorem 4.18. It does not assert that the same residual direction prices all claims. When the claim's own public compression residual is bounded, the static Hilbert selected price is the upper endpoint of a linear public-equivalent pricing family in that residual direction.

4.6. Finite public calibration. Let $\mathbf{q} \in \mathbb{R}^m$ be the vector of observed option-surface coordinates, and let

$$m(\theta) \in \mathbb{R}^m, \quad \theta \in \Theta,$$

be model-implied public coordinates. A Perron stress specification typically has

$$m_i(\theta) = a_i(\theta) C_{\mathcal{H}_i}^{\mathbb{Q}}(Z_\theta^{\mathbb{Q}}) + b_i(\theta),$$

or its date-indexed analogue.

Finite calibration is an optimization problem on public coordinates, not a recovery theorem.

Proposition 4.27 (Weighted public calibration). *Let Θ be compact, let $m : \Theta \rightarrow \mathbb{R}^m$ be continuous, and let W be a positive semidefinite $m \times m$ matrix. The objective*

$$J(\theta) := (\mathbf{q} - m(\theta))^\top W (\mathbf{q} - m(\theta))$$

admits a minimizer on Θ . Separately, if $\Theta = \mathbb{R}^p$, $m(\theta) = A\theta$, W is positive definite on the column space of A , and $A^\top W A$ is invertible, the unique unconstrained minimizer is

$$\hat{\theta} = (A^\top W A)^{-1} A^\top W \mathbf{q}.$$

If Θ is constrained, this formula solves the constrained problem only when the unconstrained minimizer belongs to Θ ; otherwise the minimizer is the constrained W -least-squares projection over Θ .

Proof. Continuity of J and compactness of Θ give existence by Weierstrass. In the unconstrained linear case, the normal equations are $A^\top W A \theta = A^\top W \mathbf{q}$, and invertibility gives the displayed solution. With a constraint set, the same quadratic objective is minimized over Θ , so the normal-equation solution is valid only if it is feasible. $\square$

5. SYNCHRONIZED STRESS AND CARRIER-ENRICHED PUBLIC SPANS

The residual fiber has a sharper form for equity-index options. Index variance contains a synchronization component linking the index to single-name variance carriers, beyond its volatility level. The index surface can be calibrated while the carrier decomposition remains incomplete.

Definition 5.1 (Synchronized-stress claim). Fix a maturity scale τ . Let $V_I^\tau \in L^2(\mathbb{Q})$ be an index variance-linked claim and let $V_1^\tau, \dots, V_n^\tau \in L^2(\mathbb{Q})$ be single-name variance-linked claims. For nonnegative index weights ω_i with $\sum_i \omega_i = 1$, define the diagonal carrier claim

$$V_{\text{diag}}^\tau(\omega) := \sum_{i=1}^n \omega_i^2 V_i^\tau$$

and the synchronized-stress claim

$$X_{\text{sync}}^\tau(\omega) := V_I^\tau - V_{\text{diag}}^\tau(\omega).$$

Its priced synchronized loading is

$$L_{\text{sync}}^\mathbb{Q} := C_{\mathcal{K}}^\mathbb{Q} X_{\text{sync}}^\tau(\omega).$$

The definition is payoff-level. In a continuous diffusion benchmark, X_{sync}^τ corresponds to the covariance part of the index variance decomposition. In the incomplete-information setting used here, index and single-name surfaces can calibrate its public price while the latent loading still retains a Perron residual.

Definition 5.2 (Implied synchronization coordinate). Let $v_I > 0$ and $v_i > 0$ be public total-variance coordinates for the index and the single names at the same date and maturity. Define

$$D_\omega(v) := 2 \sum_{1 \leq i < j \leq n} \omega_i \omega_j (v_i v_j)^{1/2}.$$

When $D_\omega(v) > 0$, the public implied synchronization coordinate is

$$\rho_{\text{sync}}^{\text{imp}} := \frac{v_I - \sum_{i=1}^n \omega_i^2 v_i}{D_\omega(v)}.$$

Remark 5.3. The coordinate $\rho_{\text{sync}}^{\text{imp}}$ is a public normalization of the index-versus-single-name variance wedge. Without additional constituent-completeness and no-arbitrage restrictions it need not lie in $[-1, 1]$, and should not be read as a literal correlation coefficient in every specification.

The coordinate is public and is built only from observed index and carrier-surface total-variance coordinates. It expresses the part of the index surface that prices synchronized variance relative to the visible single-name carrier span; the latent Perron state remains outside this coordinate.

Synchronized stress inherits price freedom exactly through its public residual loading.

Proposition 5.4 (Public-equivalent freedom for synchronized stress). *Let $X_{\text{sync}}^\tau(\omega)$ be as in Definition 5.1, and set*

$$\Delta_{\mathcal{H}}^{\text{sync}} := L_{\text{sync}}^\mathbb{Q} - C_{\mathcal{H}}^\mathbb{Q} L_{\text{sync}}^\mathbb{Q}.$$

If $\Delta_{\mathcal{H}}^{\text{sync}} \neq 0$, then public option calibration does not identify the price of $X_{\text{sync}}^\tau(\omega)$ inside the class of public-equivalent kernels. If, in addition, the normalized residual

$$R_{\text{sync}} := \frac{\Delta_{\mathcal{H}}^{\text{sync}}}{\|\Delta_{\mathcal{H}}^{\text{sync}}\|_{2,\mathbb{Q}}}$$

is bounded and $0 < \rho \|R_{\text{sync}}\|_\infty < 1$, then the public-equivalent densities

$$Z_\eta = 1 + \eta R_{\text{sync}}, \quad |\eta| \leq \rho,$$

generate the interval

$$\mathbb{E}_\mathbb{Q}[X_{\text{sync}}^\tau(\omega)] \pm \rho \|\Delta_{\mathcal{H}}^{\text{sync}}\|_{2,\mathbb{Q}}.$$

Proof. The residual $\Delta_{\mathcal{H}}^{\text{sync}}$ has zero conditional expectation given $\mathcal{H}$. If it is nonzero, Theorem 2.7, applied to $L_{\text{sync}}^\mathbb{Q}$ on $(\Omega, \mathcal{K}, \mathbb{Q})$ relative to $\mathcal{H}$, gives a bounded $\mathcal{K}$ -measurable residual direction R with $\mathbb{E}_\mathbb{Q}[L_{\text{sync}}^\mathbb{Q} R] \neq 0$. The tilt generated by R preserves public payoff prices by Theorem 3.3. Since

$$\mathbb{E}_\mathbb{Q}[X_{\text{sync}}^\tau R] = \mathbb{E}_\mathbb{Q}[L_{\text{sync}}^\mathbb{Q} R],$$

the tilted price of X_{sync}^τ moves by the same residual pairing. Under the additional bounded-normalized-residual assumption, R_{sync} itself is admissible whenever $0 < \rho \|R_{\text{sync}}\|_\infty < 1$. The displayed interval is Theorem 4.25 applied to $X = X_{\text{sync}}^\tau(\omega)$. A nonzero residual makes the interval nontrivial, so public calibration alone cannot identify a unique synchronized-stress price. $\square$

Definition 5.5 (Index, carrier, and tradeable spans). Let $\mathbf{H}_{\text{idx}} \subseteq L^2(\mathbb{Q}, \mathcal{H})$ be the finite public span generated by liquid index surface coordinates. Let $\mathbf{H}_{\text{car}}$ be the enlarged finite public span obtained by adding single-name carrier-surface coordinates, with

$$\mathbf{H}_{\text{idx}} \subseteq \mathbf{H}_{\text{car}} \subseteq L^2(\mathbb{Q}, \mathcal{H}).$$

Finally, let $\mathcal{S}_{\text{car}} \subseteq \mathbf{H}_{\text{car}}$ be the closed tradeable span generated by actually admissible public carrier instruments.

The quotient $\mathbf{H}_{\text{car}}/\mathcal{S}_{\text{car}}$ is economically operative. A single-name surface coordinate may improve public spanning even when the corresponding listed option portfolio is too costly, too sparse, or too unstable to serve as a hedge.

Carrier enrichment shrinks residual width by projection monotonicity.

Proposition 5.6 (Carrier enrichment shrinks residual price width). *Let*

$$\mathbf{U}_0 \subseteq \mathbf{U}_1 \subseteq L^2(\mathbb{Q}, \mathcal{K})$$

be closed linear subspaces, and let P_j be the $L^2(\mathbb{Q})$ -orthogonal projection onto $\mathbf{U}_j$. Let $X \in L^2(\mathbb{Q})$ have latent loading $L_X^\mathbb{Q} = C_K^\mathbb{Q} X$, and define

$$\Delta_j(X) := L_X^\mathbb{Q} - P_j L_X^\mathbb{Q}, \quad j = 0, 1.$$

Then

$$\|\Delta_0(X)\|_{2,\mathbb{Q}}^2 = \|P_1 L_X^\mathbb{Q} - P_0 L_X^\mathbb{Q}\|_{2,\mathbb{Q}}^2 + \|\Delta_1(X)\|_{2,\mathbb{Q}}^2.$$

In particular,

$$\|\Delta_1(X)\|_{2,\mathbb{Q}} \leq \|\Delta_0(X)\|_{2,\mathbb{Q}}.$$

Consequently, for any fixed stress-charge coefficient $\gamma \geq 0$, the selected residual charge weakly decreases when the public span is enriched from $\mathbf{U}_0$ to $\mathbf{U}_1$:

$$\gamma \|\Delta_1(X)\|_{2,\mathbb{Q}} \leq \gamma \|\Delta_0(X)\|_{2,\mathbb{Q}}.$$

Proof. Since $\mathbf{U}_0 \subseteq \mathbf{U}_1$, the projections satisfy $P_0 P_1 = P_0$. Hence

$$L_X^\mathbb{Q} - P_0 L_X^\mathbb{Q} = (P_1 L_X^\mathbb{Q} - P_0 L_X^\mathbb{Q}) + (L_X^\mathbb{Q} - P_1 L_X^\mathbb{Q}).$$

The first term lies in $\mathbf{U}_1$, while the second is orthogonal to $\mathbf{U}_1$. Pythagoras gives the identity and the monotonicity. The stress-charge statement follows by multiplying by γ . $\square$

The synchronized-stress interval contracts exactly when carrier coordinates add a nonzero projection of the residual stress direction.

Corollary 5.7 (Carrier-enriched synchronized-stress interval). *Let P_{idx} and P_{car} be the $L^2(\mathbb{Q})$ -orthogonal projections onto $\mathbf{H}_{\text{idx}}$ and $\mathbf{H}_{\text{car}}$, respectively. Then the carrier-enriched public span weakly reduces the synchronized-stress residual norm relative to the index-only public span. For any $\gamma > 0$, the selected residual charge reduction is strict exactly when*

$$P_{\text{car}} L_{\text{sync}}^\mathbb{Q} \neq P_{\text{idx}} L_{\text{sync}}^\mathbb{Q}.$$

Proof. Apply Proposition 5.6 to the synchronized-stress claim and to the nested spaces $\mathbf{H}_{\text{idx}} \subseteq \mathbf{H}_{\text{car}}$. Orthogonal projection gives

$$\|L_{\text{sync}}^\mathbb{Q} - P_{\text{idx}} L_{\text{sync}}^\mathbb{Q}\|_{2,\mathbb{Q}}^2 = \|P_{\text{car}} L_{\text{sync}}^\mathbb{Q} - P_{\text{idx}} L_{\text{sync}}^\mathbb{Q}\|_{2,\mathbb{Q}}^2 + \|L_{\text{sync}}^\mathbb{Q} - P_{\text{car}} L_{\text{sync}}^\mathbb{Q}\|_{2,\mathbb{Q}}^2.$$

The residual norm therefore cannot increase after carrier enrichment. It decreases exactly when $P_{\text{car}} L_{\text{sync}}^\mathbb{Q} - P_{\text{idx}} L_{\text{sync}}^\mathbb{Q} \neq 0$, which is the stated strictness condition. $\square$

A public-neutral spread can still carry latent residual price exposure.

Proposition 5.8 (Public-neutral residual spreads). *Let $X, Y \in L^2(\mathbb{Q})$ satisfy*

$$C_{\mathcal{H}}^{\mathbb{Q}} L_X^{\mathbb{Q}} = C_{\mathcal{H}}^{\mathbb{Q}} L_Y^{\mathbb{Q}}.$$

Set $U = X - Y$. Then U has zero public latent shadow:

$$C_{\mathcal{H}}^{\mathbb{Q}} L_U^{\mathbb{Q}} = 0.$$

If $L_U^{\mathbb{Q}} \neq 0$, public calibration assigns the same public latent shadow to X and Y , but public-equivalent residual kernels can move the price of U . In particular, if $L_U^{\mathbb{Q}} \in L^\infty(\mathbb{Q})$ and $0 < \rho \left\| L_U^{\mathbb{Q}} / \left\| L_U^{\mathbb{Q}} \right\|_{2, \mathbb{Q}} \right\|_\infty < 1$, then the normalized residual tilt generated by $L_U^{\mathbb{Q}}$ gives the interval

$$\mathbb{E}_{\mathbb{Q}}[U] \pm \rho \left\| L_U^{\mathbb{Q}} \right\|_{2, \mathbb{Q}}.$$

Proof. Linearity of $C_{\mathcal{K}}^{\mathbb{Q}}$ and $C_{\mathcal{H}}^{\mathbb{Q}}$ gives

$$C_{\mathcal{H}}^{\mathbb{Q}} L_U^{\mathbb{Q}} = C_{\mathcal{H}}^{\mathbb{Q}} L_X^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}} L_Y^{\mathbb{Q}} = 0.$$

Since constants are $\mathcal{H}$ -measurable, this also implies $\mathbb{E}_{\mathbb{Q}}[U] = \mathbb{E}_{\mathbb{Q}}[L_U^{\mathbb{Q}}] = 0$. Thus $L_U^{\mathbb{Q}}$ is itself a public-invisible latent residual, and the baseline spread price is zero. If $L_U^{\mathbb{Q}} \neq 0$, Theorem 2.7, applied on $(\Omega, \mathcal{K}, \mathbb{Q})$ relative to $\mathcal{H}$, supplies bounded $\mathcal{K}$ -measurable public-invisible residual directions with nonzero pairing against $L_U^{\mathbb{Q}}$. The resulting public-equivalent tilts move the spread price because $\mathbb{E}_{\mathbb{Q}}[UR] = \mathbb{E}_{\mathbb{Q}}[L_U^{\mathbb{Q}} R]$ for such R . If $L_U^{\mathbb{Q}}$ is bounded, the normalized residual $L_U^{\mathbb{Q}} / \left\| L_U^{\mathbb{Q}} \right\|_{2, \mathbb{Q}}$ is an admissible public-invisible direction under the stated positivity condition. The interval then follows from Theorem 4.2; it is the spread analogue of Proposition 5.4. $\square$

Remark 5.9. Proposition 5.8 is a pricing statement, not a trading rule. Two claims can be indistinguishable to a public calibration system while retaining different latent-stress prices. Listed-option relative-value implementations require execution, cost, liquidity, and carrier-absorption conditions. Those conditions belong to the implementation module, not to the pricing geometry.

6. VISIBLE HEDGING

Visible hedging is the projection problem on the tradeable public span.

Theorem 6.1 (Optimal visible hedge and residual decomposition). *Assume $\mathcal{S} \subseteq L^2(\mathbb{Q}, \mathcal{H})$ is closed and contains constants. For $X \in L^2(\mathbb{Q})$, the unique $L^2(\mathbb{Q})$ -optimal visible hedge is*

$$H_X^* = C_{\mathcal{S}}^{\mathbb{Q}}(X) = C_{\mathcal{S}}^{\mathbb{Q}}(L_X^{\mathbb{Q}}).$$

In addition,

$$\|X - H_X^*\|_{2, \mathbb{Q}}^2 = \left\| L_X^{\mathbb{Q}} - C_{\mathcal{S}}^{\mathbb{Q}}(L_X^{\mathbb{Q}}) \right\|_{2, \mathbb{Q}}^2 + \|X - L_X^{\mathbb{Q}}\|_{2, \mathbb{Q}}^2.$$

Proof. Since $\mathcal{S}$ is a closed subspace of $L^2(\mathbb{Q})$, the projection theorem gives a unique minimizer $H_X^* = C_{\mathcal{S}}^{\mathbb{Q}}(X)$. Since $\mathcal{S} \subseteq L^2(\mathbb{Q}, \mathcal{H}) \subseteq L^2(\mathbb{Q}, \mathcal{K})$, the orthogonal residual $X - L_X^{\mathbb{Q}} = X - C_{\mathcal{K}}^{\mathbb{Q}} X$ is orthogonal to $\mathcal{S}$. Therefore

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

Finally,

$$X - H_X^* = (L_X^{\mathbb{Q}} - C_{\mathcal{S}}^{\mathbb{Q}} L_X^{\mathbb{Q}}) + (X - L_X^{\mathbb{Q}}).$$

The first term is $\mathcal{K}$ -measurable, while the second is orthogonal to $L^2(\mathbb{Q}, \mathcal{K})$. Pythagoras gives the claimed identity. $\square$

Definition 6.2 (Bounded residual density set). Let $\mathcal{Z}$ be a nonempty family of pricing densities satisfying

$$0 < a \leq Z \leq b < \infty, \quad \mathbb{E}_{\mathbb{Q}}[Z] = 1.$$

For a closed visible hedge space $\mathcal{S} \subseteq L^2(\mathbb{Q})$, define the residual-kernel visible-hedging loss

$$\mathcal{L}_{\mathcal{Z}}(H; X) := \sup_{Z \in \mathcal{Z}} \mathbb{E}_{\mathbb{Q}}[Z(X - H)^2], \quad H \in \mathcal{S}.$$

Bounded residual-kernel uncertainty preserves the projection lower bound up to density constants.

Theorem 6.3 (Minimax visible hedge under residual-kernel uncertainty). *Let $X \in L^2(\mathbb{Q})$, let $\mathcal{S}$ be a closed linear subspace of $L^2(\mathbb{Q})$, and let $\mathcal{Z}$ satisfy Definition 6.2. There is a unique minimizer*

$$H_{\mathcal{Z}}^*(X) := \arg \min_{H \in \mathcal{S}} \mathcal{L}_{\mathcal{Z}}(H; X).$$

In addition,

$$a \|X - C_{\mathcal{S}}^{\mathbb{Q}} X\|_{2,\mathbb{Q}}^2 \leq \inf_{H \in \mathcal{S}} \mathcal{L}_{\mathcal{Z}}(H; X) \leq b \|X - C_{\mathcal{S}}^{\mathbb{Q}} X\|_{2,\mathbb{Q}}^2.$$

If $P \in L^2(\mathbb{Q})$ has $P_{\mathcal{S}}^{\perp} := P - C_{\mathcal{S}}^{\mathbb{Q}} P \neq 0$, then

$$\inf_{H \in \mathcal{S}} \mathcal{L}_{\mathcal{Z}}(H; X) \geq a \beta_{\mathcal{S}}(X; P)^2 \|P_{\mathcal{S}}^{\perp}\|_{2,\mathbb{Q}}^2,$$

where

$$\beta_{\mathcal{S}}(X; P) = \frac{\langle X - C_{\mathcal{S}}^{\mathbb{Q}} X, P_{\mathcal{S}}^{\perp} \rangle_{L^2(\mathbb{Q})}}{\|P_{\mathcal{S}}^{\perp}\|_{2,\mathbb{Q}}^2}.$$

Proof. For each fixed $Z \in \mathcal{Z}$, the map

$$H \mapsto \mathbb{E}_{\mathbb{Q}}[Z(X - H)^2]$$

is finite, continuous, weakly lower semicontinuous, and a -strongly convex in the $L^2(\mathbb{Q})$ -norm. The common lower density bound is the modulus. The supremum of functionals with this common modulus remains a -strongly convex, and a supremum of weakly lower semicontinuous functionals is weakly lower semicontinuous. The loss is proper because $\mathcal{L}_{\mathcal{Z}}(0; X) \leq b \|X\|_{2,\mathbb{Q}}^2$. Coercivity follows from

$$\mathcal{L}_{\mathcal{Z}}(H; X) \geq a \|X - H\|_{2,\mathbb{Q}}^2 \geq \frac{a}{2} \|H\|_{2,\mathbb{Q}}^2 - a \|X\|_{2,\mathbb{Q}}^2.$$

A minimizing sequence is therefore bounded in $\mathcal{S}$. Since $\mathcal{S}$ is a closed subspace of a Hilbert space, reflexivity and weak closedness give a weakly convergent subsequence whose limit remains in $\mathcal{S}$. Weak lower semicontinuity gives existence. Strong convexity gives uniqueness.

The lower bound follows from

$$\mathcal{L}_{\mathcal{Z}}(H; X) \geq a \|X - H\|_{2,\mathbb{Q}}^2$$

and minimizing over $H \in \mathcal{S}$. The upper bound follows by evaluating the minimax loss at the baseline projection:

$$\mathcal{L}_{\mathcal{Z}}(C_{\mathcal{S}}^{\mathbb{Q}} X; X) \leq b \|X - C_{\mathcal{S}}^{\mathbb{Q}} X\|_{2,\mathbb{Q}}^2.$$

For the Perron lower bound, $P_{\mathcal{S}}^{\perp} \perp \mathcal{S}$, so the scalar projection of $X - C_{\mathcal{S}}^{\mathbb{Q}} X$ onto $\text{span}\{P_{\mathcal{S}}^{\perp}\}$ has squared norm $\beta_{\mathcal{S}}(X; P)^2 \|P_{\mathcal{S}}^{\perp}\|_{2,\mathbb{Q}}^2$. Combining this one-dimensional projection bound with the first lower bound gives the claim. $\square$

The visible residual separates into claim, latent, and public-basis components.

Theorem 6.4 (Three-part visible hedge decomposition). *Assume $\mathcal{S} \subseteq L^2(\mathbb{Q}, \mathcal{H})$ is closed and contains constants. For every $X \in L^2(\mathbb{Q})$,*

$$\|X - C_{\mathcal{S}}^{\mathbb{Q}} X\|_{2,\mathbb{Q}}^2 = \|X - L_X^{\mathbb{Q}}\|_{2,\mathbb{Q}}^2 + \|L_X^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}} L_X^{\mathbb{Q}}\|_{2,\mathbb{Q}}^2 + \|C_{\mathcal{H}}^{\mathbb{Q}} L_X^{\mathbb{Q}} - C_{\mathcal{S}}^{\mathbb{Q}} L_X^{\mathbb{Q}}\|_{2,\mathbb{Q}}^2.$$

The three terms are, respectively, the claim component outside the latent Perron stress field, the latent compression residual not visible in public information, and the public basis residual not spanned by tradeable visible instruments.

Proof. Theorem 6.1 gives

$$\|X - C_S^{\mathbb{Q}}X\|_{2,\mathbb{Q}}^2 = \|X - L_X^{\mathbb{Q}}\|_{2,\mathbb{Q}}^2 + \|L_X^{\mathbb{Q}} - C_S^{\mathbb{Q}}L_X^{\mathbb{Q}}\|_{2,\mathbb{Q}}^2.$$

It remains to decompose the second term. Because

$$\mathcal{S} \subseteq L^2(\mathbb{Q}, \mathcal{H}) \subseteq L^2(\mathbb{Q}, \mathcal{K}),$$

the projection onto $\mathcal{S}$ satisfies

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

Therefore

$$L_X^{\mathbb{Q}} - C_S^{\mathbb{Q}}L_X^{\mathbb{Q}} = (L_X^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}}L_X^{\mathbb{Q}}) + (C_{\mathcal{H}}^{\mathbb{Q}}L_X^{\mathbb{Q}} - C_S^{\mathbb{Q}}L_X^{\mathbb{Q}}).$$

The first term is orthogonal to $L^2(\mathbb{Q}, \mathcal{H})$, while the second term is $\mathcal{H}$ -measurable. Pythagoras gives

$$\|L_X^{\mathbb{Q}} - C_S^{\mathbb{Q}}L_X^{\mathbb{Q}}\|_{2,\mathbb{Q}}^2 = \|L_X^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}}L_X^{\mathbb{Q}}\|_{2,\mathbb{Q}}^2 + \|C_{\mathcal{H}}^{\mathbb{Q}}L_X^{\mathbb{Q}} - C_S^{\mathbb{Q}}L_X^{\mathbb{Q}}\|_{2,\mathbb{Q}}^2.$$

Substitution proves the theorem. $\square$

Definition 6.5 (Perron hedge residual coordinate). Let $Z \in L^2(\mathbb{Q}, \mathcal{K})$ be a centered latent Perron stress coordinate with $\|Z - C_S^{\mathbb{Q}}Z\|_{2,\mathbb{Q}} > 0$. Define

$$Z_S^{\perp} := Z - C_S^{\mathbb{Q}}Z.$$

For $X \in L^2(\mathbb{Q})$, define its visible residual Perron coefficient by

$$\beta_S(X; Z) := \frac{\langle X - C_S^{\mathbb{Q}}X, Z_S^{\perp} \rangle_{L^2(\mathbb{Q})}}{\|Z_S^{\perp}\|_{2,\mathbb{Q}}^2}.$$

A nonzero Perron coefficient imposes a visible hedge-risk floor.

Theorem 6.6 (Latent Perron residual lower bound). *Under the assumptions of Definition 6.5,*

$$\inf_{H \in \mathcal{S}} \|X - H\|_{2,\mathbb{Q}}^2 = \|X - C_S^{\mathbb{Q}}X\|_{2,\mathbb{Q}}^2 \geq \beta_S(X; Z)^2 \|Z_S^{\perp}\|_{2,\mathbb{Q}}^2.$$

In particular, if $\beta_S(X; Z) \neq 0$, every visible hedge has strictly positive residual risk.

Proof. The first equality is the projection theorem. Since $Z_S^{\perp}$ is orthogonal to $\mathcal{S}$, the one-dimensional projection of $X - C_S^{\mathbb{Q}}X$ onto $\text{span}\{Z_S^{\perp}\}$ has squared norm

$$\frac{\langle X - C_S^{\mathbb{Q}}X, Z_S^{\perp} \rangle_{L^2(\mathbb{Q})}^2}{\|Z_S^{\perp}\|_{2,\mathbb{Q}}^2} = \beta_S(X; Z)^2 \|Z_S^{\perp}\|_{2,\mathbb{Q}}^2.$$

The squared norm of a vector dominates the squared norm of any of its orthogonal projections. This proves the inequality. $\square$

Corollary 6.7 (Pure Perron loading). *If*

$$L_X^{\mathbb{Q}} = a_X Z + B, \quad B - C_S^{\mathbb{Q}}B \perp Z_S^{\perp},$$

then

$$\inf_{H \in \mathcal{S}} \|X - H\|_{2,\mathbb{Q}}^2 \geq a_X^2 \|Z_S^{\perp}\|_{2,\mathbb{Q}}^2.$$

Proof. The orthogonality assumption gives $\beta_S(X; Z) = a_X$. Apply Theorem 6.6. $\square$

7. PHYSICAL RISK AND CAPITAL

The pricing model is calibrated under $\mathbb{Q}$. Hedge performance and capital are evaluated under the physical measure $\mathbb{P}$. A bounded density ratio transfers the Hilbert geometry across equivalent measures.

Assumption 7.1 (Bounded physical density). $\mathbb{P} \sim \mathbb{Q}$ and $M := d\mathbb{P}/d\mathbb{Q}$ satisfies

$$0 < m \leq M \leq \bar{m} < \infty \quad \mathbb{Q}\text{-a.s.}$$

Bounded density ratios preserve the Hilbert geometry across $\mathbb{Q}$ and $\mathbb{P}$.

Proposition 7.2 (Equivalent L^2 geometries). *Under Assumption 7.1,*

$$m \|Y\|_{2,\mathbb{Q}}^2 \leq \|Y\|_{2,\mathbb{P}}^2 \leq \bar{m} \|Y\|_{2,\mathbb{Q}}^2 \quad \text{for every } Y \in L^2(\mathbb{Q}).$$

Consequently $L^2(\mathbb{P}) = L^2(\mathbb{Q})$ as vector spaces with equivalent norms. Any subspace closed in $L^2(\mathbb{Q})$, including $\mathcal{S}$, is also closed in $L^2(\mathbb{P})$.

Proof. The inequalities follow directly from

$$\|Y\|_{2,\mathbb{P}}^2 = \mathbb{E}_{\mathbb{P}}[Y^2] = \mathbb{E}_{\mathbb{Q}}[MY^2]$$

and the bounds $m \leq M \leq \bar{m}$. They imply norm equivalence. If (Y_n) is a sequence in a $L^2(\mathbb{Q})$ -closed subspace $\mathcal{U}$ and $Y_n \rightarrow Y$ in $L^2(\mathbb{P})$, then the lower bound implies $Y_n \rightarrow Y$ in $L^2(\mathbb{Q})$. Hence $Y \in \mathcal{U}$. Thus $\mathcal{U}$ is also closed in $L^2(\mathbb{P})$. $\square$

Bayes' formula gives the exact quotient by which a conditional projection changes under an equivalent measure.

Theorem 7.3 (Bayes correction for conditional projections). *Let $\mathbb{P} \sim \mathbb{Q}$, let $M = d\mathbb{P}/d\mathbb{Q}$, and let $\mathcal{A} \subseteq \mathcal{F}$. Assume $M > 0$ and $X \in L^1(\mathbb{P}) \cap L^1(\mathbb{Q})$, and assume*

$$(M - C_{\mathcal{A}}^{\mathbb{Q}}M)(X - C_{\mathcal{A}}^{\mathbb{Q}}X) \in L^1(\mathbb{Q}).$$

Then

$$C_{\mathcal{A}}^{\mathbb{P}}X - C_{\mathcal{A}}^{\mathbb{Q}}X = \frac{\mathbb{E}_{\mathbb{Q}}[(M - C_{\mathcal{A}}^{\mathbb{Q}}M)(X - C_{\mathcal{A}}^{\mathbb{Q}}X) \mid \mathcal{A}]}{C_{\mathcal{A}}^{\mathbb{Q}}M}.$$

In particular, if M is $\mathcal{A}$ -measurable, then

$$C_{\mathcal{A}}^{\mathbb{P}}X = C_{\mathcal{A}}^{\mathbb{Q}}X.$$

If $X \in L^2(\mathbb{Q})$, $0 < m \leq M \leq \bar{m} < \infty$, and

$$\delta_{\mathcal{A}} := \|M - C_{\mathcal{A}}^{\mathbb{Q}}M\|_{\infty} < \infty,$$

then

$$\|C_{\mathcal{A}}^{\mathbb{P}}X - C_{\mathcal{A}}^{\mathbb{Q}}X\|_{2,\mathbb{Q}} \leq \frac{\delta_{\mathcal{A}}}{m} \|X - C_{\mathcal{A}}^{\mathbb{Q}}X\|_{2,\mathbb{Q}}.$$

Proof. Bayes' formula gives

$$C_{\mathcal{A}}^{\mathbb{P}}X = \frac{\mathbb{E}_{\mathbb{Q}}[MX \mid \mathcal{A}]}{\mathbb{E}_{\mathbb{Q}}[M \mid \mathcal{A}]}.$$

The denominator is strictly positive a.s. because $M > 0$ a.s. Subtracting $C_{\mathcal{A}}^{\mathbb{Q}}X$ leaves the numerator

$$\mathbb{E}_{\mathbb{Q}}[M(X - C_{\mathcal{A}}^{\mathbb{Q}}X) \mid \mathcal{A}].$$

Since $\mathbb{E}_{\mathbb{Q}}[X - C_{\mathcal{A}}^{\mathbb{Q}}X \mid \mathcal{A}] = 0$, this numerator is equal to

$$\mathbb{E}_{\mathbb{Q}}[(M - C_{\mathcal{A}}^{\mathbb{Q}}M)(X - C_{\mathcal{A}}^{\mathbb{Q}}X) \mid \mathcal{A}].$$

The displayed identity follows. If M is $\mathcal{A}$ -measurable, the numerator vanishes. For the bound, use $C_{\mathcal{A}}^{\mathbb{Q}}M \geq m$, the L^2 -contraction property of conditional expectation, and the definition of $\delta_{\mathcal{A}}$:

$$\begin{aligned} \|C_{\mathcal{A}}^{\mathbb{P}}X - C_{\mathcal{A}}^{\mathbb{Q}}X\|_{2,\mathbb{Q}} &\leq \frac{1}{m} \left\| \mathbb{E}_{\mathbb{Q}}[(M - C_{\mathcal{A}}^{\mathbb{Q}}M)(X - C_{\mathcal{A}}^{\mathbb{Q}}X) \mid \mathcal{A}] \right\|_{2,\mathbb{Q}} \\ &\leq \frac{1}{m} \left\| (M - C_{\mathcal{A}}^{\mathbb{Q}}M)(X - C_{\mathcal{A}}^{\mathbb{Q}}X) \right\|_{2,\mathbb{Q}} \\ &\leq \frac{\delta_{\mathcal{A}}}{m} \|X - C_{\mathcal{A}}^{\mathbb{Q}}X\|_{2,\mathbb{Q}}. \end{aligned}$$

□

Definition 7.4 ($\mathbb{Q}$ -visible hedge residual). For $X \in L^2(\mathbb{Q})$, define the terminal residual of the $\mathbb{Q}$ -optimal visible hedge by

$$R_X^{\mathcal{S}} := X - C_{\mathcal{S}}^{\mathbb{Q}}X.$$

Physical hedge risk inherits the risk-neutral projection floor through density bounds.

Theorem 7.5 (Physical transfer of visible hedge risk). *Under Assumption 7.1, for every $X \in L^2(\mathbb{Q})$ and every $H \in \mathcal{S}$,*

$$\|X - H\|_{2,\mathbb{P}}^2 \geq m \|X - H\|_{2,\mathbb{Q}}^2.$$

Consequently,

$$\inf_{H \in \mathcal{S}} \|X - H\|_{2,\mathbb{P}}^2 \geq m \inf_{H \in \mathcal{S}} \|X - H\|_{2,\mathbb{Q}}^2 \geq m \beta_{\mathcal{S}}(X; Z)^2 \|Z_{\mathcal{S}}^{\perp}\|_{2,\mathbb{Q}}^2$$

whenever the Perron coordinate in Definition 6.5 is available.

Proof. For any $H \in \mathcal{S}$,

$$\|X - H\|_{2,\mathbb{P}}^2 = \mathbb{E}_{\mathbb{Q}}[M(X - H)^2] \geq m \mathbb{E}_{\mathbb{Q}}[(X - H)^2].$$

Take the infimum over $H \in \mathcal{S}$, then apply Theorem 6.6. □

The $\mathbb{Q}$ -visible hedge remains controlled under $\mathbb{P}$ only up to the density-distortion factor.

Corollary 7.6 (Performance guarantee for the $\mathbb{Q}$ -visible hedge under $\mathbb{P}$). *Under Assumption 7.1, let*

$$G_{\mathbb{Q}} = C_{\mathcal{S}}^{\mathbb{Q}}X, \quad G_{\mathbb{P}} = C_{\mathcal{S}}^{\mathbb{P}}X.$$

Then

$$\|X - G_{\mathbb{Q}}\|_{2,\mathbb{P}}^2 \leq \frac{\bar{m}}{m} \|X - G_{\mathbb{P}}\|_{2,\mathbb{P}}^2.$$

Thus the baseline $\mathbb{Q}$ -projection hedge is within the density-ratio condition number of the physically optimal visible hedge.

Proof. First,

$$\|X - G_{\mathbb{Q}}\|_{2,\mathbb{P}}^2 \leq \bar{m} \|X - G_{\mathbb{Q}}\|_{2,\mathbb{Q}}^2 = \bar{m} \inf_{H \in \mathcal{S}} \|X - H\|_{2,\mathbb{Q}}^2.$$

For every $H \in \mathcal{S}$,

$$\|X - H\|_{2,\mathbb{P}}^2 \geq m \|X - H\|_{2,\mathbb{Q}}^2.$$

Taking infima gives

$$\inf_{H \in \mathcal{S}} \|X - H\|_{2,\mathbb{Q}}^2 \leq \frac{1}{m} \inf_{H \in \mathcal{S}} \|X - H\|_{2,\mathbb{P}}^2.$$

Combining the two inequalities proves the claim. □

A bounded family of residual tilts gives uniform physical-transfer constants.

Theorem 7.7 (Uniform physical transfer under bounded residual tilts). *Assume $M = d\mathbb{P}/d\mathbb{Q}$ satisfies*

$$0 < m \leq M \leq \bar{m} < \infty.$$

Let $R \in L^\infty(\mathbb{Q})$, let $\mathbb{E}_{\mathbb{Q}}[R \mid \mathcal{H}] = 0$, and suppose

$$|\eta| \|R\|_\infty \leq r < 1.$$

Set $Z_\eta = 1 + \eta R$ and $d\mathbb{Q}_\eta = Z_\eta d\mathbb{Q}$. Then

$$\frac{m}{1+r} \leq \frac{d\mathbb{P}}{d\mathbb{Q}_\eta} \leq \frac{\bar{m}}{1-r} \quad \mathbb{Q}_\eta\text{-a.s.}$$

Consequently, for every $Y \in L^2(\mathbb{Q})$,

$$\|Y\|_{2,\mathbb{P}}^2 \geq \frac{m}{1+r} \|Y\|_{2,\mathbb{Q}_\eta}^2.$$

Visible-hedge residual lower bounds therefore transfer uniformly over all such bounded residual tilts.

Proof. The bound on R gives

$$1 - r \leq Z_\eta \leq 1 + r.$$

Since $d\mathbb{P}/d\mathbb{Q}_\eta = M/Z_\eta$, the displayed density-ratio bounds follow. The measures $\mathbb{Q}_\eta$ and $\mathbb{Q}$ are equivalent, so the same tradeable residual is measured under equivalent probability structures. Hence

$$\|Y\|_{2,\mathbb{P}}^2 = \mathbb{E}_{\mathbb{Q}_\eta} \left[\frac{d\mathbb{P}}{d\mathbb{Q}_\eta} Y^2 \right] \geq \frac{m}{1+r} \mathbb{E}_{\mathbb{Q}_\eta} [Y^2].$$

□

Corollary 7.8 (Physical three-part lower bound). *Under Assumption 7.1,*

$$\inf_{H \in \mathcal{S}} \|X - H\|_{2,\mathbb{P}}^2 \geq m \left(\|X - L_X^\mathbb{Q}\|_{2,\mathbb{Q}}^2 + \|L_X^\mathbb{Q} - C_{\mathcal{H}}^\mathbb{Q} L_X^\mathbb{Q}\|_{2,\mathbb{Q}}^2 + \|C_{\mathcal{H}}^\mathbb{Q} L_X^\mathbb{Q} - C_{\mathcal{S}}^\mathbb{Q} L_X^\mathbb{Q}\|_{2,\mathbb{Q}}^2 \right).$$

Proof. Combine Theorem 7.5 with Theorem 6.4. □

The physical drift of a $\mathbb{Q}$ -orthogonal residual is governed by density misalignment.

Theorem 7.9 (Physical drift of the visible hedge residual). *Under Assumption 7.1, let $R_X^\mathcal{S} = X - C_{\mathcal{S}}^\mathbb{Q} X$. Since constants belong to $\mathcal{S}$,*

$$\mathbb{E}_{\mathbb{Q}}[R_X^\mathcal{S}] = 0.$$

Its physical mean satisfies

$$|\mathbb{E}_{\mathbb{P}}[R_X^\mathcal{S}]| \leq \|M - 1\|_{2,\mathbb{Q}} \|R_X^\mathcal{S}\|_{2,\mathbb{Q}}.$$

Consequently,

$$|\mathbb{E}_{\mathbb{P}}[R_X^\mathcal{S}]| \leq \|M - 1\|_{2,\mathbb{Q}} \left(\|X - L_X^\mathbb{Q}\|_{2,\mathbb{Q}}^2 + \|L_X^\mathbb{Q} - C_{\mathcal{H}}^\mathbb{Q} L_X^\mathbb{Q}\|_{2,\mathbb{Q}}^2 + \|C_{\mathcal{H}}^\mathbb{Q} L_X^\mathbb{Q} - C_{\mathcal{S}}^\mathbb{Q} L_X^\mathbb{Q}\|_{2,\mathbb{Q}}^2 \right)^{1/2}.$$

Proof. Because constants lie in $\mathcal{S}$, the projection residual $R_X^\mathcal{S}$ is orthogonal to constants in $L^2(\mathbb{Q})$, so $\mathbb{E}_{\mathbb{Q}}[R_X^\mathcal{S}] = 0$. Therefore

$$\mathbb{E}_{\mathbb{P}}[R_X^\mathcal{S}] = \mathbb{E}_{\mathbb{Q}}[M R_X^\mathcal{S}] = \mathbb{E}_{\mathbb{Q}}[(M - 1) R_X^\mathcal{S}].$$

Cauchy–Schwarz gives

$$|\mathbb{E}_{\mathbb{P}}[R_X^\mathcal{S}]| \leq \|M - 1\|_{2,\mathbb{Q}} \|R_X^\mathcal{S}\|_{2,\mathbb{Q}}.$$

The final bound is obtained by substituting the identity in Theorem 6.4. □

Definition 7.10 (Expected shortfall). For $\alpha \in (0, 1)$ and a terminal loss $Y \in L^2(\mathbb{P})$, define

$$\text{ES}_\alpha^\mathbb{P}(Y) := \inf_{r \in \mathbb{R}} \left\{ r + \frac{1}{1 - \alpha} \mathbb{E}_\mathbb{P}[(Y - r)_+] \right\}.$$

The formula is the Rockafellar–Uryasev representation of expected shortfall; see [35]. Coherent risk measures in the sense of [36] provide the surrounding risk-measure setting, and [37] gives a standard reference for discrete-time risk-measure and pricing conventions.

Expected shortfall supplies a convex capital input compatible with the residual geometry.

Proposition 7.11 (Expected shortfall as an admissible capital input). *For each $\alpha \in (0, 1)$, $\text{ES}_\alpha^\mathbb{P}$ is finite on $L^2(\mathbb{P})$, convex, and cash additive:*

$$\text{ES}_\alpha^\mathbb{P}(Y + c) = \text{ES}_\alpha^\mathbb{P}(Y) + c, \quad c \in \mathbb{R}.$$

In addition, if $\mathbb{E}_\mathbb{P}[Y] = 0$, then $\text{ES}_\alpha^\mathbb{P}(Y) \geq 0$.

Proof. Finiteness from above follows by taking $r = 0$:

$$\text{ES}_\alpha^\mathbb{P}(Y) \leq \frac{1}{1 - \alpha} \mathbb{E}_\mathbb{P}[Y_+] < \infty.$$

For a lower bound, use the pointwise inequality $r + (Y - r)_+ \geq Y$. Since $(1 - \alpha)^{-1} \geq 1$, for every r ,

$$r + \frac{1}{1 - \alpha} \mathbb{E}_\mathbb{P}[(Y - r)_+] \geq r + \mathbb{E}_\mathbb{P}[(Y - r)_+] \geq \mathbb{E}_\mathbb{P}[Y].$$

Thus the infimum is finite.

For cash additivity, set $s = r - c$. Then

$$r + \frac{1}{1 - \alpha} \mathbb{E}_\mathbb{P}[(Y + c - r)_+] = c + s + \frac{1}{1 - \alpha} \mathbb{E}_\mathbb{P}[(Y - s)_+],$$

and the substitution $s = r - c$ preserves the infimum.

For convexity, let $Y_1, Y_2 \in L^2(\mathbb{P})$, $\lambda \in [0, 1]$, and choose arbitrary $r_1, r_2 \in \mathbb{R}$. Put $r = \lambda r_1 + (1 - \lambda)r_2$. Since $x \mapsto x_+$ is convex,

$$(\lambda Y_1 + (1 - \lambda)Y_2 - r)_+ \leq \lambda(Y_1 - r_1)_+ + (1 - \lambda)(Y_2 - r_2)_+.$$

After taking expectations and adding r , one obtains the convexity inequality for these r_1, r_2 . Taking the infimum over r_1 and r_2 proves convexity of $\text{ES}_\alpha^\mathbb{P}$. Finally, the lower-bound argument gives $\text{ES}_\alpha^\mathbb{P}(Y) \geq \mathbb{E}_\mathbb{P}[Y] = 0$ when Y is centered. $\square$

Definition 7.12 (Compression-adjusted capital). Under Assumption 7.1, let $\rho_\mathbb{P} : L^2(\mathbb{P}) \rightarrow \mathbb{R} \cup \{+\infty\}$ be a convex cash-additive risk functional for terminal losses. For $X \in L^2(\mathbb{Q})$, set

$$\tilde{R}_X^{\mathcal{S}, \mathbb{P}} := R_X^\mathcal{S} - \mathbb{E}_\mathbb{P}[R_X^\mathcal{S}],$$

and define

$$\text{Cap}_\gamma(X) := \rho_\mathbb{P}(\tilde{R}_X^{\mathcal{S}, \mathbb{P}}) + \gamma \left\| L_X^\mathbb{Q} - C_\mathcal{S}^\mathbb{Q} L_X^\mathbb{Q} \right\|_{2, \mathbb{Q}}, \quad \gamma \geq 0.$$

Proposition 7.13 (Convexity of compression-adjusted capital). *If $\rho_\mathbb{P}$ is convex, then Cap_γ is convex on $L^2(\mathbb{Q})$. If $\rho_\mathbb{P}(Y) \geq 0$ whenever $\mathbb{E}_\mathbb{P}[Y] = 0$, then*

$$\text{Cap}_\gamma(X) \geq \gamma \left\| L_X^\mathbb{Q} - C_\mathcal{S}^\mathbb{Q} L_X^\mathbb{Q} \right\|_{2, \mathbb{Q}}.$$

Proof. The map

$$X \mapsto \tilde{R}_X^{\mathcal{S}, \mathbb{P}} = X - C_\mathcal{S}^\mathbb{Q} X - \mathbb{E}_\mathbb{P}[X - C_\mathcal{S}^\mathbb{Q} X]$$

is linear from $L^2(\mathbb{Q})$ into $L^2(\mathbb{P})$. Hence convexity of $\rho_\mathbb{P}$ gives convexity of the first term. The map

$$X \mapsto L_X^\mathbb{Q} - C_\mathcal{S}^\mathbb{Q} L_X^\mathbb{Q}$$

is linear, so its $L^2(\mathbb{Q})$ -norm is convex. The lower bound follows from the assumed nonnegativity of $\rho_\mathbb{P}$ on centered hedge residuals, since $\tilde{R}_X^{\mathcal{S}, \mathbb{P}}$ is centered under $\mathbb{P}$ in Definition 7.12. $\square$

Compression-adjusted capital inherits the latent and public-basis residual floors.

Theorem 7.14 (Compression capital lower bounds). *Assume $\rho_{\mathbb{P}}$ is convex and satisfies $\rho_{\mathbb{P}}(Y) \geq 0$ whenever $\mathbb{E}_{\mathbb{P}}[Y] = 0$. Then for every $X \in L^2(\mathbb{Q})$,*

$$\text{Cap}_{\gamma}(X) \geq \gamma \left(\|L_X^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}} L_X^{\mathbb{Q}}\|_{2,\mathbb{Q}}^2 + \|C_{\mathcal{H}}^{\mathbb{Q}} L_X^{\mathbb{Q}} - C_{\mathcal{S}}^{\mathbb{Q}} L_X^{\mathbb{Q}}\|_{2,\mathbb{Q}}^2 \right)^{1/2}.$$

If a Perron coordinate Z is available as in Definition 6.5, then

$$\text{Cap}_{\gamma}(X) \geq \gamma |\beta_{\mathcal{S}}(X; Z)| \|Z_{\mathcal{S}}^{\perp}\|_{2,\mathbb{Q}}.$$

Proof. The first lower bound in Proposition 7.13 gives

$$\text{Cap}_{\gamma}(X) \geq \gamma \|L_X^{\mathbb{Q}} - C_{\mathcal{S}}^{\mathbb{Q}} L_X^{\mathbb{Q}}\|_{2,\mathbb{Q}}.$$

Since $\mathcal{S} \subseteq L^2(\mathbb{Q}, \mathcal{H})$, the proof of Theorem 6.4 gives the orthogonal decomposition

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

with the first term orthogonal to $L^2(\mathbb{Q}, \mathcal{H})$ and the second term belonging to $L^2(\mathbb{Q}, \mathcal{H})$. Pythagoras gives the displayed square-root expression.

For the Perron bound, $Z_{\mathcal{S}}^{\perp}$ is orthogonal to $\mathcal{S}$, and

$$\langle X - C_{\mathcal{S}}^{\mathbb{Q}} X, Z_{\mathcal{S}}^{\perp} \rangle = \langle L_X^{\mathbb{Q}} - C_{\mathcal{S}}^{\mathbb{Q}} L_X^{\mathbb{Q}}, Z_{\mathcal{S}}^{\perp} \rangle.$$

By Cauchy–Schwarz and Definition 6.5,

$$|\beta_{\mathcal{S}}(X; Z)| \|Z_{\mathcal{S}}^{\perp}\|_{2,\mathbb{Q}} \leq \|L_X^{\mathbb{Q}} - C_{\mathcal{S}}^{\mathbb{Q}} L_X^{\mathbb{Q}}\|_{2,\mathbb{Q}}.$$

Combining this inequality with the capital lower bound proves the result. $\square$

Remark 7.15. Expected shortfall is one admissible choice for $\rho_{\mathbb{P}}$. Proposition 7.11 verifies the convexity and nonnegativity properties used above. The results are stated for a general convex risk functional because the projection arguments do not depend on a particular tail functional.

8. CONVEX COMPRESSION RISK

The L^2 -residual is not the only compression cost. For variance-linked and stress-amplifying claims, convexity adds a pricing and capital term.

Assumption 8.1 (Strong convexity). $\Phi : \mathbb{R} \rightarrow \mathbb{R}$ is Borel, integrable at the random variables considered in this section, and κ -strongly convex for some $\kappa > 0$:

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

for every $x, y \in \mathbb{R}$ and some subgradient $g_x \in \partial\Phi(x)$.

Strong convexity converts public compression into a quadratic latent loss.

Theorem 8.2 (Latent convex compression gap). *Under Assumption 8.1, let $X \in L^2(\mathbb{Q})$, set $L_X^{\mathbb{Q}} = C_{\mathcal{K}}^{\mathbb{Q}} X$, and assume $\Phi(L_X^{\mathbb{Q}}), \Phi(C_{\mathcal{H}}^{\mathbb{Q}} L_X^{\mathbb{Q}}) \in L^1(\mathbb{Q})$. Then*

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

If $Z \in L^2(\mathbb{Q}, \mathcal{K})$ has nonzero public residual $Z_{\mathcal{H}}^{\perp}$, then

$$\mathbb{E}_{\mathbb{Q}}[\Phi(L_X^{\mathbb{Q}})] - \mathbb{E}_{\mathbb{Q}}[\Phi(C_{\mathcal{H}}^{\mathbb{Q}} L_X^{\mathbb{Q}})] \geq \frac{\kappa}{2} a_{\mathcal{H}}(X; Z)^2 \|Z_{\mathcal{H}}^{\perp}\|_{2,\mathbb{Q}}^2.$$

Proof. Set $Y = C_{\mathcal{H}}^{\mathbb{Q}} L_X^{\mathbb{Q}}$. Strong convexity gives the convex function $\Psi(x) := \Phi(x) - \frac{\kappa}{2}x^2$. The stated integrability and $L_X^{\mathbb{Q}}, Y \in L^2(\mathbb{Q})$ make the following expectations finite. Conditional Jensen gives

$$\mathbb{E}_{\mathbb{Q}}[\Psi(L_X^{\mathbb{Q}}) \mid \mathcal{H}] \geq \Psi(Y).$$

After taking expectations and rearranging,

$$\mathbb{E}_{\mathbb{Q}}[\Phi(L_X^{\mathbb{Q}})] - \mathbb{E}_{\mathbb{Q}}[\Phi(Y)] \geq \frac{\kappa}{2} \left(\mathbb{E}_{\mathbb{Q}}[(L_X^{\mathbb{Q}})^2] - \mathbb{E}_{\mathbb{Q}}[Y^2] \right) = \frac{\kappa}{2} \|L_X^{\mathbb{Q}} - Y\|_{2,\mathbb{Q}}^2,$$

because Y is the orthogonal projection of $L_X^\mathbb{Q}$ onto $L^2(\mathbb{Q}, \mathcal{H})$. This proves the first inequality. The second follows because the one-dimensional projection of $L_X^\mathbb{Q} - C_{\mathcal{H}}^\mathbb{Q} L_X^\mathbb{Q}$ onto $\text{span}\{Z_{\mathcal{H}}^\perp\}$ has squared norm

$$a_{\mathcal{H}}(X; Z)^2 \left\| Z_{\mathcal{H}}^\perp \right\|_{2, \mathbb{Q}}^2,$$

which is bounded above by $\left\| L_X^\mathbb{Q} - C_{\mathcal{H}}^\mathbb{Q} L_X^\mathbb{Q} \right\|_{2, \mathbb{Q}}^2$. $\square$

Compression onto the public field gives the corresponding convex gap.

Theorem 8.3 (Convex public compression gap). *Under Assumption 8.1, let $X \in L^2(\mathbb{Q})$, and assume $\Phi(X), \Phi(C_{\mathcal{H}}^\mathbb{Q} X) \in L^1(\mathbb{Q})$. Then*

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

If $L_X^\mathbb{Q} = C_{\mathcal{K}}^\mathbb{Q} X$, then

$$\left\| X - C_{\mathcal{H}}^\mathbb{Q} X \right\|_{2, \mathbb{Q}}^2 = \left\| L_X^\mathbb{Q} - C_{\mathcal{H}}^\mathbb{Q} L_X^\mathbb{Q} \right\|_{2, \mathbb{Q}}^2 + \left\| X - L_X^\mathbb{Q} \right\|_{2, \mathbb{Q}}^2.$$

Proof. Set $Y = C_{\mathcal{H}}^\mathbb{Q} X$. Strong convexity gives the convex function $\Psi(x) := \Phi(x) - \frac{\kappa}{2} x^2$. The stated integrability and $X, Y \in L^2(\mathbb{Q})$ make the following expectations finite. Conditional Jensen gives

$$\mathbb{E}_{\mathbb{Q}}[\Psi(X) \mid \mathcal{H}] \geq \Psi(Y).$$

After taking expectations and rearranging,

$$\mathbb{E}_{\mathbb{Q}}[\Phi(X)] - \mathbb{E}_{\mathbb{Q}}[\Phi(Y)] \geq \frac{\kappa}{2} \left(\mathbb{E}_{\mathbb{Q}}[X^2] - \mathbb{E}_{\mathbb{Q}}[Y^2] \right) = \frac{\kappa}{2} \left\| X - Y \right\|_{2, \mathbb{Q}}^2,$$

because $Y = C_{\mathcal{H}}^\mathbb{Q} X$ is the orthogonal projection of X onto $L^2(\mathbb{Q}, \mathcal{H})$. The second identity is Pythagoras for the nested projections

$$L^2(\mathbb{Q}, \mathcal{H}) \subseteq L^2(\mathbb{Q}, \mathcal{K}) \subseteq L^2(\mathbb{Q}, \mathcal{F}).$$

$\square$

Corollary 8.4 (Convex Perron residual contribution). *Under the assumptions of Theorem 8.3, if Z is a Perron stress coordinate as in Definition 6.5 with $\mathcal{S} = L^2(\mathbb{Q}, \mathcal{H})$, write $\beta_{\mathcal{H}}(X; Z) := \beta_{\mathcal{S}}(X; Z)$ for this specialization. Then*

$$\mathbb{E}_{\mathbb{Q}}[\Phi(X)] - \mathbb{E}_{\mathbb{Q}}[\Phi(C_{\mathcal{H}}^\mathbb{Q} X)] \geq \frac{\kappa}{2} \beta_{\mathcal{H}}(X; Z)^2 \left\| Z - C_{\mathcal{H}}^\mathbb{Q} Z \right\|_{2, \mathbb{Q}}^2.$$

Proof. Combine Theorem 8.3 with Theorem 6.6, taking $\mathcal{S} = L^2(\mathbb{Q}, \mathcal{H})$. $\square$

9. OBSERVABLE ESTIMANDS

Fixed-date pricing statements generate date-indexed estimands. The objects below specify what the pricing geometry makes observable; they are neither estimates nor trading rules, and they do not recover the latent Perron state.

Fix a finite set of dates $\mathcal{D}$. Let P_d be a Perron stress coordinate available at date $d \in \mathcal{D}$, and let $Q_d \in \mathbb{R}^m$ be a finite vector of public option-surface coordinates. The finite public span is

$$\mathbf{H}_m = \text{span}\{1, Q_{\cdot, 1}, \dots, Q_{\cdot, m}\}.$$

The finite public projection is

$$P_d = \hat{P}_{\mathbf{H}_m, d} + \hat{P}_{\mathbf{H}_m, d}^\perp,$$

with projection and residual shares, all sums being over $d \in \mathcal{D}$,

$$\hat{\mathbf{C}}_m := \frac{\sum_{d \in \mathcal{D}} (\hat{P}_{\mathbf{H}_m, d} - \bar{\hat{P}}_{\mathbf{H}_m})^2}{\sum_{d \in \mathcal{D}} (P_d - \bar{P})^2}, \quad \hat{\mathbf{R}}_m := 1 - \hat{\mathbf{C}}_m.$$

These ratios are reported only when the denominator is strictly positive. They are finite-span quantities. They describe the public shadow and residual component associated with the chosen public coordinates, not the full latent sigma-field.

Carrier enrichment is the second finite-span object. If H_{idx} is the span generated by liquid index-surface coordinates and $H_{\text{car}} \supseteq H_{\text{idx}}$ is the span enlarged by single-name carrier-surface coordinates, the carrier increment is

$$\Delta \hat{C}_{\text{car}} := \hat{C}(H_{\text{car}}) - \hat{C}(H_{\text{idx}}).$$

Proposition 5.6 gives the population counterpart. Enlarging the public span weakly reduces residual interval width. In a date-indexed implementation, the carrier increment asks whether single-name public surfaces compress synchronized stress beyond the index surface alone.

Residual hedge risk is the third object. For a future or terminal stress claim Y and a visible public span $\mathcal{S}_m$, the finite visible hedge is the least-squares projection $\hat{C}_{\mathcal{S}_m} Y$, and the observable residual floor is

$$\hat{F}(Y; \mathcal{S}_m) := \frac{\sum_{d \in \mathcal{D}} (Y_d - \hat{C}_{\mathcal{S}_m} Y_d)^2}{\sum_{d \in \mathcal{D}} (Y_d - \bar{Y})^2}.$$

Again the ratio is defined only when the denominator is strictly positive. Tail versions replace the squared loss by an upper-tail or spectral risk functional. These quantities correspond to the visible-hedging lower bounds in Section 6 and the capital functionals in Section 7.

The synchronized-stress wedge is the fourth object. Given public total-variance coordinates $v_{I,d,\tau}$ for an index and $v_{i,d,\tau}$ for single-name carriers, Definition 5.2 gives

$$\rho_{\text{sync},d,\tau}^{\text{imp}} = \frac{v_{I,d,\tau} - \sum_i \omega_i^2 v_{i,d,\tau}}{2 \sum_{i < j} \omega_i \omega_j (v_{i,d,\tau} v_{j,d,\tau})^{1/2}}.$$

If $\hat{\rho}_{\text{sync},d,\tau}^{\text{real}}$ is a realized synchronization proxy available at date d , the observable pricing wedge is

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

This wedge is an observable coordinate induced by the synchronized-stress claim of Definition 5.1. Execution, liquidity, carrier absorption, and information-timing controls are required before it becomes an executable pricing or portfolio rule.

A listed-option synchronization implementation can turn this coordinate into the candidate source for a finite public trading rule. Such an implementation adds contract-level execution, carrier admissibility, liquidity filters, lower-bound sellability, and adaptedness audits to the estimands defined here. The two objects live in different quotients. The Hilbert geometry identifies the residual-risk objects. A trading implementation tests whether the public listed-option span makes those objects executable.

The estimands separate the public-surface span, the carrier-enrichment increment, the residual hedge-risk floor, and the wedge between priced index synchronization and realized carrier synchronization. The pricing and hedging geometry stops at this finite observable layer. Numerical answers belong to a data or trading implementation.

10. SCOPE AND EXTENSIONS

The theory is stated in L^2 , so pricing, hedging, and residual risk share one projection geometry. The appendices record extensions needed at implementation boundaries. Appendix A replaces static terminal hedge spaces by stochastic-integral hedge spaces. Appendix B treats a low-dimensional Perron stress state. Appendix C gives density-band and entropic public-equivalent pricing sets beyond local linear tilts. Appendix D extends the capital functional from expected shortfall to spectral risk measures. Appendix E incorporates bid-ask and liquidity constraints into the public calibration system.

These extensions preserve the projection order. Public calibration, latent stress pricing, visible hedging, and residual capital are distinct projections of one incomplete-information problem. They also mark the boundary of the analysis. The synchronized-stress and carrier-enrichment results identify the residual pricing object and the public spans that compress it. Executable

listed-option strategies, contract-level portfolio optimizers, and full risk-neutral dynamics for the latent Perron state belong to separate implementation or modeling modules.

11. CONCLUSION

Public option surfaces were the starting object. They calibrate admitted payoff prices, but they do not determine every latent stress claim consistent with those prices. The paper formalizes that failure as a residual pricing fiber. The latent Perron loading is projected onto public information; the component outside the projection remains available to public-equivalent pricing kernels.

The core object is the Hilbert projection chain. Public-equivalent kernels preserve public option prices while moving latent-stress claims. Public prices therefore identify compressed priced loadings, not full latent loadings. Visible hedging is a different projection onto the tradeable public span. Its residual inherits a lower bound from the latent Perron component whenever that component remains outside the span. Under an equivalent physical measure, the same residual becomes hedge risk and capital.

The finite layer audits the obstruction rather than removing it. Public projection shares, carrier-enrichment increments, residual hedge-risk floors, and implied-versus-realized synchronization wedges ask whether the public surface spans the stress component, whether carrier surfaces add public information, and whether a visible hedge leaves a residual. They are pricing and hedging estimands. Contract-level execution is a separate implementation module.

A compression-consistent pricing model must therefore do more than fit the surface. It must state which quotient the surface identifies, which synchronized-stress component lies outside the carrier span, and which latent stress risk survives public projection.

APPENDIX A. DYNAMIC VISIBLE HEDGE SPACES

The main text uses a closed terminal subspace $\mathcal{S} \subseteq L^2(\mathbb{Q}, \mathcal{H})$. The dynamic version replaces it by the terminal closed span of public stochastic integrals. The projection geometry is unchanged. The martingale arguments use standard conditional-expectation facts; see, for example, [38].

Definition A.1 (Dynamic public hedge space). Let $(\mathcal{F}_t)_{0 \leq t \leq T}$ be a filtration under $\mathbb{Q}$, and let

$$S = (S^1, \dots, S^d)$$

be a discounted d -dimensional square-integrable $\mathbb{Q}$ -martingale adapted to the public filtration. Let $\mathcal{I}_T(S)$ be the $L^2(\mathbb{Q})$ -closure of terminal stochastic integrals

$$\int_0^T \vartheta_t dS_t$$

over bounded predictable simple strategies ϑ . The dynamic visible hedge space is

$$\mathcal{S}_{\text{dyn}} := \mathbb{R} + \mathcal{I}_T(S).$$

Assumption A.2 (Public dynamic spanning). The terminal gains in $\mathcal{S}_{\text{dyn}}$ are $\mathcal{H}$ -measurable, so

$$\mathcal{S}_{\text{dyn}} \subseteq L^2(\mathbb{Q}, \mathcal{H}).$$

Dynamic-safe public-terminal kernels preserve public martingales.

Theorem A.3 (Dynamic public martingale preservation). *Let S be a discounted square-integrable $\mathbb{Q}$ -martingale adapted to a public filtration $(\mathcal{G}_t)_{0 \leq t \leq T}$. Let $Z > 0$ satisfy $\mathbb{E}_{\mathbb{Q}}[Z] = 1$ and*

$$\mathbb{E}_{\mathbb{Q}}[Z \mid \mathcal{G}_T] = 1.$$

If $d\mathbb{Q}^Z = Z d\mathbb{Q}$, then S is also a $\mathbb{Q}^Z$ -martingale with respect to $(\mathcal{G}_t)_{0 \leq t \leq T}$.

Proof. The variable S_t is $\mathbb{Q}^Z$ -integrable because S_t is $\mathcal{G}_T$ -measurable and

$$\mathbb{E}_{\mathbb{Q}^Z}[|S_t|] = \mathbb{E}_{\mathbb{Q}}[Z|S_t|] = \mathbb{E}_{\mathbb{Q}}[\mathbb{E}_{\mathbb{Q}}[Z \mid \mathcal{G}_T]|S_t|] = \mathbb{E}_{\mathbb{Q}}[|S_t|] < \infty.$$

For $s \leq t \leq T$, Bayes' formula gives

$$\mathbb{E}_{\mathbb{Q}^Z}[S_t \mid \mathcal{G}_s] = \frac{\mathbb{E}_{\mathbb{Q}}[Z S_t \mid \mathcal{G}_s]}{\mathbb{E}_{\mathbb{Q}}[Z \mid \mathcal{G}_s]}.$$

The denominator is strictly positive a.s. The assumption and the tower property imply $\mathbb{E}_{\mathbb{Q}}[Z | \mathcal{G}_u] = 1$ for every $u \leq T$. Therefore

$$\mathbb{E}_{\mathbb{Q}}[ZS_t | \mathcal{G}_s] = \mathbb{E}_{\mathbb{Q}}[\mathbb{E}_{\mathbb{Q}}[Z | \mathcal{G}_t]S_t | \mathcal{G}_s] = \mathbb{E}_{\mathbb{Q}}[S_t | \mathcal{G}_s] = S_s,$$

and $\mathbb{E}_{\mathbb{Q}}[Z | \mathcal{G}_s] = 1$. Hence $\mathbb{E}_{\mathbb{Q}}[S_t | \mathcal{G}_s] = S_s$. $\square$

The dynamic hedge space reduces the terminal problem to the same projection.

Theorem A.4 (Dynamic projection hedge). *Under Assumption A.2, $\mathcal{S}_{\text{dyn}}$ is a closed linear subspace of $L^2(\mathbb{Q}, \mathcal{H})$ containing constants. For every $X \in L^2(\mathbb{Q})$, the optimal dynamic visible hedge is*

$$H_{X, \text{dyn}}^* = C_{\mathcal{S}_{\text{dyn}}}^{\mathbb{Q}} X,$$

and

$$\|X - H_{X, \text{dyn}}^*\|_{2, \mathbb{Q}}^2 = \|X - L_X^{\mathbb{Q}}\|_{2, \mathbb{Q}}^2 + \|L_X^{\mathbb{Q}} - C_{\mathcal{H}}^{\mathbb{Q}} L_X^{\mathbb{Q}}\|_{2, \mathbb{Q}}^2 + \|C_{\mathcal{H}}^{\mathbb{Q}} L_X^{\mathbb{Q}} - C_{\mathcal{S}_{\text{dyn}}}^{\mathbb{Q}} L_X^{\mathbb{Q}}\|_{2, \mathbb{Q}}^2.$$

Proof. By definition, $\mathcal{I}_T(S)$ is closed in $L^2(\mathbb{Q})$, and adding constants preserves closedness. Assumption A.2 places $\mathcal{S}_{\text{dyn}}$ inside $L^2(\mathbb{Q}, \mathcal{H})$. The projection theorem gives the unique minimizer $C_{\mathcal{S}_{\text{dyn}}}^{\mathbb{Q}} X$. Applying Theorem 6.4 with $\mathcal{S} = \mathcal{S}_{\text{dyn}}$ gives the displayed decomposition. $\square$

Dynamic trading removes only the Perron component contained in the stochastic-integral span.

Corollary A.5 (Dynamic Perron residual lower bound). *Let $Z \in L^2(\mathbb{Q}, \mathcal{K})$ be centered, and assume*

$$Z_{\text{dyn}}^{\perp} := Z - C_{\mathcal{S}_{\text{dyn}}}^{\mathbb{Q}} Z \neq 0.$$

Define

$$\beta_{\text{dyn}}(X; Z) := \frac{\langle X - C_{\mathcal{S}_{\text{dyn}}}^{\mathbb{Q}} X, Z_{\text{dyn}}^{\perp} \rangle_{L^2(\mathbb{Q})}}{\|Z_{\text{dyn}}^{\perp}\|_{2, \mathbb{Q}}^2}.$$

Then

$$\inf_{H \in \mathcal{S}_{\text{dyn}}} \|X - H\|_{2, \mathbb{Q}}^2 \geq \beta_{\text{dyn}}(X; Z)^2 \|Z_{\text{dyn}}^{\perp}\|_{2, \mathbb{Q}}^2.$$

Proof. Apply Theorem 6.6 with $\mathcal{S} = \mathcal{S}_{\text{dyn}}$. $\square$

APPENDIX B. LOW-DIMENSIONAL PERRON STRESS STATES

The scalar Perron coordinate is the baseline case. Applications may retain a small Perron block rather than a single residual direction. The projection geometry extends by replacing one residual coordinate with a finite-dimensional residual factor space.

Definition B.1 (Vector Perron stress coordinate). Let

$$Z = (Z_1, \dots, Z_r)^{\top} \in L^2(\mathbb{Q}, \mathcal{K}; \mathbb{R}^r)$$

be centered componentwise. Define its visible residual vector by

$$Z_{\mathcal{S}}^{\perp} := Z - C_{\mathcal{S}}^{\mathbb{Q}} Z,$$

where the projection is applied componentwise. Its residual Gram matrix is

$$G_{\mathcal{S}}(Z) := \mathbb{E}_{\mathbb{Q}}[Z_{\mathcal{S}}^{\perp} (Z_{\mathcal{S}}^{\perp})^{\top}].$$

Assumption B.2 (Nondegenerate residual Perron block). The matrix $G_{\mathcal{S}}(Z)$ is positive definite.

Definition B.3 (Vector residual Perron coefficient). Under Assumption B.2, define

$$\beta_{\mathcal{S}}^{(r)}(X; Z) := G_{\mathcal{S}}(Z)^{-1} \mathbb{E}_{\mathbb{Q}}[Z_{\mathcal{S}}^{\perp} (X - C_{\mathcal{S}}^{\mathbb{Q}} X)].$$

A nonzero vector exposure imposes a residual floor in the Perron block.

Theorem B.4 (Vector Perron residual lower bound). *Under Assumption B.2,*

$$\inf_{H \in \mathcal{S}} \|X - H\|_{2, \mathbb{Q}}^2 \geq (\beta_{\mathcal{S}}^{(r)}(X; Z))^{\top} G_{\mathcal{S}}(Z) \beta_{\mathcal{S}}^{(r)}(X; Z).$$

If

$$L_X^{\mathbb{Q}} = Y_X + b_X^{\top} Z + B_X, \quad Y_X \in \mathcal{S},$$

and

$$\mathbb{E}_{\mathbb{Q}}[Z_{\mathcal{S}}^{\perp}(B_X - C_{\mathcal{S}}^{\mathbb{Q}} B_X)] = 0,$$

then $\beta_{\mathcal{S}}^{(r)}(X; Z) = b_X$ and

$$\inf_{H \in \mathcal{S}} \|X - H\|_{2, \mathbb{Q}}^2 \geq b_X^{\top} G_{\mathcal{S}}(Z) b_X.$$

Proof. Let $R_X = X - C_{\mathcal{S}}^{\mathbb{Q}} X$. The projection of R_X onto the finite-dimensional span of $\{Z_1^{\perp}, \dots, Z_r^{\perp}\}$ has coefficient vector $\beta_{\mathcal{S}}^{(r)}(X; Z)$, because the normal equations are

$$G_{\mathcal{S}}(Z) \beta = \mathbb{E}_{\mathbb{Q}}[Z_{\mathcal{S}}^{\perp} R_X].$$

The squared norm of this finite-dimensional projection is

$$(\beta_{\mathcal{S}}^{(r)}(X; Z))^{\top} G_{\mathcal{S}}(Z) \beta_{\mathcal{S}}^{(r)}(X; Z).$$

Since the squared norm of R_X dominates the squared norm of any orthogonal projection, the first inequality follows.

For the pure vector loading statement, $Z_{\mathcal{S}}^{\perp}$ is $\mathcal{K}$ -measurable and orthogonal to $\mathcal{S}$, so

$$\mathbb{E}_{\mathbb{Q}}[Z_{\mathcal{S}}^{\perp} R_X] = \mathbb{E}_{\mathbb{Q}}[Z_{\mathcal{S}}^{\perp} (L_X^{\mathbb{Q}} - C_{\mathcal{S}}^{\mathbb{Q}} L_X^{\mathbb{Q}})].$$

The Y_X term contributes zero because $Y_X \in \mathcal{S}$. The B_X term contributes zero by assumption. For the Perron block, linearity of $C_{\mathcal{S}}^{\mathbb{Q}}$ gives $C_{\mathcal{S}}^{\mathbb{Q}}(b_X^{\top} Z) = b_X^{\top} C_{\mathcal{S}}^{\mathbb{Q}} Z$, so the residual is $b_X^{\top} Z_{\mathcal{S}}^{\perp}$ and its moment with $Z_{\mathcal{S}}^{\perp}$ is $G_{\mathcal{S}}(Z) b_X$. Hence $\beta_{\mathcal{S}}^{(r)}(X; Z) = b_X$. $\square$

Whitened vector budgets produce coordinate-free Perron intervals.

Corollary B.5 (Vector Perron ellipsoidal interval). *Assume $Z_{\mathcal{H}}^{\perp} := Z - C_{\mathcal{H}}^{\mathbb{Q}} Z$ is componentwise bounded and has positive definite Gram matrix*

$$G_{\mathcal{H}}(Z) := \mathbb{E}_{\mathbb{Q}}[Z_{\mathcal{H}}^{\perp} (Z_{\mathcal{H}}^{\perp})^{\top}].$$

Suppose

$$L_X^{\mathbb{Q}} = Y_X + b_X^{\top} Z + B_X, \quad Y_X \in L^2(\mathbb{Q}, \mathcal{H}),$$

and

$$\mathbb{E}_{\mathbb{Q}}[Z_{\mathcal{H}}^{\perp} (B_X - C_{\mathcal{H}}^{\mathbb{Q}} B_X)] = 0.$$

Let A and ρ satisfy the bounded-positivity condition in Definition 4.5 for $R = Z_{\mathcal{H}}^{\perp}$. Then the ellipsoidal residual set generated by $R = Z_{\mathcal{H}}^{\perp}$ gives

$$\Pi_{A, \rho}^+(X) - \Pi_{A, \rho}^-(X) = 2\rho (b_X^{\top} G_{\mathcal{H}}(Z) A^{-1} G_{\mathcal{H}}(Z) b_X)^{1/2}.$$

Proof. In Theorem 4.6, the residual moment vector has entries

$$\mathbb{E}_{\mathbb{Q}}[X Z_{i, \mathcal{H}}^{\perp}] = \mathbb{E}_{\mathbb{Q}}[L_X^{\mathbb{Q}} Z_{i, \mathcal{H}}^{\perp}],$$

because $Z_{i, \mathcal{H}}^{\perp}$ is $\mathcal{K}$ -measurable. The public term Y_X is orthogonal to $Z_{i, \mathcal{H}}^{\perp}$, the B_X term is removed by the stated orthogonality condition, and the block term contributes $(G_{\mathcal{H}}(Z) b_X)_i$. Thus the moment vector is $G_{\mathcal{H}}(Z) b_X$. Substitution into the interval width formula of Theorem 4.6 gives the result. $\square$

APPENDIX C. ALTERNATIVE PUBLIC-EQUIVALENT PRICING SETS

The main text uses local linear residual tilts as tangent objects and conditional exponential kernels as the positive Gaussian residual family. Density bands and entropy penalties give alternative public-equivalent pricing sets that preserve public option prices.

Definition C.1 (Density-band public-equivalent set). For constants $0 < a < 1 < b < \infty$, define

$$\mathcal{Z}_{\mathcal{H}}^{[a,b]} := \{Z \in L^\infty(\mathbb{Q}) : a \leq Z \leq b, \mathbb{E}_{\mathbb{Q}}[Z \mid \mathcal{H}] = 1\}.$$

The associated pricing measures are $d\mathbb{Q}^Z/d\mathbb{Q} = Z$.

Density bands convert residual pricing freedom into an order interval.

Proposition C.2 (Density-band residual interval). For $X \in L^1(\mathbb{Q})$, the set

$$\{\mathbb{E}_{\mathbb{Q}}[ZX] : Z \in \mathcal{Z}_{\mathcal{H}}^{[a,b]}\}$$

is a bounded interval whenever $\mathcal{Z}_{\mathcal{H}}^{[a,b]}$ is nonempty and closed under convex combinations. Every $Z \in \mathcal{Z}_{\mathcal{H}}^{[a,b]}$ preserves bounded public payoff prices.

Proof. Public payoff preservation follows from

$$\mathbb{E}_{\mathbb{Q}}[ZY] = \mathbb{E}_{\mathbb{Q}}[\mathbb{E}_{\mathbb{Q}}[Z \mid \mathcal{H}]Y] = \mathbb{E}_{\mathbb{Q}}[Y]$$

for bounded $\mathcal{H}$ -measurable Y . The map $Z \mapsto \mathbb{E}_{\mathbb{Q}}[ZX]$ is linear, and $\mathcal{Z}_{\mathcal{H}}^{[a,b]}$ is convex. Therefore the image is a convex subset of $\mathbb{R}$, hence an interval. It is bounded because

$$|\mathbb{E}_{\mathbb{Q}}[ZX]| \leq b \mathbb{E}_{\mathbb{Q}}[|X|].$$

□

Definition C.3 (Conditional upper-tail support). Let $U \in L^1(\mathbb{Q})$, let $p \in (0, 1)$, and let $\mathcal{H} \subseteq \mathcal{F}$. Define

$$T_p(U \mid \mathcal{H}) := \text{ess sup} \{\mathbb{E}_{\mathbb{Q}}[U\xi \mid \mathcal{H}] : 0 \leq \xi \leq 1, \mathbb{E}_{\mathbb{Q}}[\xi \mid \mathcal{H}] = p\}.$$

The exact density-band support reallocates residual mass across conditional tails.

Theorem C.4 (Exact density-band residual support). Let $0 < a < 1 < b < \infty$, and set

$$p := \frac{1-a}{b-a}.$$

Let $X \in L^1(\mathbb{Q})$, and write

$$U := X - C_{\mathcal{H}}^{\mathbb{Q}}X.$$

Then

$$\begin{aligned} \sup_{\substack{a \leq Z \leq b \\ \mathbb{E}_{\mathbb{Q}}[Z \mid \mathcal{H}] = 1}} \mathbb{E}_{\mathbb{Q}}[ZX] &= \mathbb{E}_{\mathbb{Q}}[X] + (b-a)\mathbb{E}_{\mathbb{Q}}[T_p(U \mid \mathcal{H})], \\ \inf_{\substack{a \leq Z \leq b \\ \mathbb{E}_{\mathbb{Q}}[Z \mid \mathcal{H}] = 1}} \mathbb{E}_{\mathbb{Q}}[ZX] &= \mathbb{E}_{\mathbb{Q}}[X] - (b-a)\mathbb{E}_{\mathbb{Q}}[T_p(-U \mid \mathcal{H})]. \end{aligned}$$

If the density set is restricted to $\mathcal{K}$ -measurable densities, replace U by

$$C_{\mathcal{K}}^{\mathbb{Q}}X - C_{\mathcal{H}}^{\mathbb{Q}}C_{\mathcal{K}}^{\mathbb{Q}}X$$

and take the essential supremum over $\mathcal{K}$ -measurable ξ .

Proof. The feasible set is nonempty because $\xi \equiv p$ corresponds to $Z \equiv 1$. For any feasible Z , define

$$\xi := \frac{Z-a}{b-a}.$$

Then $0 \leq \xi \leq 1$, and

$$\mathbb{E}_{\mathbb{Q}}[\xi \mid \mathcal{H}] = \frac{1-a}{b-a} = p.$$

Since $\mathbb{E}_{\mathbb{Q}}[U \mid \mathcal{H}] = 0$,

$$\mathbb{E}_{\mathbb{Q}}[ZU] = a\mathbb{E}_{\mathbb{Q}}[U] + (b-a)\mathbb{E}_{\mathbb{Q}}[\xi U] = (b-a)\mathbb{E}_{\mathbb{Q}}[\xi U].$$

The affine change of variables therefore converts the supremum over feasible Z into the supremum over feasible ξ . It is reversible, so no admissible density is lost. The feasible family is stable under $\mathcal{H}$ -measurable pasting: if ξ_1, ξ_2 are feasible and $A \in \mathcal{H}$, then

$$\xi = \mathbf{1}_A \xi_1 + \mathbf{1}_{A^c} \xi_2$$

is feasible. Hence the conditional values $\mathbb{E}_{\mathbb{Q}}[U\xi \mid \mathcal{H}]$ are upward directed, and the lattice argument for essential suprema gives a sequence of feasible ξ_n such that

$$\mathbb{E}_{\mathbb{Q}}[U\xi_n \mid \mathcal{H}] \uparrow T_p(U \mid \mathcal{H}) \quad \text{a.s.}$$

Since $|\mathbb{E}_{\mathbb{Q}}[U\xi_n \mid \mathcal{H}]| \leq \mathbb{E}_{\mathbb{Q}}[|U| \mid \mathcal{H}]$, dominated convergence gives

$$\sup_{\xi} \mathbb{E}_{\mathbb{Q}}[U\xi] = \mathbb{E}_{\mathbb{Q}}[T_p(U \mid \mathcal{H})].$$

This gives the first display. Replacing U by $-U$ gives the infimum. If densities are required to be $\mathcal{K}$ -measurable, then $\mathbb{E}_{\mathbb{Q}}[ZX] = \mathbb{E}_{\mathbb{Q}}[ZC_{\mathcal{K}}^{\mathbb{Q}}X]$. The same calculation applies to the residual $C_{\mathcal{K}}^{\mathbb{Q}}X - C_{\mathcal{H}}^{\mathbb{Q}}C_{\mathcal{K}}^{\mathbb{Q}}X$, with the optimizer ξ restricted to be $\mathcal{K}$ -measurable. $\square$

Definition C.5 (Entropy-penalized public-equivalent price). For bounded $X \in L^{\infty}(\mathbb{Q})$ and $\lambda > 0$, define

$$\Pi_{\lambda}^{\text{ent}}(X) := \sup_{\substack{Z > 0 \\ \mathbb{E}_{\mathbb{Q}}[Z \mid \mathcal{H}] = 1}} \{ \mathbb{E}_{\mathbb{Q}}[ZX] - \lambda \mathbb{E}_{\mathbb{Q}}[Z \log Z] \}.$$

Entropy penalization gives a monotone conditional selector on the public-equivalent fiber.

Theorem C.6 (Conditional entropic residual price). For $X \in L^{\infty}(\mathbb{Q})$ and $\lambda > 0$,

$$\Pi_{\lambda}^{\text{ent}}(X) = \mathbb{E}_{\mathbb{Q}}[\lambda \log \mathbb{E}_{\mathbb{Q}}[\exp(X/\lambda) \mid \mathcal{H}]].$$

The optimizer is

$$Z_{\lambda}^X = \frac{\exp(X/\lambda)}{\mathbb{E}_{\mathbb{Q}}[\exp(X/\lambda) \mid \mathcal{H}]}.$$

Equivalently, with $\Delta_{\mathcal{H}}X = X - C_{\mathcal{H}}^{\mathbb{Q}}X$,

$$Z_{\lambda}^X = \frac{\exp(\Delta_{\mathcal{H}}X/\lambda)}{\mathbb{E}_{\mathbb{Q}}[\exp(\Delta_{\mathcal{H}}X/\lambda) \mid \mathcal{H}]}.$$

Proof. The displayed Z_{λ}^X is positive and satisfies $\mathbb{E}_{\mathbb{Q}}[Z_{\lambda}^X \mid \mathcal{H}] = 1$. For any admissible Z , conditional Gibbs' inequality, with the value understood in the extended sense when $\mathbb{E}_{\mathbb{Q}}[Z \log Z] = +\infty$, gives

$$\mathbb{E}_{\mathbb{Q}}\left[Z \frac{X}{\lambda} - Z \log Z \mid \mathcal{H}\right] \leq \log \mathbb{E}_{\mathbb{Q}}[\exp(X/\lambda) \mid \mathcal{H}].$$

Multiplying by λ and taking expectations gives the upper bound. Equality holds when

$$Z = \frac{\exp(X/\lambda)}{\mathbb{E}_{\mathbb{Q}}[\exp(X/\lambda) \mid \mathcal{H}]},$$

which proves the value and optimizer. The final formula follows because

$$\exp(X/\lambda) = \exp(C_{\mathcal{H}}^{\mathbb{Q}}X/\lambda) \exp(\Delta_{\mathcal{H}}X/\lambda),$$

and the $\mathcal{H}$ -measurable factor cancels between numerator and conditional denominator. $\square$

The large-penalty limit recovers a variance residual premium.

Corollary C.7 (Large-penalty entropic expansion). *Let $X \in L^\infty(\mathbb{Q})$, set*

$$\Delta_{\mathcal{H}}X = X - C_{\mathcal{H}}^{\mathbb{Q}}X, \quad M_X = \|\Delta_{\mathcal{H}}X\|_\infty.$$

Then

$$\Pi_\lambda^{\text{ent}}(X) = \mathbb{E}_{\mathbb{Q}}[X] + \frac{1}{2\lambda} \mathbb{E}_{\mathbb{Q}}[\text{Var}_{\mathbb{Q}}(X \mid \mathcal{H})] + O\left(\frac{M_X^3}{\lambda^2}\right) \quad \lambda \rightarrow \infty.$$

In particular, the entropic residual premium is nonnegative and vanishes for every $\mathcal{H}$ -measurable X .

Proof. The value formula in Theorem C.6 gives

$$\Pi_\lambda^{\text{ent}}(X) = \mathbb{E}_{\mathbb{Q}}[C_{\mathcal{H}}^{\mathbb{Q}}X] + \mathbb{E}_{\mathbb{Q}}[\lambda \log \mathbb{E}_{\mathbb{Q}}[\exp(\Delta_{\mathcal{H}}X/\lambda) \mid \mathcal{H}]].$$

Since $\mathbb{E}_{\mathbb{Q}}[\Delta_{\mathcal{H}}X \mid \mathcal{H}] = 0$ and $|\Delta_{\mathcal{H}}X| \leq M_X$, Taylor's formula with a uniform remainder gives

$$\mathbb{E}_{\mathbb{Q}}[\exp(\Delta_{\mathcal{H}}X/\lambda) \mid \mathcal{H}] = 1 + \frac{1}{2\lambda^2} \mathbb{E}_{\mathbb{Q}}[(\Delta_{\mathcal{H}}X)^2 \mid \mathcal{H}] + O\left(\frac{M_X^3}{\lambda^3}\right),$$

uniformly for large λ . The inside of the logarithm is then bounded away from zero, and $\log(1+u) = u + O(u^2)$ gives, after multiplying by λ ,

$$\lambda \log \mathbb{E}_{\mathbb{Q}}[\exp(\Delta_{\mathcal{H}}X/\lambda) \mid \mathcal{H}] = \frac{1}{2\lambda} \mathbb{E}_{\mathbb{Q}}[(\Delta_{\mathcal{H}}X)^2 \mid \mathcal{H}] + O\left(\frac{M_X^3}{\lambda^2}\right),$$

again uniformly for large λ . Taking expectations proves the expansion because $\mathbb{E}_{\mathbb{Q}}[(\Delta_{\mathcal{H}}X)^2 \mid \mathcal{H}] = \text{Var}_{\mathbb{Q}}(X \mid \mathcal{H})$. Nonnegativity follows from conditional Jensen applied to the logarithmic moment generating function; if X is $\mathcal{H}$ -measurable, then $\Delta_{\mathcal{H}}X = 0$ and the premium term is zero. $\square$

APPENDIX D. SPECTRAL CAPITAL FUNCTIONALS

Expected shortfall is one admissible physical capital input. Spectral risk measures average expected shortfall levels and leave the projection results unchanged.

Definition D.1 (Admissible spectrum). A function $\varphi : (0, 1) \rightarrow [0, \infty)$ is an admissible L^2 -spectrum if

$$\int_0^1 \varphi(\alpha) d\alpha = 1, \quad \int_0^1 \frac{\varphi(\alpha)}{\sqrt{1-\alpha}} d\alpha < \infty.$$

For $Y \in L^2(\mathbb{P})$, define

$$\rho_\varphi^{\mathbb{P}}(Y) := \int_0^1 \text{ES}_\alpha^{\mathbb{P}}(Y) \varphi(\alpha) d\alpha.$$

The spectral extension needs only an L^2 bound for expected shortfall.

Lemma D.2 (Expected shortfall L^2 bound). *For $Y \in L^2(\mathbb{P})$,*

$$\text{ES}_\alpha^{\mathbb{P}}(Y) \leq \mathbb{E}_{\mathbb{P}}[Y] + \frac{\|Y - \mathbb{E}_{\mathbb{P}}[Y]\|_{2,\mathbb{P}}}{\sqrt{1-\alpha}}, \quad \alpha \in (0, 1).$$

If $\mathbb{E}_{\mathbb{P}}[Y] = 0$, then

$$0 \leq \text{ES}_\alpha^{\mathbb{P}}(Y) \leq \frac{\|Y\|_{2,\mathbb{P}}}{\sqrt{1-\alpha}}.$$

Proof. By cash additivity, it is enough to treat centered Y . For centered Y , take $r = 0$ in Definition 7.10 to get finiteness from above. For the stated bound, use, for $r > 0$,

$$(Y - r)_+ = (Y_+ - r)_+ \leq \frac{Y_+^2}{4r}.$$

This gives

$$\text{ES}_\alpha^{\mathbb{P}}(Y) \leq r + \frac{\mathbb{E}_{\mathbb{P}}[Y_+^2]}{4r(1-\alpha)}.$$

Optimizing over $r > 0$ gives

$$\text{ES}_\alpha^\mathbb{P}(Y) \leq \frac{\|Y_+\|_{2,\mathbb{P}}}{\sqrt{1-\alpha}} \leq \frac{\|Y\|_{2,\mathbb{P}}}{\sqrt{1-\alpha}}.$$

The lower bound for centered Y follows from Proposition 7.11. Restoring the mean gives the displayed inequality for general Y . $\square$

Admissible spectra preserve the capital properties required by the projection bounds.

Proposition D.3 (Spectral capital is admissible). *If φ is an admissible L^2 -spectrum, then $\rho_\varphi^\mathbb{P}$ is finite on $L^2(\mathbb{P})$, convex, cash additive, and satisfies*

$$\rho_\varphi^\mathbb{P}(Y) \geq 0 \quad \text{whenever} \quad \mathbb{E}_\mathbb{P}[Y] = 0.$$

Consequently $\rho_\varphi^\mathbb{P}$ can be used as $\rho_\mathbb{P}$ in Definition 7.12, Proposition 7.13, and Theorem 7.14.

Proof. Finiteness follows from Lemma D.2 and the defining integrability condition on φ . Convexity and cash additivity follow by integrating the corresponding properties of $\text{ES}_\alpha^\mathbb{P}$ from Proposition 7.11. The nonnegativity on centered losses follows in the same way from $\text{ES}_\alpha^\mathbb{P}(Y) \geq 0$ when $\mathbb{E}_\mathbb{P}[Y] = 0$. $\square$

APPENDIX E. BID-ASK AND LIQUIDITY-CONSTRAINED CALIBRATION

Observed option data contain bid-ask spreads, quote-density variation, and liquidity differences across maturity and moneyness. These constraints enter the public calibration system as finite feasibility restrictions.

Definition E.1 (Bid-ask public calibration system). Let $Y_1, \dots, Y_m \in L^\infty(\mathbb{Q}, \mathcal{H})$ be public payoff coordinates, and let

$$\underline{p}_i \leq \bar{p}_i, \quad i = 1, \dots, m,$$

be discounted bid and ask price bounds. A density Z is bid-ask feasible if

$$Z > 0, \quad \mathbb{E}_\mathbb{Q}[Z] = 1, \quad \underline{p}_i \leq \mathbb{E}_\mathbb{Q}[ZY_i] \leq \bar{p}_i, \quad i = 1, \dots, m.$$

Bid-ask constraints turn finite residual calibration into a compact convex feasibility problem.

Proposition E.2 (Finite residual calibration with bid-ask bounds). *Let $R_1, \dots, R_d \in L^\infty(\mathbb{Q})$, and define*

$$Z_\vartheta = 1 + \vartheta^\top R.$$

Assume a compact convex set $K \subset \mathbb{R}^d$ satisfies $Z_\vartheta > 0$ for every $\vartheta \in K$. The feasible parameter set

$$K_{\text{ba}} := \left\{ \vartheta \in K : \mathbb{E}_\mathbb{Q}[Z_\vartheta] = 1, \underline{p}_i \leq \mathbb{E}_\mathbb{Q}[Z_\vartheta Y_i] \leq \bar{p}_i, \quad i = 1, \dots, m \right\}$$

is compact and convex. For every $X \in L^1(\mathbb{Q})$, the bid-ask residual price bounds

$$\inf_{\vartheta \in K_{\text{ba}}} \mathbb{E}_\mathbb{Q}[Z_\vartheta X], \quad \sup_{\vartheta \in K_{\text{ba}}} \mathbb{E}_\mathbb{Q}[Z_\vartheta X]$$

are attained whenever $K_{\text{ba}} \neq \emptyset$.

Proof. The normalization condition is affine because

$$\mathbb{E}_\mathbb{Q}[Z_\vartheta] = 1 + \vartheta^\top \mathbb{E}_\mathbb{Q}[R].$$

For each i ,

$$\mathbb{E}_\mathbb{Q}[Z_\vartheta Y_i] = \mathbb{E}_\mathbb{Q}[Y_i] + \vartheta^\top \mathbb{E}_\mathbb{Q}[RY_i],$$

so each bid-ask constraint is a closed affine half-space condition in ϑ . Intersecting the normalization hyperplane and finitely many bid-ask half-spaces with compact convex K gives compact convex K_{ba} . The objective

$$\vartheta \mapsto \mathbb{E}_\mathbb{Q}[Z_\vartheta X] = \mathbb{E}_\mathbb{Q}[X] + \vartheta^\top \mathbb{E}_\mathbb{Q}[RX]$$

is continuous and affine. A continuous function attains its infimum and supremum on a nonempty compact set. $\square$

Definition E.3 (Liquidity-weighted calibration objective). For observed public coordinates $q \in \mathbb{R}^m$, model coordinates $m(\theta)$, and positive liquidity weights $w_i > 0$, define

$$J_{\text{liq}}(\theta) := \sum_{i=1}^m w_i (q_i - m_i(\theta))^2.$$

If bid-ask bounds are present, the constrained calibration set is

$$\Theta_{\text{ba}} := \{\theta \in \Theta : \underline{p}_i \leq m_i(\theta) \leq \bar{p}_i, i = 1, \dots, m\}.$$

Liquidity-weighted calibration has a minimizer once the feasible bid-ask set is compact.

Proposition E.4 (Existence of liquidity-constrained public calibration). *Let Θ be compact, let $m : \Theta \rightarrow \mathbb{R}^m$ be continuous, and let Θ_{ba} be nonempty. Then J_{liq} attains a minimum on Θ_{ba} .*

Proof. The set Θ_{ba} is closed in Θ , because it is defined by finitely many closed inequalities involving continuous functions. Hence it is compact. The objective J_{liq} is continuous, so Weierstrass' theorem gives a minimizer. $\square$

REFERENCES

- [1] J. Vidal Llauradó (2026). *Risk-Neutral Compression and the Geometry of Latent Market Stress*. SSRN. <https://doi.org/10.2139/ssrn.6680938>.
- [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. Vidal Llauradó (2026). *A Rough Theory of Markets: Latent Contagion, Market Phases, and Market-Variance Risk*. SSRN. <https://doi.org/10.2139/ssrn.6616121>.
- [6] 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.
- [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] S. A. Ross, *Options and efficiency*, Quarterly Journal of Economics **90** (1976), no. 1, 75–89, doi:10.2307/1886087.
- [9] G. Bakshi and D. Madan, *Spanning and derivative-security valuation*, Journal of Financial Economics **55** (2000), no. 2, 205–238, doi:10.1016/S0304-405X(99)00050-1.
- [10] B. Dumas, J. Fleming, and R. E. Whaley, *Implied volatility functions: Empirical tests*, Journal of Finance **53** (1998), no. 6, 2059–2106, doi:10.1111/0022-1082.00083.
- [11] G. Bakshi, C. Cao, and Z. Chen, *Empirical performance of alternative option pricing models*, Journal of Finance **52** (1997), no. 5, 2003–2049, doi:10.1111/j.1540-6261.1997.tb02749.x.
- [12] J. Pan, *The jump-risk premia implicit in options: Evidence from an integrated time-series study*, Journal of Financial Economics **63** (2002), no. 1, 3–50, doi:10.1016/S0304-405X(01)00088-5.
- [13] A. Buraschi and J. Jackwerth, *The price of a smile: Hedging and spanning in option markets*, Review of Financial Studies **14** (2001), no. 2, 495–527, doi:10.1093/rfs/14.2.495.
- [14] J. D. Coval and T. Shumway, *Expected option returns*, Journal of Finance **56** (2001), no. 3, 983–1009, doi:10.1111/0022-1082.00352.
- [15] G. Bakshi and N. Kapadia, *Delta-hedged gains and the negative market volatility risk premium*, Review of Financial Studies **16** (2003), no. 2, 527–566, doi:10.1093/rfs/hhg002.
- [16] C. S. Jones, *A nonlinear factor analysis of S&P 500 index option returns*, Journal of Finance **61** (2006), no. 5, 2325–2363, doi:10.1111/j.1540-6261.2006.01059.x.
- [17] P. Carr and L. Wu, *Variance risk premiums*, Review of Financial Studies **22** (2009), no. 3, 1311–1341, doi:10.1093/rfs/hhn038.
- [18] T. Bollerslev, G. Tauchen, and H. Zhou, *Expected stock returns and variance risk premia*, Review of Financial Studies **22** (2009), no. 11, 4463–4492, doi:10.1093/rfs/hhp008.
- [19] T. G. Andersen, N. Fusari, and V. Todorov, *The risk premia embedded in index options*, Journal of Financial Economics **117** (2015), no. 3, 558–584, doi:10.1016/j.jfineco.2015.06.005.
- [20] R. Cont and J. da Fonseca, *Dynamics of implied volatility surfaces*, Quantitative Finance **2** (2002), no. 1, 45–60, doi:10.1088/1469-7688/2/1/304.
- [21] P. Christoffersen, S. Heston, and K. Jacobs, *The shape and term structure of the index option smirk: Why multifactor stochastic volatility models work so well*, Management Science **55** (2009), no. 12, 1914–1932, doi:10.1287/mnsc.1090.1065.

- [22] P. Collin-Dufresne and R. S. Goldstein, *Do bonds span the fixed income markets? Theory and evidence for unspanned stochastic volatility*, Journal of Finance **57** (2002), no. 4, 1685–1730, doi:10.1111/1540-6261.00475.
- [23] A. B. Trolle and E. S. Schwartz, *Unspanned stochastic volatility and the pricing of commodity derivatives*, Review of Financial Studies **22** (2009), no. 11, 4423–4461, doi:10.1093/rfs/hhp036.
- [24] J. H. Cochrane and J. Saá-Requejo, *Beyond arbitrage: Good-deal asset price bounds in incomplete markets*, Journal of Political Economy **108** (2000), no. 1, 79–119, doi:10.1086/262112.
- [25] 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.
- [26] J. Gatheral, T. Jaisson, and M. Rosenbaum, *Volatility is rough*, Quantitative Finance **18** (2018), no. 6, 933–949, doi:10.1080/14697688.2017.1393551.
- [27] 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.
- [28] 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.
- [29] 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.
- [30] 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.
- [31] 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.
- [32] 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.
- [33] M. Schweizer, *Mean-variance hedging for general claims*, The Annals of Applied Probability **2** (1992), no. 1, 171–179, doi:10.1214/aoap/1177005776.
- [34] 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.
- [35] R. T. Rockafellar and S. Uryasev, *Optimization of conditional value-at-risk*, Journal of Risk **2** (2000), no. 3, 21–41, doi:10.21314/JOR.2000.038.
- [36] P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath, *Coherent measures of risk*, Mathematical Finance **9** (1999), no. 3, 203–228, doi:10.1111/1467-9965.00068.
- [37] H. Föllmer and A. Schied, *Stochastic Finance: An Introduction in Discrete Time*, 4th ed., De Gruyter, 2016, doi:10.1515/9783110463453.
- [38] D. Williams, *Probability with Martingales*, Cambridge Mathematical Textbooks, Cambridge University Press, 1991, doi:10.1017/CBO9780511813658.

Email address: `vidal.llaurado@gmail.com`